

Quantum-Resilient Federated Intrusion Detection: A Hybrid Quantum Classical Framework For Safeguarding U.S. Critical Infrastructure In The Post-Quantum Era

Abidul Alam¹, Md. Salahuddin Gazi², Salman Mohammad Abdullah³, Masrufa Tasnim⁴, Md Himeluzzaman⁵, Mahbub Ahmed Nabil⁶, Kabita Ahamia Akter⁷, Sonia Akter⁸

¹Bachelor of Business Administration (BBA) in Marketing, International University of Business Agriculture and Technology (IUBAT), Dhaka, Bangladesh

²B.S.S. (Honours) in Sociology, Jagannath University, Dhaka, Bangladesh

³B.Sc. in Computer Science and Engineering (CSE), Daffodil International University, Bangladesh

⁴Master of Business Administration (MBA) in Human Resource Management, University of Dhaka, Dhaka, Bangladesh

⁵B.Sc. in Computer Science and Engineering, Leading University, Sylhet, Bangladesh

⁶Bachelor's in Computer Science and Engineering, North South University, Dhaka, Bangladesh

⁷Bachelor's Degree in Human Resource Management, American International University-Bangladesh (AIUB), Dhaka, Bangladesh

⁸Bachelor of Business Administration in Management, Narsingdi Government College, Dhaka, Bangladesh

Abstract—Large scale quantum computation poses a threat to the public-key based encryption that protects the privacy of communications traversing U.S. critical infrastructure and defense networks. This property strongly motivates detection systems against “harvest-now, decrypt-later” adversaries and is not compatible with approaches that limit themselves to either centralizing sensitive traffic or searching for classical decision boundaries. We introduce a quantum-resilient federated intrusion-detection solution that combines a hybrid quantum-classical classifier—a classical feature encoder feeding the expectation-value outputs of a variational quantum circuit (VQC) into a classical decision head—with the federated averaging scheme operating across physically distributed infrastructure sites, such that raw traffic never leaves its site of origin. You are the eighth of its kind, specifically designed to receive a post-quantum key-encapsulation mechanism.

The proposed federated hybrid model achieves a 77.65% accuracy and a macro-averaged F1 of 0.6200 on the NSL-KDD benchmark (5-class attack taxonomy), outperforming both centralized and federated classical multilayer-perceptron baselines on macro-F1 (0.5975 and 0.5949 respectively) and significantly outperforms a Random-Forest baseline (0.4734), indicative of improved discrimination of rare Remote-to-Local (R2L) and User-to-Root (U2R) attacks. The accuracy-cost frontier and the degradation due to non-independent, identically distributed partitions is quantified via ablations over qubit count, circuit depth, client population and data heterogeneity. We honestly report training-cost overheads, and discuss examples where most of the performance improvements come from traversing a barren plateau at greater depth circuits. The results indicate that near-term variational quantum models can be part of an intrusion detection system that is privacy-preserving and sensitive to rare classes while the surrounding protocol is hardened to the post-quantum world.

Keywords: quantum machine learning; federated learning; intrusion detection; variational quantum circuits; post-quantum cryptography; critical infrastructure security; NSL-KDD.

1. Introduction

The security of cryptographic primitives that underpin critical infrastructure—power grids, water systems, transportation, and defense communication

networks—is based on the classical intractability of integer factorization or discrete logarithms. On a gapless, fault-free quantum computer, Shor’s algorithm does precisely that, lowering the cost of breaking RSA and elliptic-curve schemes from sub-exponential to polynomial time. [1] Because of Grover’s search, the effective security of symmetric primitives is also effectively halved [2]. While fault-tolerant machines are still not online, the adversaries harvesting high-value, long-lived defense data today using “harvest-now, decrypt-later” methodology can break the encryption when that quantum capability matures [3,4].

This entails two consequences for those who defend such networks. As a first step, the cryptographic envelope securing inter-site communication should migrate to post-quantum key establishment [3,4]. Second, and less evidently, the intrusion-detection systems (IDS) which observe these networks also need to do so without aggregating up raw, sensitive traffic into a single targetable source of attack and by working on information representations that are difficult for the as-classically-limited-adversary to predict. The first is allayed by federated learning which retains the data physically local and only exchanges updates to models [5,6] while quantum machine learning (QML) uniquely provides feature maps in exponentially large Hilbert spaces unexplored by classical models [7,8,9].

