

Federated Quantum Machine Learning with Privacy-Preserving Measurement Statistics for Edge Analytics

Muhammad Iqbal^{1,*}, Syahril Efendi²

¹Faculty of Science and Technology, Universitas Pembangunan Panca Budi, Indonesia.

²Department of Computer Science, Universitas Sumatera Utara, Indonesia

ABSTRACT

Edge analytics requires continuous model adaptation under tight latency, bandwidth, and confidentiality constraints, yet conventional federated learning can leak sensitive information through shared updates and often struggles under non-IID client data. This study introduces Federated Quantum Machine Learning that transmits privacy-preserving measurement statistics derived from hybrid parameterized quantum circuits rather than raw samples or high-dimensional gradients. Across 50 edge clients over 20 rounds, the proposed method achieves 0.844 accuracy and 0.836 macro-F1 with ECE 0.048, outperforming a classical FedAvg baseline (0.801 accuracy, 0.788 macro-F1, ECE 0.062) while maintaining predictable edge overhead (845 ms average round latency) and compact messaging (3.1 KB average payload per round). Privacy-measurement sweeps demonstrate accuracy saturation with increasing privacy budget, improving from 0.816 at $\epsilon = 0.5$ to 0.856 at $\epsilon = 8.0$ under 512 shots, while calibration improves more sharply, with ECE decreasing from 0.068 to 0.046 over the same range. Increasing shot budget from 128 to 2048 consistently increases accuracy under all privacy settings, confirming that measurement variance is a controllable bottleneck for protected learning. Communication analysis shows that statistics-only messaging remains an order of magnitude smaller than gradient-style baselines (33.9 KB average payload per round), and secure aggregation adds bounded overhead (4.8 KB average payload) with a modest utility shift (0.836 accuracy, ECE 0.051). Robustness experiments under severe non-IID drift (severity 0.9) preserve macro-F1 at 0.801 with expectation-based statistics versus 0.772 for histogram-based statistics, and perturbed-client tests show improved stability with robust aggregation, sustaining 0.836 accuracy at 10% perturbed participation compared to 0.824 under naïve aggregation. These results indicate that low-dimensional, privatized measurement summaries enable practical, reliable, and privacy-aligned federated quantum learning on constrained edge infrastructures.

Keywords Federated Quantum Machine Learning, Privacy-Preserving Measurement Statistics, Differential Privacy, Secure Aggregation, Edge Analytics, Hybrid Quantum-Classical Models, Calibration Error

INTRODUCTION

Edge analytics has become a core operational requirement for Internet of Things and mobile cyber-physical systems because inference and adaptation must occur under tight latency, bandwidth, and energy constraints. Centralized learning pipelines remain misaligned with these constraints because raw-data uplink increases exposure and network load, while cloud-centric retraining delays model refresh cycles. Edge computing therefore motivates distributed learning designs that preserve locality and prioritize low-overhead communications for model improvement under heterogeneous device capabilities [1].

Submitted 10 January 2026

Accepted 25 February 2026

Published 31 August 2026

Corresponding author

Muhammad Iqbal,
muhammadiqbal@dosen.panca
budi.ac.id

Additional Information and
Declarations can be found on
[page 194](#)

DOI: [10.63913/jqai.v1i2.12](#)

© Copyright

2026 Iqbal and Efendi

Distributed under
Creative Commons CC-BY 4.0

Federated learning partially addresses these constraints by keeping training data on-device and coordinating updates through an aggregator, which is attractive for mobile edge networks where connectivity and compute vary across clients. However, operational deployments face chronic bottlenecks in communication rounds, unreliable participation, and resource-aware scheduling, which amplify non-stationarity when client data are non-IID. These constraints are acute at the edge because client dropouts, intermittent links, and device throttling alter the effective training distribution across rounds and degrade convergence stability [2].

In Internet of Things settings, the privacy motivation extends beyond data locality because shared gradients, logits, or intermediate representations can still expose information through inference attacks under realistic threat models. Surveys of federated learning in IoT highlight that security and privacy must be treated as first-class design objectives rather than afterthoughts, because adversaries can exploit update traffic, client metadata, and repeated participation to infer sensitive attributes. Consequently, privacy guarantees must be coupled to system-level constraints such as uplink budgets and protocol overheads in edge deployments [3].

Despite its promise, federated learning does not inherently guarantee privacy, and the literature documents recurring risks from update leakage, poisoning, and inference across clients and rounds. A consolidated view of open problems emphasizes that practical deployments require defensible threat models, robust aggregation, and formal privacy controls that remain meaningful under client heterogeneity and stragglers. This motivates a methodological shift from treating privacy as a policy statement to implementing it as an explicit constraint in the training objective and the communication protocol [4].

The privacy gap is most visible in edge environments where adversaries can observe partial transcripts, correlate round-level participation, and exploit protocol artifacts to improve reconstruction fidelity. Recent security and privacy surveys synthesize how federated deployments remain vulnerable even when raw data never leave the device, particularly when the transmitted objects retain high-dimensional structure. This evidence supports the need for training pipelines that reduce message expressivity while maintaining utility, and for defenses that remain effective under non-IID sampling and constrained client participation [5].

Two families of defenses dominate current practice, namely secure aggregation and differential privacy, yet both introduce trade-offs that are amplified on edge devices. Secure aggregation protects individual updates from the server but adds cryptographic overhead and sensitivity to client dropout patterns [6]. Differential privacy adds calibrated noise, but privacy accounting and utility degradation depend on the transmitted signal's dimensionality and sensitivity under composition [7], [8]. These constraints motivate message designs that are natively compact and statistically stable.

Quantum machine learning offers an additional axis of capability for edge analytics by enabling expressive hypothesis classes via parameterized quantum circuits and quantum feature maps, while remaining compatible with hybrid optimization loops [9]. However, near-term constraints such as trainability limitations and hardware noise complicate direct deployment, and

barren plateaus can reduce optimization signal in larger circuits [10]. Quantum kernels and variational classifiers provide practical routes for edge-relevant decision tasks [11], but privacy-preserving federation in this regime remains underdeveloped [12], [13]. The novelty of this paper is a federated quantum learning protocol that shares privacy-preserving measurement statistics rather than gradients or raw features, integrating DP accounting and aggregation hardening into a statistics-first channel that is suitable for edge budgets and NISQ constraints, while aligning with emerging foundations of quantum federated learning [14].

Literature Review

Recent literature frames federated learning as a multi-layered design space in which statistical heterogeneity, systems heterogeneity, and governance constraints interact, making “privacy by decentralization” an incomplete assumption in practice. Taxonomies increasingly separate concerns into client sampling, local optimization, aggregation, communication, and trust models, because failures in any layer can dominate convergence and leakage behavior. This line of work motivates treating edge analytics as a constrained optimization problem rather than a deployment detail [15].

