

Modeling Quantum Federated Autoencoder for Anomaly Detection in IoT Networks

Devashish Chaudhary*, Sutharshan Rajasegarar†, Shiva Raj Pokhrel‡

School of Information Technology, Deakin University, Geelong, Australia

*s224281473@deakin.edu.au, †srajas@deakin.edu.au, ‡shiva.pokhrel@deakin.edu.au

Abstract—We propose a Quantum Federated Autoencoder for Anomaly Detection, a framework that leverages quantum federated learning for efficient, secure, and distributed processing in IoT networks. By harnessing quantum autoencoders for high-dimensional feature representation and federated learning for decentralized model training, the approach transforms localized learning on edge devices without requiring transmission of raw data, thereby preserving privacy and minimizing communication overhead. The model leverages quantum advantage in pattern recognition to enhance detection sensitivity, particularly in complex and dynamic IoT network traffic. Experiments on a real-world IoT dataset show that the proposed method delivers anomaly detection accuracy and robustness comparable to centralized approaches, while ensuring data privacy.

Index Terms—IoT, quantum federated learning, anomaly detection, quantum autoencoder, network security.

I. INTRODUCTION

With the recent accelerated growth of interconnected devices, securing the network from attacks and timely detection of emerging anomalies have become increasingly challenging [1]. Conventional approaches require transmission of raw data to a centralized server for model training, which introduces severe privacy risks and creates a single point of failure; any compromise of the server might expose the entire dataset. Federated Learning (FL) [2], [3] mitigates this issue by training models locally on devices and sharing only the model parameters with a server, which then aggregates updates to form a global model [4].

Recent advances in quantum computing provide new capabilities for machine learning by exploiting superposition and entanglement to process complex data more efficiently. Integrating these capabilities into FL yields Quantum Federated Learning (QFL) [5], where quantum models are trained locally, and only parameter updates are communicated for aggregation. QFL holds significant promise for enhancing both performance and security in large-scale distributed networks, particularly for anomaly detection.

The primary contributions of this work are summarized below:

- We construct a fully operational hierarchical IoT network using Raspberry Pi 3B+ devices and XBee transceivers, enabling the generation of realistic, multilevel traffic patterns. This testbed provides high-fidelity data streams from which salient features for anomaly detection are systematically extracted.

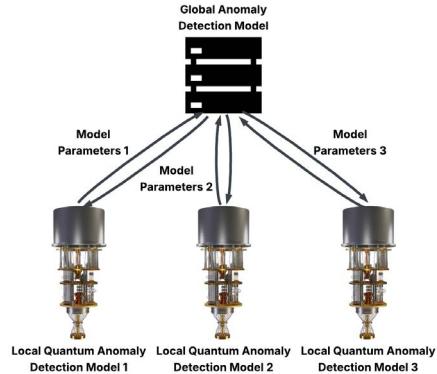


Fig. 1: Quantum Federated Autoencoder Framework for Anomaly Detection.

- We develop a new QFL architecture (Fig. 1) that enables the local training of quantum autoencoders on edge devices and supports both hierarchical and canonical FedAvg aggregation strategies. The framework is rigorously benchmarked against an equivalent centralized quantum-training baseline.
- Comprehensive experiments across heterogeneous IoT devices show that the proposed QFL framework preserves data privacy without degrading model utility, achieving detection performance statistically indistinguishable from fully centralized quantum training.

To the best of our knowledge, this work constitutes the first demonstration of a quantum autoencoder adapted for anomaly detection and deployed within a fully operational QFL setting.

II. RELATED WORK

Classical deep learning approaches have been extensively used for anomaly detection in networks [6], leveraging both supervised and unsupervised models. Many works have combined autoencoders with temporal architectures to capture dynamic traffic behavior. For instance, a PSO-Autoencoder-LSTM framework [7] integrates autoencoders for feature extraction with LSTMs for temporal dependency modeling, while Particle Swarm Optimization (PSO) tunes hyperparameters to improve detection accuracy and resilience against class imbalance in intrusion detection tasks. Despite their effectiveness, these approaches [6], [7] remain inherently centralized,

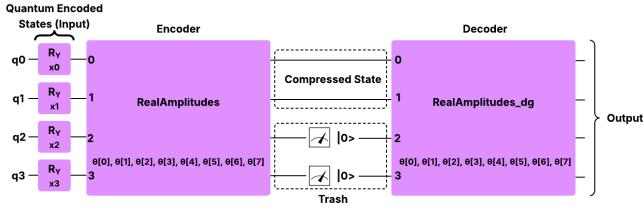


Fig. 2: Quantum Autoencoder with 4 qubits: 3 latent, 1 trash.

requiring raw data aggregation at a server, which exposes sensitive traffic to privacy and security risks.