We unify these ideas. Using federated averaging, our framework trains a hybrid quantum-classical classifier across simulated infrastructure sites, sends updates using key-encapsulation channel that is capable of accepting post-quantum mechanism and tests on the well-known NSL-KDD [10] benchmark. The variational quantum circuit is made to be shallow and relatively low-width so that it can be trained on near-term (NISQ) hardware while its classical simulation remains efficient [11,12]. The contributions are:

- An architecture for a federated IDS that is resilient to quantum computing threats: A variational quantum classifier combined with federated averaging and a post-quantum-ready secure-aggregation channel.
- My main contributions to this paper are (iii) an empirical investigation of adversarial robustness on the confirmation bias and voxnet models, using a matchmaking strategy which returns classical baselines matched in terms of transfer characteristics (i.e.
- (iii) systematic ablations over qubit count, circuit depth, client population and Dirichlet data heterogeneity combined with an accurate assessment of quantum-simulation av cloud-based obfuscation costs across the nodal workload.

The rest of the paper is organized as follows. Section 2 reviews related work. Section 3 describes the threat model, as well as the quantum and federated preliminaries. Section 4 describes the framework we propose. Section 5 explains the experimental setup followed by Section 6 outlining results and discussion, Section 7 specifying limitations, and finally Section 8 concludes the article.

2. Related Work

2.1 Machine learning for intrusion detection

A wide variety of survey literature exists covering general machine-learning approaches to intrusion-detection [13,14]. The NSL-KDD dataset [10] was

introduced with the aim of eliminating redundancy and imbalance in the original KDD'99 corpus, which is still a commonly used benchmark as it provides suitably sized and reproducible train and test partitions. A challenging task is the imbalanced class distribution between high-frequency Denial-of-Service (DoS) traffic, and rare classes of R2L and U2R types, which lowers mega-averaged metrics for classifiers designed for accuracy optimization [13].

2.2 Federated learning

Federated averaging (FedAvg) [5] with localisation of the data and weights averaged relative to client size; communication-efficiency strategies [6]; and finally a systems/strategy approach [15] laid the foundations for what is now considered canonical in sensitive domains. Unfortunately, secure aggregation protocols [16] permit a server to aggregate individual client updates without reading them individually — we preserve this and extend it towards post-quantum key establishment.

2.3 Quantum machine learning

Quantum machine learning [7] includes kernel methods on feature Hilbert spaces [8,17], parameterized (variational) quantum circuits as trainable models [9,18–19], and quantum neural networks on near-term processors [20]. Gradient-based training machinery is provided by quantum-circuit learning [19] and parameterized circuits [9] which we implemented here with automatic differentiation [21]. McClean et al. Exponential, vanishing gradients: [22] found barren plateaus in deep quantum networks — we see this signature when increasing circuit depth. Although variational circuits have been applied to reinforcement learning and classification tasks on nearterm devices [23], their combination with federated learning for post-quantum error-intrusion detection has also not received attention, which this work addresses.

3. Preliminaries and Threat Model

3.1 Post-quantum threat model

Let us assume a defense or critical-infrastructure network whose inter-site sessions are protected by a public-key scheme—commonly RSA or an elliptic-curve variant—whose security reduces to the assumed classical difficult of integer factorization or the discrete-logarithm problem. These are precisely the assumptions that a fault-tolerant quantum adversary having implemented Shor's algorithm violates, with recovery of the underlying private keys in time which grows polynomially with key length:

$$T_{\text{Shor}}(N) = O((\log N)^3) \ll T_{\text{classical}}(N), (1)$$

Thus, all ciphertext collected today would be vulnerable to decryption with such a machine [1,3]. The assumptions made are therefore an (a) honest-but-curious aggregation server, (b) clients unwilling to reveal raw traffic and (c) a network-level adversary able to record update traffic for later quantum-assisted cryptanalysis. It mixes data locality with a key-encapsulation channel, which can be instantiated with a post quantitic mechanism [4].