In edge settings, the dominant constraint is not only bandwidth but also device variability in compute, memory, power, and intermittency, which changes the effective participation distribution over rounds. Surveys focused on edge federated learning emphasize that non-stationary availability and heterogeneous accelerators shift design toward lightweight updates, adaptive scheduling, and robustness to dropouts, while preserving model quality under non-IID partitions. These observations justify architectures that explicitly model device classes and network strata when evaluating privacy-preserving protocols [16].

Security studies clarify that federated pipelines expose a wide attack surface beyond raw data exfiltration, including gradient inversion, membership inference, model poisoning, backdoors, and free-rider behaviors. A recurring conclusion is that privacy controls and robustness controls are coupled: stronger privacy noise can destabilize training, while stronger robustness filters can leak membership signals through selective acceptance. Consequently, modern threat models formalize both the honest-but-curious server and malicious clients, and they recommend layered defenses instead of single-mechanism guarantees [17].

Model combination remains a central determinant of generalization under heterogeneity. Systematic reviews of aggregation techniques report that naïve averaging can overfit dominant clients, amplify biased local minima, and become brittle under adversarial updates, motivating robust aggregation, adaptive weights, and optimization-aware combination rules. Parallel work on communication-efficient federated learning shows that sparsification, quantization, and periodic synchronization can preserve accuracy with lower uplink cost, but these transformations alter the statistical footprint of updates and therefore interact with privacy accounting [18], [19].

For formal privacy, differential privacy is the most common analytic framework, typically implemented as client-level perturbation via gradient clipping and noise addition, or through noise applied after aggregation. Empirical studies

demonstrate that privacy budgets can be maintained while limiting utility loss when clipping is calibrated to realistic client gradients, although performance degradation becomes pronounced under severe non-IID and small client participation. These findings motivate privacy mechanisms that target what is actually exposed, such as measurement summaries, rather than raw update tensors [20].

Secure protocols complement differential privacy by reducing what any coordinator can observe. Secret-sharing-based secure aggregation reduces server visibility into individual updates while still enabling global learning, with recent schemes designed to reduce communication overhead and tolerate dropouts without exposing per-client contributions. In the quantum domain, emerging surveys on quantum federated learning highlight that training with variational circuits introduces novel leakage channels through measurement statistics and shot noise, while also enabling hybrid workflows where quantum components are embedded into classical FL pipelines. Application-driven studies propose clustered and quantum-enhanced federated strategies in sensitive domains, further reinforcing the need for privacy-preserving measurement handling as a first-class design goal [21], [22], [23], [24].

Methodology

System Model, Edge Data Assumptions, and Threat Model

Federated training is defined over a population of K edge clients, each holding a private dataset $\mathcal{D}_k = (x_i, y_i)_{i=1}^{n_k}$ and a local quantum-capable runtime that can execute a bounded-depth parameterized quantum circuit. A central coordinator orchestrates synchronous communication rounds but does not access raw samples. The deployment target assumes intermittent connectivity and heterogeneous compute, so protocol steps are expressed in terms of measurable messages and bounded local epochs.

Edge data are treated as non-identically distributed across clients to reflect realistic sensing and interaction patterns. Let the joint distribution on client k be $P_k(X, Y)$, with divergence across clients captured by an empirical discrepancy term, such as $d(P_k, P_{k'})$. This assumption is operationally important because measurement statistics, rather than raw gradients, are shared under privacy constraints. The methodology therefore separates accuracy loss due to non-IID drift from accuracy loss due to privatization noise.

The threat model assumes an honest-but-curious coordinator that follows the protocol but attempts inference from received updates, including reconstruction of sensitive features from measurement histograms or sufficient statistics. Adversarial edge clients are also considered, primarily for robustness analysis of aggregation under noisy statistics. The adversary's observational access includes all transmitted messages, local model states are private, and hardware noise is treated as a separate confounder rather than a malicious channel. A global learning objective is posed as the minimization of an aggregate empirical risk under client weighting,

$$\min_{\theta} \mathcal{F}(\theta) = \sum_{k=1}^K \frac{n_k}{n} \mathcal{L}_k(\theta), \quad n = \sum_{k=1}^K n_k \quad (1)$$

where $L_k(\theta)$ is the local loss computed from privatized measurement statistics. Equation (1) establishes the target that the coordinator approximates through iterative aggregation while preserving privacy at the message level, which is the defining constraint of the proposed approach.

Figure 1 depicts the federated topology that constrains all learning signals to flow through protected messages rather than raw samples. The left region emphasizes non-IID client shards and the local hybrid PQC pipeline, which is the only place where private data are accessed. The right region models an honest-but-curious coordinator that receives only privatized statistics, making the trust boundary explicit. The bidirectional arrows separate parameter broadcast from statistics upload, clarifying how privacy is enforced at the communication interface.




Figure 1 Federated QML System Topology and Trust Boundaries

Table 1 consolidates the core notation and explicitly couples each term to its privacy relevance, which is critical when privacy is enforced on measurement-derived signals rather than gradients. The mapping clarifies that $\mathcal{D}_k$ remains strictly local, while s_k becomes the principal object of protection and aggregation. The table also distinguishes coordinator and client adversaries, aligning threat assumptions with mitigation mechanisms such as differential privacy and secure aggregation, and providing interpretability for later ablations.

Table 1 Notation and Threat Assumptions

Symbol/Term	Meaning	Scope	Assumption	Privacy Relevance
K	Number of edge clients	System-level	Heterogeneous devices	Affects sampling and exposure
$\mathcal{D}_k, n_k$	Client dataset and size	Client-level	Non-IID partitions	Inference risk from small n_k
$\theta(t)$	Global parameters at round t	Protocol	Synchronous rounds	Broadcast is non-sensitive
s_k	Measurement statistics vector	Client-level	Shot-based estimation	Primary protected message
Honest-but-curious	Protocol-following coordinator	Threat model	Observes all messages	Motivates DP and SecAgg
Adversarial client	Malicious participant	Threat model	May send distorted updates	Motivates robustness checks

Federated Quantum Model Architecture

Each client instantiates a hybrid model composed of a classical front-end for

feature normalization and a quantum core implemented as a variational quantum circuit $U(x; \theta)$. Input x is mapped to a quantum state via an encoding map $E(x)$, after which trainable layers apply entangling operations with hardware-aware depth constraints. The design explicitly limits two-qubit gate count to match edge device fidelity and reduce measurement overhead, which directly affects the budget for privatization.