FL addresses these limitations by enabling decentralized training without exposing raw data. Instead, clients train models locally and transmit only parameter updates to a central server for aggregation. Recent research has explored FL-based anomaly detection in network data. For example, Fed-ANIDS [8] deploys simple, variational, and adversarial autoencoders across distributed clients, employing FedAvg or FedProx for aggregation. Clients share model updates derived from local benign traffic, while the global model uses reconstruction error thresholds to detect anomalies. Such approaches effectively preserve privacy while maintaining strong intrusion detection performance.

Nevertheless, IoT traffic often exhibits highly complex, nonlinear and high-dimensional patterns that classical models struggle to capture efficiently. QML offers a promising alternative by exploiting superposition and entanglement to model such complexity more effectively. Hdaib et al. [9] demonstrated the potential of autoencoders for anomaly detection, benchmarking their integration with quantum one-class SVM, quantum k-nearest neighbor, and quantum random forest classifiers. Their results showed that an autoencoder combined with quantum kNN achieved superior performance, underscoring the advantage of quantum models. Nevertheless, existing QML methods are centralized and thus suffer from the same privacy and scalability issues as classical deep learning. To date, no study has explored a federated paradigm that combines the privacy-preserving benefits of FL with the representational power of quantum autoencoders.

III. PROPOSED QUANTUM FEDERATED AUTOENCODER

Fig. 2 shows a minimal 4-qubit Quantum autoencoder (QAE) for illustration. The experiments use a 10-qubit model (8 latent, 2 trash) as described in Section IV-A. Different from classical autoencoders, which reduce the dimensionality of input data by learning a compact representation that can be decoded to reconstruct the original input, QAEs extend to the quantum domain, efficiently compressing quantum states stored on n qubits into a smaller set of $m < n$ qubits.

The proposed framework leverages QFL to enable distributed training of QAEs across local devices while preserving data privacy. In the proposed QFL, each device trains a QAE locally and transmits only model parameters to a central server for aggregation. Aggregation can follow a standard

Algorithm 1 Federated Quantum Autoencoder Training

```

1: procedure TRAINQAE( $X, d, n, R, F, I$ )
2:   Input: Network traffic data  $X$ , PCA components  $d =$ 
   10, qubits  $n = 10$ , routers  $R = 3$ , federated rounds  $F = 5$ ,
   local iterations  $I = 50$ 
3:   Output: Trained global QAE parameters  $\theta_{\text{global}}$ 
4:   Perform PCA on  $X$  to reduce dimensionality to  $d$ 
   components
5:   for  $r = 1$  to  $R$  do
6:     Encode PCA features into  $n$  qubits using angle
   encoding ( $R_y$  rotations)
7:     Initialize parameterized quantum circuit (RealAm-
   plitude) as encoder
8:   end for
9:   for  $t = 1$  to  $F$  do                                 $\triangleright$  Federated rounds
10:    for  $r = 1$  to  $R$  do     $\triangleright$  Routers process in parallel
11:      for  $i = 1$  to  $I$  do     $\triangleright$  Local training iterations
12:        Forward pass through QAE to compute
   probabilities
13:        Compute loss  $L(\theta)$ : Equation (1)
14:        Update parameters  $\theta_r$  to minimize  $L(\theta)$ 
15:      end for
16:    end for
17:    Routers send local parameters to coordinator
18:    Coordinator performs hierarchical federated aver-
   aging (hierarchical FL) or FedAvg to obtain  $\theta_{\text{global}}$ 
19:    Coordinator sends updated global parameters back
   to all routers
20:  end for
21:  return  $\theta_{\text{global}}$ 
22: end procedure

```

FedAvg or a *hierarchical FL* protocol. In the *hierarchical FL*, updates from child nodes are first combined at parent nodes before progressively forming a global model (See Fig. 1, Algorithm 1).