3.2 Variational quantum classifier

The classifier applies to a preprocessed feature vector that is first mapped by an

efficient classical encoding into n real-only angles of rotation – one for each qubit. These angles are inputted in the n -qubit quantum register via angle embedding, which consists of preparing a product state by applying the single-qubit Pauli-Y rotation gate on each qubit; additionally, the prior scaling of inputs to interval $[0, \pi]$ guarantees a unique and non-saturating relationship between feature size and its respective rotation angle

$$|\psi(x)\rangle = \bigotimes_{i=1}^n RY(\pi x_i) |0\rangle^{\otimes n}, (2)$$

This embedded state is processed with a strongly-entangling variational ansatz of depth L , where in each layer, every qubit first performs a general single-qubit rotation $R(\theta) = RZ(\alpha)RY(\beta)RZ(\gamma)$ —three trainable parameters per qubit per layer—before a ring of CNOT gates entangle neighbouring qubits to enable correlations to propagate across the entire register. Using this simple building block for entangling rotations, the trainable unitary using L such blocks is given

$$U(\theta) = \prod_{l=1}^L \left[\bigotimes_i R(\theta_{l,i}) \right] \cdot \text{Cent}, (3)$$

Each Cent represents the ring of CNOT gates that connect neighbouring qubits and thus couple the individual rotations into a genuine multiqubit transformation, $R(\theta_{l,i})$ is a general single-qubit rotation acting on qubit i in layer l and L is the depth of circuit. Then, by considering L independent L such entangling layers in a series of gates to make the parameterized state more expressive while keeping both gate counts linear with number of qubit and depth.

You do not read amplitudes directly, but extract information from the evolved quantum state by measurement. We measure each qubit in the computational basis and estimate its Pauli-Z expectation value from repeated compilations of this circuit (shots), producing a bounded real-valued feature per qubit that encodes how the variational circuit has pushed the input tensor along this direction:

$$f_i(x; \theta) = \langle \psi(x) | U^\dagger(\theta) Z_i U(\theta) | \psi(x) \rangle \in [-1, 1], (4)$$

The circuit gives n expectation values, forming a compact quantum feature vector which is used by a classical linear layer plus softmax to output a normalized probability distribution over the five attack classes. While enabling high-dimensional feature transformation through class planarization of the quantum circuit, delegating the final decision to a lightweight classical head gives the model an inexpensive way to calibrate its class boundaries.

$$\hat{y} = \text{softmax}(W \cdot f(\phi_{\text{enc}}(x); \theta) + b). (5)$$

The complete trainable parameter set $\Theta = \{\phi_{\text{enc}}, \theta, W, b\}$ includes all three coupled components: the classical encoder weights ϕ_{enc} , quantum rotation angles θ and the linear-head weights and biases $\{W, b\}$. These are jointly and end-to-end optimized with gradient descent, where we use automatic differentiation to propagate gradients through the quantum layer so that the classical and quantum parts can co-adapt rather than tune in isolation.

3.3 Federated averaging

In federated setting, K participating sites do not share raw traffic with each other, but instead are locally holding a private dataset of size n_k and train the joint model based on its own local data. With every communication round central server reconstructs new global model by averaging locally-updated parameters, weighting the contribution of each client with respect to amount of data it holds in order to larger sites have more influence over this aggregate:

$$\Theta^{t+1} = \sum_{k=1}^K (n_k/n) \Theta_k^{t+1}, \quad n = \sum_k n_k, (6)$$

At each round every client re-initializes with the broadcasted global parameters and performs E epochs of local optimization minimizing across minibatches the multi-class cross-entropy between predicted and true labels. This is captured and returned to the server only as the resulting parameter update — never raw records themselves — so that sensitive traffic probably remains on-site throughout training.

$$L(\Theta) = -(1/B) \sum_{b=1}^B \sum_{c=1}^5 y_{b,c} \log \hat{y}_{b,c}. (7)$$

4. Proposed Framework

Contextualising the design in Figure 1, each critical-infrastructure site trains its own traffic on a hybrid quantumclassical model, and FedAvg aggregates encapsulated updates and broadcasts the new global model from the server. We generate quantum gradients via the adjoint method for efficiency when simulating; on hardware they can be calculated using the parameter-shift rule [9,19].




Figure 1. Variational quantum circuit of the hybrid classifier.

4.1 Hybrid model and circuit resources

Our classical encoder is a tiny multilayer perceptron mapping 122-dimensional preprocessed NSL-KDD feature vector to n qubit angles, the quantum layer returns n expectation values and we add a linear head who gives us five logits. The ($n=4$, $L=2$) configuration is intentionally small: it contains only 24 quantum-trainable parameters embedded in ≈ 2022 hybrid parameters, circuit depth 9 and 20 gates (Table 1). This ensures that the model remains in the optimal regime to train on currently available devices while keeping its simulation feasible [11, 22].

Table 1. Quantum-circuit resource summary for representative configurations.