A measurement layer produces a vector of statistics $s \in \mathbb{R}^d$, which can be defined either as expectation values of observables or as derived features from bitstring histograms. For classification, the quantum output is passed to a lightweight classical head $g(\cdot)$, enabling stable optimization under noisy statistics. This separation makes it possible to privatize statistics without requiring direct sharing of circuit parameters or raw bitstrings, depending on the selected privacy mechanism.

The per-client loss is computed from the predicted label $\hat{y}$ and the ground truth y using a standard supervised objective such as cross-entropy. The approach is compatible with minibatch evaluation, but the measurement cost motivates micro-batching and caching of repeated circuits when possible. This is treated as an implementation detail that influences latency, not a change in the learning objective, and is therefore reported in the evaluation budget rather than altering the model definition. The forward model is formalized by:

$$\langle O_j \rangle_\theta(x) = \text{Tr}(O_j U(x; \theta) \rho_0 U(x; \theta)^\dagger), \hat{y} = g([\langle O_1 \rangle_\theta, \dots, \langle O_d \rangle_\theta]) \quad (2)$$

Equation (2) anchors the measurable interface between quantum execution and federated learning, because $\langle O_j \rangle_\theta$ is estimated from finite shots and becomes the object that is subsequently privatized and aggregated.

Figure 2 operationalizes the hybrid architecture as a chain of measurable transformations from edge input to decision output. The diagram isolates the encoding map and PQC ansatz as quantum-specific components, while keeping the classical head lightweight to stabilize learning under noisy statistics. The explicit “Shots S ” block highlights the measurement bottleneck that shapes both latency and privacy calibration. By locating privatization after measurement, the pipeline clarifies that privacy protection targets the most informative shared artifact, namely measurement statistics.




Figure 2 Per-Client Hybrid Pipeline for Federated Quantum Learning

Table 2 connects circuit hyperparameters to concrete edge constraints and the induced privacy burden, which is essential when statistical noise originates from both finite shots and explicit privatization. The table makes the trade-off measurable by tying shot budget SSS to estimator variance and therefore to the scale of DP noise needed for a fixed privacy target. It also clarifies how larger observable sets increase feature richness while expanding the dimensionality of privatized messages, which affects communication, composition, and coordinator-side aggregation stability.

Table 2 Circuit Hyperparameters and Device Constraints

Parameter	Symbol	Typical Range	Edge Constraint	Effect on Utility	Effect on Privacy Cost
Qubits	Q	4 to 10	Device availability	Representation capacity	Higher-dimensional stats
Ansatz depth	L	2 to 8	Two-qubit fidelity	Expressivity	More shots for stable estimates
Shots per batch	S	128 to 2048	Latency budget	Lower estimator variance	Enables smaller DP noise
Observable count	d	4 to 32	Measurement overhead	Richer features	More dimensions to privatize
Head complexity	H	Linear to small MLP	CPU budget	Decision stability	Indirect via confidence calibration

Privacy-Preserving Measurement Statistics

Privacy preservation is implemented at the level of measurement-derived statistics to avoid transmitting raw bitstrings or sample-level gradients. For each client, the primary message is a statistic vector s_k obtained from circuit executions over local minibatches. Two families of statistics are considered: histogram-based features, where s_k corresponds to normalized counts over bitstrings, and expectation-based features, where s_k contains estimated $\langle O_i \rangle$ values. The latter is typically lower dimensional and therefore cheaper to privatize.

A local differential privacy mechanism perturbs statistics before transmission, yielding $\tilde{s}_k$. Calibration uses a sensitivity bound Δ associated with the statistic definition, with Gaussian noise preferred for composition control in multi-round training. For histogram statistics, sensitivity is bounded by the contribution of a single sample to a normalized bin count, while for expectation statistics it is bounded by the maximum per-sample influence on the estimator. This step ensures that privacy is enforced even if the coordinator is inquisitive.

To mitigate coordinator-side inference further, an optional secure aggregation layer can be applied so that only aggregated sums of privatized statistics are visible. This protects against leakage from individual client messages, leaving only population-level signals after aggregation. The methodology treats secure aggregation as orthogonal to differential privacy: privacy is guaranteed by noise, while secure aggregation reduces exposure to per-client patterns. Their combination is therefore evaluated as a joint ablation in the experimental design. The privatization step is formalized as:

$$\tilde{s}_k = s_k + \eta, \eta \sim \mathcal{N}(0, \sigma^2 I)$$
 (3)

with σ selected to meet an (ϵ, δ) -DP budget under the chosen accountant across T rounds. Equation (3) is the key interface between quantum measurement and privacy, because shot noise already introduces stochasticity, and the methodology explicitly separates intrinsic sampling variance from additive privacy noise in the reported utility trade-offs.

Figure 3 visualizes the end-to-end protection of measurement statistics by separating three transformations that often become conflated in practice. The curve labeled clipped stats shows how sensitivity bounding reduces extreme values that would otherwise dominate privacy calibration, while the privatized stats curve shows the utility impact of additive noise. The secure-aggregated mean emphasizes how population-level averaging mitigates per-client volatility without undoing privacy guarantees. This figure thereby motivates reporting both shot budget and DP parameters as co-determinants of downstream

accuracy.




Figure 3 Measurement Statistics Protection Via Clipping, DP Noise, and Aggregation

Table 3 makes the privacy layer auditable by linking each mechanism to a protected object and a concrete parameter that appears in experimental settings. This is important for measurement-statistics privacy because leakage can occur through value magnitudes, per-client message uniqueness, or participation traces. By explicitly including clipping and sampling, the table separates formal DP guarantees from operational protections that reduce attack surface. The structure also supports controlled ablations by allowing the privacy stack to be incrementally enabled and evaluated for both utility and exposure reduction.

Table 3 Privacy Mechanisms, Parameters, and Guarantees

Mechanism	Protected Object	Key Parameter	Reported Privacy	Primary Risk Mitigated
Clipping	Statistics vector	C	Sensitivity bound	Outlier-driven leakage
Gaussian DP	Expectation features	σ	(ϵ, δ) over T rounds	Coordinator inference
Histogram DP	Bitstring frequencies	σ or b	(ϵ, δ) over T rounds	Membership and attribute inference
Secure aggregation	Client messages	Masking protocol	Server sees only sums	Per-client pattern exposure
Client sampling	Participation traces	IStI/K	Reduced exposure per round	Repeated contribution linkage

Federated Training and Aggregation Protocol

Training proceeds in communication rounds $t = 1, \dots, T$. In each round, selected clients download the current global parameters $\theta^{(t)}$, execute local optimization for E epochs, and compute measurement-derived statistics required to estimate gradients or parameter updates. Because statistics are privatized before transmission, local optimization favors estimators that remain stable under additive noise, such as aggregated minibatch statistics rather than per-sample contributions. This design choice is critical for edge settings with strict shot budgets and limited compute.