A QAE compresses quantum states by mapping an input $|\psi\rangle$ into a tensor product of compressed and “trash” qubits:

$$U|\psi\rangle = |\phi\rangle_C \otimes |0\rangle_T,$$

where $|\phi\rangle_C$ retains relevant information and $|0\rangle_T$ is the fixed state for the trash qubits. Reconstruction is achieved by applying the inverse unitary:

$$U^\dagger(|\phi\rangle_C \otimes |0\rangle_T) = |\psi\rangle.$$

Unlike classical autoencoders, QAEs cannot discard qubits, so the network explicitly minimizes residual information in trash qubits. For a system with n qubits, where the last k qubits form the trash subsystem, we define a cost function as the total probability of measuring “bad” states in the trash qubits (all basis states except $|0\rangle_T$):

$$L(\theta) = \frac{1}{N} \sum_{i=1}^N \sum_{b \in \text{bad states}} \Pr[\text{trash qubits in state } b \mid X_{\text{train}}^{(i)}, \theta], \quad (1)$$

TABLE I: Attack scenarios: C – Coordinator, A – Attacker, Rx – Router, Ex – Edge/Leaf device.

Scenario I				Scenario II				Scenario III						
Normal	E1 → R1	E2 → R1	E3 → R3	E4 → R2	R1 → C	R2 → C	R3 → R2	E1 → R1	E2 → R1	E3 → R3	E4 → R2	R1 → C	R2 → C	R3 → C
Attack	E1 → R2	E2 → R2	E3 → R1	E4 → R1	R1 → R2	R2 → R1	R3 → R1	E1 → A	E2 → A	E3 → A	E4 → A	R1 → A	R2 → A	R3 → A
	E1 → R3	E2 → R3	E3 → R2	E4 → R3	R1 → R3		R3 → C							
	E1 → C	E2 → C	E3 → C	E4 → C										

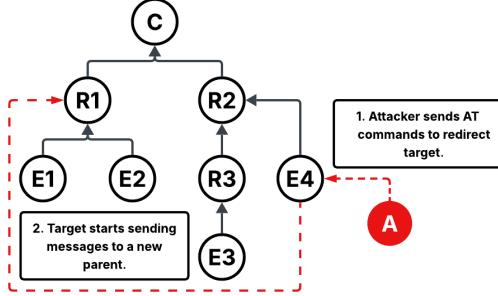


Fig. 3: Network topology with Coordinator (C), Routers (R), Edge nodes (E) and Attack node (A).

with $\text{Pr}[\cdot]$ obtained from the parameterized quantum circuit and N the batch size. Minimizing this loss ensures that the compressed qubits retain meaningful information while the trash qubits converge to $|0\rangle_T$.

By adapting QAE into the QFL framework, our approach achieves privacy-preserving, distributed quantum compression and anomaly detection across IoT devices, combining the representational power of quantum models with the scalability of FL.

IV. EXPERIMENTAL SETUP

We have implemented a real-world IoT testbed consisting of Raspberry Pi 3B+ devices equipped with XBee S2C ZigBee modules for wireless communication [10]. The network was modelled hierarchically, as illustrated in Fig. 3. One node was designated as an attacker to launch security attacks and generate both benign and malicious network traffic. Each edge device (E) sends data packets with timestamps. Intermediate routers (R) record reception timestamps before forwarding packets to the next hop, while the coordinator (C) maintains a complete log of packet paths and timings.

Initially, network traffic logs were recorded under normal operating conditions. Relevant features were extracted for each one-minute time window, including mean and first-hop delay, delay quartiles, Shannon entropy, per-type and overall communication counts, and average hops per communication. Subsequently, redirection attacks were executed by exploiting ZigBee attention (AT) commands, generating malicious traffic for evaluation. Normal traffic was collected for 5 hours for training, with an additional 1 hour for validation. Each attack session included 20 minutes of normal traffic, 5 minutes of attack traffic, and 10 minutes of normal traffic, in a sequence. Three scenarios were created to generate attacks in the network, as shown in Table I.