Qubits	Layers	Circuit depth	Total gates	Quantum params	Hybrid params
4	1	6	12	12	2010
4	2	9	20	24	2022
4	3	14	28	36	2034
6	2	12	30	36	2072
8	2	15	40	48	2122
8	4	22	72	96	2170

4.2 Quantum-resilient federated averaging

Algorithm 1 specifies the training loop. The only cryptographic requirement is that Encapsulate/Decapsulate be instantiated with a key-encapsulation mechanism; substituting a post-quantum mechanism (e.g. a lattice-based KEM) for the classical baseline yields quantum-resilient update transport without altering the learning procedure.

Algorithm 1 Quantum-Resilient Federated Averaging (QR-FedAvg) Input: clients $k=1..K$ with local data D_k ; rounds T ; local epochs E ; lr η
Initialize global parameters Θ^0 (encoder, quantum θ , head)
for $t = 0$ to $T-1$ do
 server
 broadcasts Θ^t to all clients
 for each client k in parallel do
 $\Theta_k \leftarrow \Theta^t$
 for $e = 1$ to E do
 for minibatch $B \subseteq D_k$ do
 $g \leftarrow \nabla L(\Theta_k; B)$ // quantum grads via adjoint / parameter-shift
 $\Theta_k \leftarrow \text{Adam}(\Theta_k, g, \eta)$
 $c_k \leftarrow \text{KEM.Encapsulate}(\Theta_k)$ // post-quantum-ready secure channel
 send (c_k, n_k) to server
 server: $\Theta_k \leftarrow \text{KEM.Decapsulate}(c_k)$
 $\Theta^{t+1} \leftarrow \sum_k (n_k / n) \Theta_k$ // FedAvg aggregation return Θ^T

4.3 Dataset and preprocessing

We are using the NSL-KDD standard train (KDDTrain+, 125973 records) and test (KDDTest+, 22544 records) partitions[10]. There are 41 fields total, of which three categorical fields (protocol_type, service, flag) combined with one-hot encoding for the union of train and test vocabularies yielding a 122-dimensional numeric vector. All zero-variance features are dropped and all features are min-max scaled into the field range $[0,1]$ to comply with the angle-embedding domain. Attack labels are mapped to five classes:—Normal, DoS, Probe, R2L, and U2R (Table II). The well-known imbalance (Figure 2) is the main challenge for macro-averaged evaluation.

Table 2. NSL-KDD class composition of the standard train and test partitions.

Class	Train count	Test count	Role
Normal	67,343	9,711	benign
DoS	45,927	7,458	frequent
Probe	11,656	2,421	moderate
R2L	995	2,754	rare
U2R	52	200	rare
Total	125,973	22,544	—




Figure 2. Per-class record counts for the train and test splits, illustrating the dominance of Normal and DoS traffic and the scarcity of R2L and U2R attacks.

5. Experimental Setup

For balanced class-aware cross entropy, the Adam optimizer [21] is used for training. By default, federated configurations always refer to configuration with $K=3$ clients, circuit in production — ($n=4$, $L=2$) and FedAvg communication rounds with validation macro-F1 early stopping. Clients exist in separate independent partitions which mimic different infrastructure sites, and the heterogeneity study discussed below controls partition with a Dirichlet concentration parameter α (smaller α = more non-IID). We repeat all configurations across random seeds and report mean $\pm$ standard deviation. Lightning controls the simulation of quantum circuits. qubit backend using adjoint differentiation. Metric name Metric Accuracy, Macro/Wighted Precision, Recall and F1, Matthews Correlation Coefficient (MCC); One-vs-Rest ROC-AUC (Macro/Micro) A non-IID client partition of the order is shown in Fig 3.




Figure 3. Stacked per-client class distribution under a non-IID partition, showing the

label skew each site experiences.


6. Results and Discussion

6.1 Comparative evaluation

The six methods, as shown in Table 3. The MLP, hybrid QC and support-vector variants have indiscernibly similar raw accuracy ($\approx 77.6\text{--}78.0\%$), with Random Forest lowest at 73.7%. Differences in macro-F1 which gives equal weight to classifying the rare classes are informative. The centralized hybrid QC model obtains the best macro-F1 (0.6285), and proposed federated hybrid QC keeps 0.6200—above both of the centralised (0.5975) and federated (0.5949) MLP baselines, as well as Random Forest (0.4734). The centralized and federated case tells us that the macro-F1 gap (0.6285 $\rightarrow$ 0.6200) lets us quantify the small utility cost of preserving data locality. To summarize, the quantum feature map provides only what accuracy cannot see: a solution to the minority-class discrimination.