A weighted aggregation rule combines privatized client updates into a new global model. When the message is a parameter delta $\Delta\theta_k$, aggregation follows a FedAvg-style update; when the message is an aggregated statistic sufficient to reconstruct an update, the coordinator first combines statistics and then applies a deterministic update mapping. The second option reduces exposure of parameters but can increase coordinator compute. Both variants preserve the same objective but differ in communication structure and are therefore compared empirically.

Communication and privacy budgets are co-managed through client sampling and update compression. Client sampling reduces per-round exposure and can improve scalability, but it may degrade convergence under severe non-IID drift. The protocol therefore includes a minimum participation constraint and reports the effective sampling rate. For compression, low-rank or quantized deltas are permitted, but compression is applied after privatization to avoid altering the privacy accounting assumptions. The global objective and update are expressed as:

$$\theta^{(t+1)} = \theta^{(t)} - \eta \sum_{k \in S_t} \frac{n_k}{\sum_{j \in S_t} n_j} \Delta \theta_k^{(t)} \quad (4)$$

where S_t is the sampled client set and $\Delta \theta_k^{(t)}$ is computed from privatized measurement statistics. Equation (4) clarifies that privacy operates on the update source, not the aggregation algebra, so utility degradation can be traced to noise magnitude, shot budgets, and client heterogeneity. Figure 4 highlights the systems-side constraints that materially influence federated quantum optimization on the edge.




Figure 4 Federated Training Workflow: Runtime and Participation Dynamics

The stacked latency decomposition reveals that local quantum execution dominates round time, while privacy overhead and uplink remain non-trivial and vary with message size and runtime conditions. The companion participation plot evidences-controlled client sampling, which directly affects privacy exposure and convergence stability under non-IID drift. Together, these plots support a reproducible interpretation of accuracy results in terms of communication, measurement cost, and participation dynamics.

Algorithm 1 Federated QML with Privacy-Preserving Measurement Statistics

Input: rounds T , local epochs E , client weights $\{n_k\}$, DP budget (ϵ, δ) , noise σ , step size η
 Initialize global parameters $\theta(0)$

```
For t = 0 to T-1 do
  Server selects client subset S_t
  Server broadcasts  $\theta(t)$  to all k in S_t
  For each client k in S_t (in parallel) do
     $\theta_k \leftarrow \theta(t)$ 
    For e = 1 to E do
      Sample minibatch from  $\mathcal{D}_k$ 
      Execute PQC  $U(x; \theta_k)$  with shot budget S
      Compute measurement statistics  $s_k$  from outcomes (expectations or histograms)
      Update  $\theta_k$  using local optimizer driven by  $s_k$ 
    End for
    Compute update  $\Delta\theta_k = \theta_k - \theta(t)$  (or compute sufficient statistics  $u_k$ )
    Clip statistics/update to bound sensitivity
    Privatize:  $\tilde{u}_k \leftarrow u_k + N(0, \sigma^2 I)$  (or privatize  $\Delta\theta_k$  similarly)
    Send  $\tilde{u}_k$  to server using secure aggregation if enabled
  End for
  Server aggregates:  $\bar{u} \leftarrow \sum_k (n_k / \sum_{j \in S_t} n_j) * \tilde{u}_k$ 
  Server updates  $\theta(t+1)$  using  $\eta$  and aggregated message  $\bar{u}$ 
End for
Output: final global parameters  $\theta(T)$ 
```

Table 4 summarizes the operational training regime and makes the privacy relevant levers explicit in one place. The values align rounds TTT, participation rate, and message dimension ddd with a clearly interpretable privacy composition horizon and exposure profile. The shot budget SSS is included as a first-class systems variable because it determines estimator variance, which, in turn, governs how strongly DP noise can be applied without collapsing signal. This table therefore anchors reproducibility for latency and utility comparisons.

Table 4 Training Schedule and Communication Budget				
Setting	Symbol	Value	Unit	Interpretation
Rounds	T	20	rounds	Privacy composition horizon
Local epochs	E	2	epochs	Compute-communication trade-off
Shots per minibatch	S	512	shots	Estimator stability on edge
Participation	$ S_t /K$	0.30 to 0.50	fraction	Exposure and robustness control
Message dimension	d	16	features	Statistics vector size

Experimental Design, Metrics, and Baselines

Evaluation targets edge analytics scenarios where privacy constraints and measurement costs jointly limit learning capacity. Datasets are partitioned to create controlled non-IID regimes, such as label skew and feature skew, ensuring that client distributions are meaningfully distinct. Device profiles are simulated using per-client limits on shots, circuit depth, and communication bandwidth. This yields a realistic operating envelope in which improvements must hold under both statistical and systems constraints.

Baselines are selected to isolate the contribution of privacy-preserving measurement statistics. Classical federated baselines include FedAvg with standard gradient exchange and DP-SGD-style noise, while quantum baselines include federated training without DP and centralized quantum training where data are pooled. Ablations further separate expectation-based statistics from histogram-based statistics, and secure aggregation is toggled independently.

This design provides attribution for utility loss and clarifies whether benefits arise from the quantum model, the privacy mechanism, or their interaction.

Primary model utility is measured by accuracy and macro-F1, while reliability is measured via Expected Calibration Error to quantify overconfidence under noise. Systems metrics include per-round latency, total communicated bytes, and effective shot consumption. Privacy reporting includes the final (ϵ, δ) for the full training horizon and a discussion of how composition interacts with client sampling. All metrics are computed on held-out client-local test sets to preserve the federated evaluation discipline. Calibration is computed as:

$$\text{ECE} = \sum_{m=1}^M \frac{|B_m|}{N} |\text{acc}(B_m) - \text{conf}(B_m)| \quad (5)$$

where B_m denotes a confidence bin, $\text{acc}(B_m)$ is empirical accuracy, and $\text{conf}(B_m)$ is mean predicted confidence. Equation (5) is essential in this methodology because DP noise can alter probability estimates even when accuracy remains stable, so calibration is treated as a first-class outcome for privacy-preserving edge deployment.

Figure 5 provides a compact empirical view of the privacy–utility–measurement trade space by mapping accuracy across privacy budgets and shot budgets. The heatmap reveals the expected saturation behavior where utility improves rapidly at low ϵ and then plateaus, while higher shot budgets reduce estimation variance and stabilize learning under noise. The companion calibration plot underscores that privacy noise affects not only accuracy but also probabilistic reliability. This visualization directly supports the edge deployment argument by linking utility to both privacy and measurement constraints.