TABLE II: Performance metrics across routers for FL methods

Routers	Method	Accuracy	Precision	Recall	F1
R1	Hierarchical FL	0.91	0.94	0.91	0.92
	FedAvg FL	0.93	0.94	0.93	0.93
	Centralized	0.88	0.93	0.88	0.90
R2	Hierarchical FL	0.91	0.91	0.91	0.91
	FedAvg FL	0.91	0.91	0.91	0.91
	Centralized	0.91	0.91	0.91	0.91
R3	Hierarchical FL	0.99	0.99	0.99	0.99
	FedAvg FL	0.94	0.96	0.94	0.95
	Centralized	0.99	0.99	0.99	0.99

Features were computed locally at each router for federated learning, whereas centralized training combined and shuffled data from all routers (R1, R2, R3). Preprocessing included MinMax scaling and principal component analysis (PCA), fitted on training data and applied to validation and test sets. For federated training, each router's training set was partitioned into 5 subsets corresponding to 5 FL rounds.

A. Training Configuration

Training was conducted in Python 3.11.7 using Qiskit on a quantum simulator. The QAE utilized 10 qubits (8 latent, 2 trash), Ry rotations for feature transformation, and a RealAmplitude circuit as the parameterized encoder. Optimization employed COBYLA with 50 iterations for FL experiments and 100 iterations for centralized training. Federated learning was conducted over 5 rounds across 3 clients, with MinMax scaling and PCA for reducing 31 original features to 10 principal components. Experiments were conducted on an Apple MacBook Pro (M2, 8-core CPU, 10-core GPU).

B. Thresholding for Anomaly Detection

Anomaly detection is performed using a threshold (τ) computed from the reconstruction fidelity on the validation set: $\tau = \mu - 4\sigma$, where μ and σ are the mean and standard deviation of fidelities. Test samples with fidelity below τ are classified as anomalies, providing a statistically robust cutoff that balances sensitivity and variability in reconstruction performance.

V. EVALUATION: RESULTS AND DISCUSSION

We compared two different types of FL methods with a centralized approach for anomaly detection in IoT networks.

Table II presents the performance metrics across the three routers for different learning methods. Overall, federated learning approaches deliver performance on par with or exceeding centralized training, demonstrating that the privacy-preserving distributed learning can maintain high anomaly detection effectiveness without requiring raw data aggregation.

Examining individual routers provides further insight into the strengths of hierarchical QFL. For Router 1, standard

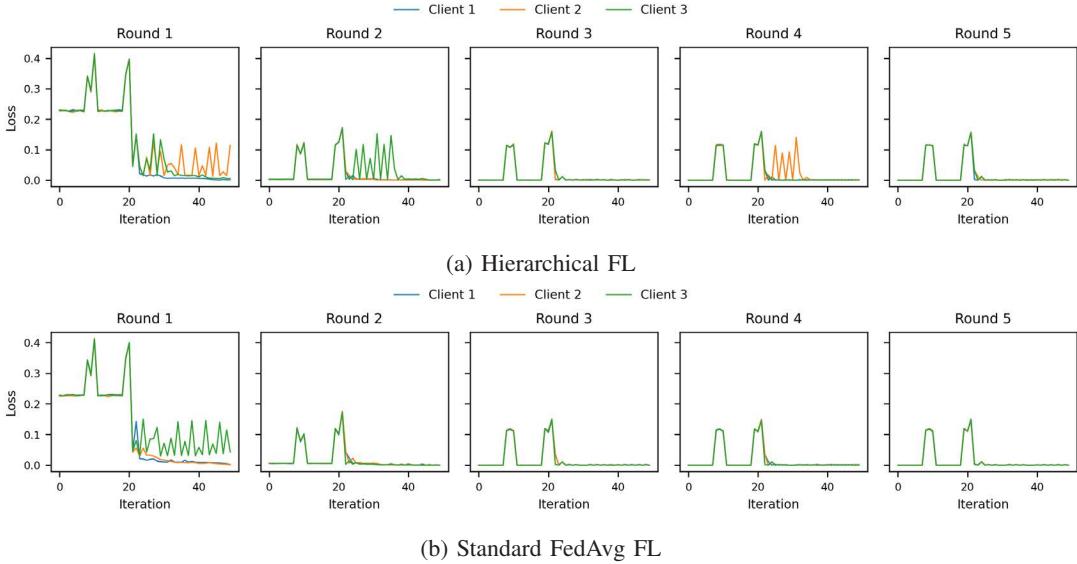


Fig. 4: Comparison of FL methods for 3 devices over 5 rounds: (a) Hierarchical FL and (b) Standard FedAvg FL.