Table 3. Comparative performance on NSL-KDD (five-class, test partition; mean $\pm$ std across seeds). Best macro-F1 in each block in bold; the proposed method is highlighted.

Method	Accuracy	F1-macro	F1-weighted	MCC	AUC-macro
Centralized Classical MLP	0.7779 $\pm$ 0.013	0.5975 $\pm$ 0.011	0.7576	0.6733	0.9198
Centralized Hybrid QC	0.7758 $\pm$ 0.010	0.6285$\pm$0.037	0.7549	0.6683	0.9152
Federated Classical MLP	0.7797 $\pm$ 0.007	0.5949 $\pm$ 0.045	0.7601	0.6711	0.9210
Federated Hybrid QC (Ours)	0.7765 $\pm$ 0.013	0.6200$\pm$0.030	0.7561	0.6686	0.9062
Random Forest	0.7371 $\pm$ 0.0	0.4734 $\pm$ 0.0	0.6893	0.6213	0.9234



	01	04			
SVM (RBF)	0.7757	0.6149	0.7591	0.6701	0.8928

Figure 4. Method comparison across centralized/federated and classical/hybrid-quantum settings. Hybrid quantum models lead on macro-F1 while matching classical accuracy.

6.2 Convergence and per-client behavior

As seen in Figure 5, the federated hybrid model converges in few training rounds, following classical-FL curve evolution on global accuracy while building up its rare-class advantage in loss. Local label distributions, albeit unbalanced per client, are confirmed as positive aggregators when evaluated on a per-client basis (Figure 8 panel communication).




Figure 5. FedAvg convergence: global test accuracy and loss versus communication round for the proposed hybrid model and the classical-FL baseline.

6.3 Per-class analysis

The remaining error is localized in the normalized confusion matrix (Figure 6) and per-class F1 (Figure 7). Normal, DoS and Probe are reliably recognised; R2L and U2R remain the hardest classes—an inherent difficulty for NSL-KDD since many R2L/U2R attack types only occur in the test partition. Indeed, the macro-F1 gain of the hybrid model over both MLP and Random-Forest baselines is accounted for by quantifiable improvements in precisely these minority classes corresponding to what we see in class-separation (Figure 9) by quantum feature embedding.




Figure 6. Row-normalized five-class confusion matrix for the proposed federated hybrid model.




Figure 7. Per-class F1 scores, highlighting the rare R2L and U2R categories that dominate macro-averaged performance.




Figure 8. One-vs-rest ROC curves with per-class and macro/micro AUC; the proposed model attains a micro-averaged AUC of 0.923.





Figure 9. t-SNE projection of the post-circuit quantum feature embeddings, coloured by class, showing emergent separation of the majority classes and partial structure among rare classes.

6.4 Ablation studies

Circuit width. Increasing qubit count from four to eight raises macro-F1 from 0.6200 to 0.6364 (Table 4) but roughly triples simulated training cost; six qubits underperforms four, indicating a non-monotonic width–difficulty interaction rather than a clean scaling law at this data scale.

Table 4. Ablation over qubit count (L=2; mean $\pm$ std).

Qubits	Accuracy	F1-macro	AUC-macro
4	0.7765 $\pm$ 0.013	0.6200 $\pm$ 0.030	0.9062
6	0.7756 $\pm$ 0.007	0.6061 $\pm$ 0.012	0.9246
8	0.7892 $\pm$ 0.007	0.6364$\pm$0.013	0.9334

**Figure 10.** Macro-F1 versus qubit count (mean $\pm$ std across seeds).

Circuit depth. Depth exhibits the barren-plateau signature [22]: macro-F1 is highest at L=1 (0.6381), essentially flat at L=2–3, and degrades at L=4 (0.5845) as gradients vanish and optimization stalls (Table 5). Shallow circuits are therefore both cheaper and more accurate here.

Table 5. Ablation over variational circuit depth (n=4; mean $\pm$ std).

Layers	Accuracy	F1-macro	AUC-macro
1	0.7825 $\pm$ 0.006	0.6381$\pm$0.011	0.9106
2	0.7765 $\pm$ 0.013	0.6200 $\pm$ 0.030	0.9062
3	0.7892 $\pm$ 0.006	0.6259 $\pm$ 0.025	0.9274
4	0.7824 $\pm$ 0.011	0.5845 $\pm$ 0.013	0.9214




Figure 11. Macro-F1 versus variational circuit depth; performance peaks at shallow depth and degrades as gradients flatten.