Figure 5 Evaluation Matrix over Privacy and Measurement Budgets

Table 5 defines the comparison set used to attribute gains specifically to privacy-preserving measurement statistics rather than to federation alone or to quantum modeling alone. The baseline separation between centralized and federated settings establishes upper bounds on utility and isolates federation overhead. The privacy-enabled variants then localize the effect of DP and secure aggregation on the exact shared signal, namely statistics versus deltas. The classical FedAvg baseline provides a non-quantum reference that is necessary for judging whether quantum advantages persist under edge and privacy constraints.

Table 5 Baselines and Ablation Settings				
Method	Model Type	Federated	Privacy	Shared Signal
Centralized QML	PQC + lightweight head	No	No	Full-data training
FQML (no privacy)	PQC + lightweight head	Yes	No	Measurement stats or parameter deltas
FQML + DP	PQC + lightweight head	Yes	Gaussian DP	Privatized measurement statistics
FQML + DP + SecAgg	PQC + lightweight head	Yes	DP + secure aggregation	Aggregated privatized statistics
Classical FedAvg	Classical model	Yes	Optional DP	Gradients or parameter deltas

Result and Discussion

Overall Predictive Performance Against Baselines

Across all evaluation splits, the proposed Federated Quantum Machine Learning approach with privacy-preserving measurement statistics maintains competitive accuracy while imposing strict edge and privacy constraints. The dominant empirical pattern is that federated quantum learning without privacy yields the strongest federated quantum baseline, while privacy-enabled variants incur a controlled performance degradation that remains smaller than the gap between classical baselines and quantum baselines under the same device limitations. The calibration profile remains stable when privacy noise is calibrated to the shot budget, indicating that reliability loss is not proportional to accuracy loss.

Performance differences become most visible on non-IID client partitions where local label skew amplifies the effect of noise injected into measurement statistics. In these settings, the privacy-preserving variant preserves macro-F1 more effectively than a naïve privacy approach applied to high-dimensional histogram signals, indicating that compact expectation-based statistics reduce both message dimensionality and instability. The centralized quantum baseline remains an upper bound, confirming that the residual gap is explained primarily by federated heterogeneity and message protection rather than quantum expressivity limits.

Figure 6 shows that the proposed privacy-preserving federated quantum approach achieves a strong balance between accuracy and macro-F1, with a smaller degradation than the classical federated baseline under the same edge constraints. The gap between “FQML (no privacy)” and “FQML + DP (proposed)” quantifies the direct cost of protecting measurement statistics, and its magnitude indicates that the retained signal is sufficient for stable optimization. The centralized quantum results act as a meaningful upper bound rather than a trivial reference.




Figure 6 Overall Predictive Performance Across Baselines

A key insight in figure 6 is the tight coupling between accuracy and macro-F1 for privacy-enabled quantum variants, which indicates limited class-specific collapse under non-IID drift. The “DP + SecAgg” condition slightly reduces performance relative to the proposed DP-only condition, consistent with more conservative messaging and reduced per-client specificity after aggregation. The classical FedAvg baseline remains lower, indicating that the quantum feature space retains discriminative structure even when messages are restricted to protected statistics.

Table 6 consolidates utility, reliability, and latency in a single view, which is essential for edge analytics where predictive gains are only actionable if the runtime profile remains feasible. The proposed method improves both accuracy and calibration error compared with the classical baseline, demonstrating that the privacy constraint does not induce uncontrolled overconfidence. Latency increases for privacy-enabled variants are consistent with added local privacy processing and the messaging protocol, rather than indicating a fundamentally slower quantum execution path.

Table 6 Summary Metrics Under Edge Constraints

Method	Accuracy	Macro-F1	ECE	Avg Round Latency (ms)
Centralized QML	0.893	0.881	0.034	620
FQML (no privacy)	0.871	0.86	0.041	790
FQML + DP (proposed)	0.844	0.836	0.048	845
FQML + DP + SecAgg	0.836	0.829	0.051	875
Classical FedAvg	0.801	0.788	0.062	710

The calibration column supports the claim that privacy-preserving measurement statistics can maintain reliable probability outputs when shot budgets and noise parameters are co-tuned. The proposed method’s ECE remains within a narrow band relative to the non-private federated quantum setting, indicating that the injected noise behaves like controlled regularization rather than destabilizing the decision boundary. The DP plus secure aggregation setting shows a modest latency increase, consistent with extra coordination overhead, while the utility cost remains bounded.

Privacy Utility and Calibration Trade-offs

The privacy utility analysis varies the privacy budget and shot budget to quantify how much predictive signal survives protection at the measurement statistics level. Empirical trends show a clear saturation regime where increasing the privacy budget yields diminishing accuracy improvements after a moderate threshold, while increasing shots continues to improve stability until measurement noise becomes secondary to non-IID drift. This behavior indicates that measurement-based privatization is viable because the protected statistics remain low-dimensional and retain enough structure for aggregation to reconstruct consistent global updates.

Calibration trends exhibit a distinct sensitivity to privacy, even when accuracy changes are modest. Lower privacy budgets increase uncertainty in measurement summaries, which propagates into less confident and occasionally miscalibrated outputs. However, larger shot budgets systematically mitigate this effect by reducing estimation variance before privatization, enabling the privacy mechanism to preserve relative ranking signals that drive stable decision boundaries. This joint dependence confirms that edge viability is determined by the coupled allocation of privacy and measurement resources.

Figure 7 consolidates the privacy–measurement analysis into a single two-panel visualization, where the upper-level message is that utility and reliability are jointly controlled by the privacy budget and the shot budget. In the left panel, the accuracy heatmap shows a clear regime shift when moving from very strict privacy to moderate privacy, after which accuracy gains saturate as ϵ increases. The vertical gradients across shot budgets indicate that increasing shots consistently improves accuracy at every privacy level, confirming that measurement variance is a dominant and controllable bottleneck under protected learning.




Figure 7 Privacy-Measurement Trade-offs: Utility and Calibration

The right panel complements this by isolating reliability behavior through Expected Calibration Error trajectories across ϵ for multiple shot budgets. The curves decrease monotonically with larger privacy budgets and also shift downward as shots increase, indicating that stabilization before privatization improves confidence alignment even when the shared signal is noisy. The combined figure therefore supports an actionable interpretation: modest relaxations of privacy budget recover most of the accuracy, while shot allocation is a robust lever for reducing miscalibration without weakening privacy.

guarantees, which is critical for threshold-based decisions in edge analytics.

Table 7 translates the surfaces into concrete operating points that can be used as deployment presets, showing how accuracy and calibration shift as privacy is relaxed while measurement budget is fixed. The results reflect the saturation phenomenon, where accuracy gains between moderate and high ϵ are smaller than the gains achieved when moving from strict privacy to moderate privacy. The latency column remains stable because shot budget is held constant, isolating privacy effects from measurement effects.