FedAvg slightly outperforms hierarchical QFL, while the centralized model shows marginally lower accuracy and F1-score, suggesting that uniform local patterns allow standard aggregation to perform well. For Router 2, all methods perform equivalently, indicating a simple or homogeneous data distribution across clients. In contrast, Router 3 highlights the advantage of hierarchical aggregation: both hierarchical FL and centralized training achieve near-perfect scores, whereas standard FedAvg falls slightly short. This demonstrates that hierarchical FL more effectively captures local variations and inter-node dependencies, which is particularly important in heterogeneous or non-uniform IoT networks. Collectively, these findings justify the adoption of hierarchical federated learning as a robust, privacy-preserving approach capable of matching or surpassing centralized methods.

Figures 4a and 4b show the evolution of training loss over iterations. Initial spikes in the first round reflect the model adapting to diverse local data distributions. In subsequent rounds, the loss begins to decrease and converges more smoothly, indicating that the models progressively internalize the underlying patterns across clients. This trend confirms that both standard and hierarchical QFL enable stable and effective distributed learning, with hierarchical aggregation providing additional resilience against local data heterogeneity.

Table II reveals that the federated approaches achieve accuracy and F1-scores comparable to the centralized baseline, demonstrating that privacy-preserving training does not compromise effectiveness. Router-specific outcomes further validate data heterogeneity: while FedAvg slightly outperforms in Router 1, hierarchical FL excels in Router 3, confirming the benefit of topology-aware aggregation. Figures 4a and 4b show a smooth convergence of training loss across rounds, evidencing the stability of federated optimization.

VI. CONCLUSION

We evaluated anomaly detection using quantum autoencoders in a federated learning framework with both hierarchical and standard aggregation. Using data collected from a real Raspberry Pi 3B+ testbed, we showed that the QFL approaches achieve performance comparable to a centralized approach while preserving data privacy.

REFERENCES

- [1] S. Singh, G. Madaan, A. Singh, D. Pandey, A. S. George, B. K. Pandey *et al.*, “Empowering connectivity: exploring the internet of things,” in *Interdisciplinary Approaches to AI, Internet of Everything, and Machine Learning*, 2025, pp. 89–116.
- [2] D. C. Nguyen, M. Ding, P. N. Pathirana, A. Seneviratne, J. Li, and H. V. Poor, “Federated learning for internet of things: A comprehensive survey,” *IEEE comms. sys. & tut.*, vol. 23, no. 3, pp. 1622–1658, 2021.
- [3] D. Chaudhary, S. Rajasegarar, and S. R. Pokhrel, “Towards adapting federated & quantum machine learning for network intrusion detection: A survey,” *arXiv preprint arXiv:2509.21389*, 2025.
- [4] S. I. Nanayakkara, S. R. Pokhrel, and G. Li, “Understanding global aggregation and optimization of federated learning,” *Future Generation Computer Systems*, vol. 159, pp. 114–133, 2024.
- [5] D. Gurung, S. R. Pokhrel, and G. Li, “Quantum federated learning for metaverse: Analysis, design, and implementation,” *IEEE Trans. on Network and Service Management*, vol. 22, no. 3, pp. 2595–2606, 2025.
- [6] Y. Wu, D. Wei, and J. Feng, “Network attacks detection methods based on deep learning techniques: A survey,” *Security and Communication Networks*, vol. 2020, no. 1, p. 8872923, 2020.
- [7] S. Narmadha and N. Balaji, “Improved network anomaly detection system using optimized autoencoder- lstm,” *Expert Systems with Applications*, vol. 273, p. 126854, 2025.
- [8] M. J. Idrissi, H. Alami, A. El Mahdaouy, A. El Mekki, S. Oualil, Z. Yartaoui, and I. Berrada, “Fed-anids: Federated learning for anomaly-based network intrusion detection systems,” *Expert Systems with Applications*, vol. 234, p. 121000, 2023.
- [9] M. Hdaib, S. Rajasegarar, and L. Pan, “Quantum deep learning-based anomaly detection for enhanced network security,” *Quantum Machine Intelligence*, vol. 6, no. 1, p. 26, 2024.
- [10] D. Chaudhary, S. Rajasegarar, S. R. Pokhrel, L. Pan, and R. D., “In-network attack detection with federated deep learning in IoT networks: Real implementation and analysis,” in *2025 IEEE Conference on Engineering Informatics (ICEI)*, 2025, pp. 1–9.