Table 7 Representative Operating Points for The Proposed Method				
ϵ	Shots	Accuracy	ECE	Avg Round Latency (ms)
0.5	512	0.816	0.068	820
1	512	0.828	0.061	835
2	512	0.844	0.054	845
4	512	0.852	0.049	858
8	512	0.856	0.046	870

The calibration column in table 7 clarifies that reliability continues to improve even in regions where accuracy has largely plateaued. This indicates that privacy tuning can be justified not only for small accuracy gains but also for meaningful reductions in miscalibration, which affects decision thresholds in edge analytics workflows. The table also supports the methodological claim that measurement-statistics privatization enables stable training without excessive runtime growth, because privacy adjustments primarily alter noise scale rather than requiring additional quantum execution.

Communication Cost and Edge Latency Decomposition

The communication analysis indicates that privacy-preserving measurement statistics reduce message payload volatility by transmitting compact vectors rather than high-dimensional gradients or raw bitstring logs. The dominant cost driver is the per-round frequency of participation and the statistics dimensionality, while secure aggregation adds a bounded overhead that scales with the number of participating clients. Under the proposed design, communication remains stable across rounds, which is operationally important for edge links with variable throughput and strict energy budgets.

Latency results show that round time is primarily governed by local quantum execution, with privacy processing and uplink acting as secondary but non-negligible components. This decomposition clarifies why increasing the shot budget improves utility but must be treated as a systems decision, because it increases local execution time directly. Secure aggregation raises coordinator-side processing and adds handshake overhead, yet the incremental latency remains smaller than the cost of increasing shots to compensate for severe privacy budgets.

Figure 8 is presented as a single two-panel visualization that links the protocol’s bandwidth footprint to its end-to-end runtime behavior under realistic participation variability. The left panel reports per-round communication payload, showing that statistics-only messaging remains consistently small because payload size scales with the number of participating clients and the

fixed dimensionality of the measurement-statistics vector, rather than with model parameter count. The separation between the “Stats-only” and “Stats + SecAgg” curves indicate a predictable, bounded overhead from secure aggregation, while the large gap to the classical FedAvg-like baseline emphasizes that gradient-style updates impose substantially higher uplink demand in edge networks.




Figure 8 Communication Payload and Latency Decomposition

The right panel decomposes per-round latency into additive components, making the dominant bottleneck explicit. The stacked areas show that local quantum execution contributes the largest share of latency across rounds, while privacy processing and uplink add secondary but stable overhead that does not introduce large variability. The total-latency curve remains smooth, indicating that communication cost fluctuations in the left panel do not translate into unstable runtime excursions. Taken together, the combined figure supports a systems-level conclusion: the proposed protocol is bandwidth-efficient and runtime-predictable, and optimization efforts should prioritize shot efficiency and circuit execution rather than coordinator-side overhead.

Table 8 quantifies the practical advantage of measurement-statistics communication by contrasting average and tail payload behavior. The proposed method remains within a few kilobytes per round even at high participation, supporting feasibility for constrained uplinks and reducing variability that would otherwise cause queueing delays. The tail latency values remain bounded, consistent with the plot-level finding that local quantum execution dominates runtime and that protocol overhead does not produce large excursions.

Table 8 Communication and Latency Summary by Method				
Method	Avg Payload/Round (KB)	P95 Payload/Round (KB)	Avg Round Latency (ms)	P95 Round Latency (ms)
Stats-only (proposed)	3.1	3.8	845	930
Stats + SecAgg	4.8	5.6	875	965
Classical FedAvg-like	33.9	38.4	710	780

The comparison also clarifies a deployment trade-off: the classical baseline can appear faster per round due to the absence of quantum execution, yet it consumes an order of magnitude more bandwidth and yields lower predictive utility in earlier sections. This trade-off is central for edge analytics, where communication and energy budgets often dominate feasibility. Secure

aggregation increases payload and latency modestly, providing stronger exposure reduction without changing the systems regime.

Robustness Under Non-IID Drift, Client Sampling, and Perturbed Updates

Robustness experiments vary non-IID severity and participation rates to evaluate whether privacy-preserving measurement statistics remain aggregatable under distribution shift. Results show that moderate client sampling improves stability by reducing repeated exposure and implicitly regularizing against spurious client-specific patterns, but overly aggressive sampling slows convergence and increases variance in the global model. Under high non-IID severity, the proposed approach retains macro-F1 more effectively than histogram-based messaging, indicating that compact expectation statistics are less prone to noise amplification.

Perturbed-update tests emulate malicious or unreliable clients by injecting biased statistics into a fraction of participants. The proposed method remains resilient when the coordinator applies simple robust aggregation heuristics, because DP noise and clipping already bound extreme contributions and limit the influence of any single client. Secure aggregation preserves privacy but does not directly add robustness, so robustness improvements primarily arise from bounded statistics and aggregation rules rather than from cryptographic concealment.

Figure 9 is presented as a single two-panel visualization that summarizes robustness from two complementary angles: distribution shift and unreliable participation. The left panel plots macro-F1 against a non-IID severity index, showing a controlled and monotonic degradation as heterogeneity increases. The expectation-statistics configuration consistently stays above histogram-statistics under the same privacy constraints, and the gap widens at higher severity, indicating that low-dimensional expectation-based measurement summaries preserve class-level signal more reliably when aggregation is stressed by drift and privatization noise. The classical FedAvg baseline remains lower across the full range, reinforcing that the quantum feature space retains discriminative structure even under severe heterogeneity.




Figure 9 Robustness Under Drift and Perturbed Participation

The right panel evaluates sensitivity to perturbed participation by plotting accuracy against the fraction of clients that send biased or unreliable messages. The robust aggregation curve degrades slowly and remains stable at moderate

perturbation rates, whereas naïve aggregation drops more sharply, demonstrating that privacy alone is not a substitute for robustness controls. The combined figure therefore supports a clear operational conclusion: robustness under realistic edge conditions requires both bounded, privacy-protected statistics and an aggregation rule that suppresses anomalous contributions, especially when client behavior is partially adversarial or noisy.

Table 9 consolidates robustness findings by reporting representative points under distribution shift and adversarial perturbations. The results support the design choice of expectation-based statistics by showing higher macro-F1 under the same non-IID severity, which implies better preservation of minority-class behavior when noise is present. The perturbation rows confirm that robust aggregation is a necessary operational component, because privacy alone does not prevent biased contributions from influencing the global model.

Table 9 Robustness Summary Across Regimes					
Regime	Setting	Method	Accuracy	Macro-F1	Notes
Non-IID	Severity 0.7	FQML + DP (expectation)	0.834	0.818	Stable aggregation under noise
Non-IID	Severity 0.7	FQML + DP (histogram)	0.819	0.796	Higher variance due to dimension
Non-IID	Severity 0.7	Classical FedAvg	0.792	0.762	Lower utility under drift
Perturbation	10% perturbed	Robust aggregation	0.836	0.829	Bounded influence
Perturbation	10% perturbed	Naive aggregation	0.824	0.814	Faster degradation

The table also clarifies the mechanism behind robustness: clipping and DP reduce update amplitude, while robust aggregation controls contribution weighting at the server. This combination yields bounded influence for anomalous clients without requiring direct inspection of raw data or unprotected messages. In edge analytics, this is practically important because devices can be unreliable or partially compromised, yet the learning system must remain stable while maintaining privacy guarantees and predictable communication costs.

Ablation Study: Measurement Statistics Type, Secure Aggregation, and Client Sampling

The ablation study isolates which methodological choices account for the retained utility under privacy constraints. Results indicate that expectation-based statistics are the primary driver of stability because they reduce message dimensionality and concentrate signal into low-variance summaries that tolerate DP perturbation. Histogram-based messaging can approach expectation-based utility only when shot budgets are substantially increased, which is not operationally attractive for edge devices because it shifts cost into quantum execution time and energy.

Secure aggregation primarily improves exposure reduction rather than predictive utility. The measured utility difference between DP-only and DP plus secure aggregation is consistently small, suggesting that the masking protocol introduces modest protocol overhead and additional aggregation conservatism without materially changing the learning signal. Client sampling exerts a clearer

performance effect, where moderate sampling improves generalization and privacy exposure, but very low participation yields noisier global updates and slower convergence, especially under non-IID distributions.

Figure 10 is presented as a single two-panel visualization that isolates which design choices most strongly determine utility and reliability under privacy-preserving federated quantum learning. The left panel reports accuracy across ablation settings, showing that expectation-based statistics deliver the best utility when combined with differential privacy, while histogram-based statistics consistently underperform due to higher dimensionality and greater variance under privatization. The small downward shift when enabling secure aggregation indicates that its primary contribution is exposure reduction rather than utility gain, whereas the sampling variants show that participation rate is a stability lever: low sampling weakens aggregation signal and reduces accuracy, while higher sampling improves averaging and convergence.




Figure 10 Ablations on Statistics Type, Secure Aggregation, and Sampling

The right panel reports Expected Calibration Error for the same ablations and demonstrates that reliability is more sensitive than accuracy to the choice of statistics representation. Histogram-based messaging yields higher ECE, consistent with noisier protected signals that distort confidence estimates, while expectation-based messaging improves calibration by preserving relative score structure after privatization. Secure aggregation introduces a modest increase in ECE, aligning with slightly more conservative aggregation behavior, and higher sampling yields the best calibration by smoothing client-level noise through broader participation. Taken together, the combined figure supports a practical configuration: expectation-based statistics with moderate-to-high sampling, and secure aggregation enabled when exposure reduction is prioritized over a small reliability and utility shift.

Table 10 provides a deployment-oriented interpretation of the ablation results by linking each configuration choice to measurable outcomes and operational consequences. The strongest evidence favors expectation-based statistics as the default messaging format because it achieves higher accuracy and lower calibration error under the same privacy conditions, while also reducing communication burden. Secure aggregation is best interpreted as a privacy hardening layer rather than a performance enhancement, since its utility impact is small but consistently negative.

Table 10 Ablation Summary and Operational Interpretation				
Ablation Factor	Setting	Accuracy	ECE	Operational Implication
Statistics type	Expectation + DP	0.844	0.054	Preferred for stable, compact messaging
Statistics type	Histogram + DP	0.829	0.062	Requires higher shots to recover stability
Secure aggregation	Expectation + DP + SecAgg	0.836	0.057	Exposure reduction with small utility cost
Secure aggregation	Histogram + DP + SecAgg	0.822	0.065	Added overhead amplifies instability
Client sampling	Low sampling (0.20)	0.831	0.06	Higher variance, slower convergence
Client sampling	High sampling (0.60)	0.846	0.053	More stable aggregation, better calibration

The client sampling rows show that participation is not merely a scalability lever but a stability control that shapes both utility and reliability under privatized messaging. Low sampling increases global update variance and amplifies the effect of DP noise, while high sampling improves averaging and stabilizes confidence outputs. In aggregate, the ablations support a recommended configuration for edge analytics consisting of expectation-based statistics, moderate-to-high participation, and secure aggregation enabled when exposure reduction is prioritized over the small additional utility cost.

Conclusion

Chapter 5 consolidates the empirical evidence that Federated Quantum Machine Learning can be operationalized for edge analytics when learning signals are restricted to privacy-preserving measurement statistics rather than raw samples or high-dimensional gradients. The results show that the proposed approach sustains strong predictive utility relative to classical federated baselines under the same edge constraints, while maintaining competitive macro-F1 and stable reliability. The end-to-end behavior confirms that the hybrid quantum model retains discriminative structure even under non-IID client distributions and communication limitations that commonly degrade federated performance.

A central conclusion is that the deployable performance regime is governed by a coupled allocation of privacy and measurement resources. Accuracy exhibits saturation with increasing privacy budget, while calibration remains more sensitive, making reliability metrics necessary for operational decision-making in privacy-constrained edge settings. The communication and latency analyses further demonstrate that compact statistics messaging keeps per-round payloads predictable and small, and that end-to-end runtime is dominated by local quantum execution rather than protocol overhead. Robustness experiments indicate that bounded statistics and robust aggregation mitigate the influence of unreliable or perturbed clients without weakening privacy protections.

The ablation findings support clear design recommendations for practical deployments. Expectation-based statistics should be treated as the default message type because they improve stability under privatization and reduce bandwidth demands relative to histogram-based alternatives. Secure

aggregation should be enabled when exposure reduction is prioritized, as its utility cost is modest and its overhead remains predictable under variable participation. Future work should extend the approach to adaptive privacy accounting across rounds, hardware-aware ansatz selection under dynamic shot budgets, and stronger defenses against targeted poisoning under encrypted aggregation, thereby advancing scalable and trustworthy edge federated quantum learning for real-world analytics.

Declarations

Author Contributions

Conceptualization: Conceptualization: M.I., S.E.; Methodology: M.I., S.E.; Software: M.I., S.E.; Validation: M.I., S.E.; Formal Analysis: M.I.; Investigation: M.I., S.E.; Resources: M.I., S.E.; Data Curation: M.I.; Writing – Original Draft Preparation: M.I.; Writing – Review and Editing: M.I., S.E.; Visualization: S.E.; All authors have read and agreed to the published version of the manuscript.

Data Availability Statement

The data presented in this study are available on request from the corresponding author.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge Computing: Vision and Challenges," *IEEE Internet Things J.*, vol. 3, no. 5, pp. 637–646, 2016, doi: 10.1109/JIOT.2016.2579198.
- [2] W. Y. B. Lim et al., "Federated Learning in Mobile Edge Networks: A Comprehensive Survey," *IEEE Commun. Surv. Tutorials*, vol. 22, no. 3, pp. 2031–2063, 2020, doi: 10.1109/COMST.2020.2986024.
- [3] D. C. Nguyen, M. Ding, P. N. Pathirana, A. Seneviratne, J. Li, and H. Vincent Poor, "Federated Learning for Internet of Things: A Comprehensive Survey," *IEEE Commun. Surv. Tutorials*, vol. 23, no. 3, pp. 1622–1658, 2021, doi: 10.1109/COMST.2021.3075439.
- [4] P. Kairouz and H. B. McMahan, "Advances and Open Problems in Federated Learning," *Foundations and Trends in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021, doi: 10.1561/22000000083.

- [5] B. Rao, J. Zhang, D. Wu, C. Zhu, X. Sun, and B. Chen, "Privacy Inference Attack and Defense in Centralized and Federated Learning: A Comprehensive Survey," *IEEE Trans. Artif. Intell.*, vol. 6, no. 2, pp. 333–353, 2025, doi: 10.1109/TAI.2024.3363670.
- [6] K. Bonawitz et al., "Practical Secure Aggregation for Privacy-Preserving Machine Learning," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, Dallas Texas USA: ACM, vol. 2017, no. October, pp. 1175–1191, 2017, doi: 10.1145/3133956.3133982.
- [7] M. Abadi et al., "Deep Learning with Differential Privacy," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, Vienna Austria: ACM, vol. 2016, no. October, pp. 308–318, 2016, doi: 10.1145/2976749.2978318.
- [8] I. Mironov, "Rényi Differential Privacy," in *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, Santa Barbara, CA: IEEE, vol. 2017, no. September, pp. 263–275, 2017, doi: 10.1109/CSF.2017.11.
- [9] M. Cerezo et al., "Variational quantum algorithms," *Nat Rev Phys*, vol. 3, no. 9, pp. 625–644, 2021, doi: 10.1038/s42254-021-00348-9.
- [10] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, "Barren plateaus in quantum neural network training landscapes," *Nat Commun*, vol. 9, no. 1, p. 4812, 2018, doi: 10.1038/s41467-018-07090-4.
- [11] M. Schuld and N. Killoran, "Quantum Machine Learning in Feature Hilbert Spaces," *Phys. Rev. Lett.*, vol. 122, no. 4, p. 040504, 2019, doi: 10.1103/PhysRevLett.122.040504.
- [12] S. Y.-C. Chen and S. Yoo, "Federated Quantum Machine Learning," *Entropy*, vol. 23, no. 4, p. 460, 2021, doi: 10.3390/e23040460.
- [13] M. Senekane, M. Mafu, and B. M. Taele, "Privacy-preserving quantum machine learning using differential privacy," in *2017 IEEE AFRICON*, Cape Town: IEEE, vol. 2017, no. November, pp. 1432–1435, 2017, doi: 10.1109/AFRCON.2017.8095692.
- [14] M. Chehimi, S. Y.-C. Chen, W. Saad, D. Towsley, and M. Debbah, "Foundations of Quantum Federated Learning Over Classical and Quantum Networks," *IEEE Network*, vol. 38, no. 1, pp. 124–130, 2024, doi: 10.1109/MNET.2023.3327365.
- [15] M. Arbaoui, M.-A. Brahmia, A. Rahmoun, and M. Zghal, "Federated Learning Survey: A Multi-Level Taxonomy of Aggregation Techniques, Experimental Insights, and Future Frontiers," *ACM Trans. Intell. Syst. Technol.*, vol. 15, no. 6, pp. 1–69, 2024, doi: 10.1145/3678182.
- [16] A. Brecko, E. Kajati, J. Koziorek, and I. Zolotova, "Federated Learning for Edge Computing: A Survey," *Applied Sciences*, vol. 12, no. 18, p. 9124, 2022, doi: 10.3390/app12189124.
- [17] V. Mothukuri, R. M. Parizi, S. Pouriyeh, Y. Huang, A. Dehghantanha, and G. Srivastava, "A survey on security and privacy of federated learning," *Future Generation Computer Systems*, vol. 115, no. February, pp. 619–640, 2021, doi: 10.1016/j.future.2020.10.007.
- [18] P. Qi, D. Chiaro, A. Guzzo, M. Ianni, G. Fortino, and F. Piccialli, "Model aggregation techniques in federated learning: A comprehensive survey," *Future Generation Computer Systems*, vol. 150, no. January, pp. 272–293, 2024, doi: 10.1016/j.future.2024.100000.

10.1016/j.future.2023.09.008.

- [19] M. Chen, N. Shlezinger, H. V. Poor, Y. C. Eldar, and S. Cui, "Communication-efficient federated learning," *Proc. Natl. Acad. Sci. U.S.A.*, vol. 118, no. 17, p. e2024789118, 2021, doi: 10.1073/pnas.2024789118.
- [20] M. Ahmadzai and G. Nguyen, "Federated Learning with Differential Privacy on Personal Opinions: A Privacy-Preserving Approach," *Procedia Computer Science*, vol. 225, no. 2023, pp. 543–552, 2023, doi: 10.1016/j.procs.2023.10.039.
- [21] X. Jin, Y. Yao, N. Yu, "Efficient secure aggregation for privacy-preserving federated learning based on secret sharing," *zgkxjsdxxb*, vol. 54, no. 1, p. 0104, 2024, doi: 10.52396/JUSTC-2022-0116.
- [22] W. M. Watkins, S. Y.-C. Chen, and S. Yoo, "Quantum machine learning with differential privacy," *Sci Rep*, vol. 13, no. 1, p. 2453, 2023, doi: 10.1038/s41598-022-24082-z.
- [23] R. Ballester, J. Cerquides, and L. Artilles, "Quantum federated learning: a comprehensive literature review of foundations, challenges, and future directions," *Quantum Mach. Intell.*, vol. 7, no. 2, p. 73, 2025, doi: 10.1007/s42484-025-00292-2.
- [24] A. Gupta, M. Kumar Maurya, K. Dhere, and V. Kumar Chaurasiya, "Privacy-Preserving Hybrid Federated Learning Framework for Mental Healthcare Applications: Clustered and Quantum Approaches," *IEEE Access*, vol. 12, no. September, pp. 145054–145068, 2024, doi: 10.1109/ACCESS.2024.3464240.