Client population. The expected decrease in macro-F1 ($0.6200 \rightarrow 0.5731$) when going from 3 to 5 clients per site is due to the increasingly imbalanced nature of the data seen by each site, a natural federated statistical-efficiency trade-off.

Data heterogeneity. Table 6. Sweeping the Dirichlet concentration α . Macro-F1 0.5966: Near-IID ($\alpha=100$), Performance collapses to 0.4074 for extreme heterogeneity ($\alpha=0.1$) with large variance (accuracy 0.581 ± 0.141): A quantification of realistic label skew cost, motivating heterogeneity-aware aggregation in future work

Table 6. IID-to-non-IID sweep via the Dirichlet concentration α (smaller = more heterogeneous; mean $\pm$ std).

α (heterogeneity)	Accuracy	F1-macro	MCC
0.1 (severe)	0.5810 ± 0.141	0.4074 ± 0.061	0.3664
0.5	0.7367 ± 0.013	0.5048 ± 0.016	0.6111
1.0	0.7677 ± 0.022	0.5913 ± 0.046	0.6570
100 ($\approx$ IID)	0.7687 ± 0.008	0.5966 ± 0.014	0.6555




Figure 12. Effect of data heterogeneity: performance degrades sharply as the Dirichlet α decreases toward severe non-IID partitions.

6.5 Computational cost and secure aggregation

Table 7: We Report Cost Transparently. As the variational circuit is classically simulated, federated hybrid training becomes the most expensive approach (≈ 6.2 hours of wall-clock simulation) on a par with kernel SVM (≈ 6.4 hours), but

orders of magnitude larger than classical MLPs. This overhead is an artefact of statevector simulation, not of quantum execution: inference has negligible cost and on quantum hardware the cost of a circuit scales with its shallow depth—as opposed to exponentially in qubit count. The classic key-encapsulation handshake for the secure-aggregation channel only adds ≈ 2 ms per round with an ≈ 11 KB payload, which is negligible relative to training, so substituting a lattice-based post-quantum KEM in place of quantum-resilient transport incurs no concrete penalty on the learning pipeline.

Table 7. Training wall-clock time by method (seconds). Quantum-simulated methods dominate cost; inference is negligible for all methods.

Method	Train time (s)	Relative to MLP
Random Forest	24.4	0.5×
Centralized Classical MLP	47.6	1×
Federated Classical MLP	79.7	1.7×
Centralized Hybrid QC	459.9	9.7×
SVM (RBF)	23,163	486×
Federated Hybrid QC (Ours)	22,377	470×




Figure 13. Training and inference wall-clock time, quantum-simulated versus classical—an honest overhead comparison.

7. Limitations

Three limitations temper the conclusions. Finally, we obtain our results under small-circuit classical simulation; shot statistics and hardware noise after circuit execution are not modelled, and the overhead is observed in simulation rather than simulated execution. Second, though a reproducible benchmark and NSL-KDD is much less than a live defense-network capture (the framing critical-infrastructure makes for a threat-model motivation instead of an operational deployment with the challenge to transfer to real telemetry). Third, the deepest circuits show barren-plateau

behaviour [22], limiting the depth at which the current ansatz is useful and indicating structured or problem-informed ansatzes as a direction for future work. Post-quantum transport is architected and overhead assessed at the classical-KEM baseline, an end-to-end test against a standard lattice-based KEM makes for a logical next step [4].

8. Conclusion and Future Work

Then, we illustrated a post-quantum secure quantum federated intrusion-detection framework that combines both shallow variational quantum classifier with a federated averaging and a post-quantum-ready secure-aggregation channel. For NSL-KDD, the federated hybrid model achieves classical accuracy while improving macro-F1, with its advantage concentrated on the rare R2L and U2R attacks prominent under balanced evaluation. The results show that a short-lived quantum model can serve as a privacy-preserving, rare-class-aware part of an intrusion detection system while the rest of the protocol is robust to future quantum adversaries. Future work will implement a standardized post-quantum KEM from end-to-end, develop hardware-noise models along with problem-informed ansatz to counter barren plateaus, introduce heterogeneity-aware aggregation, and test on real critical-infrastructure telemetry.

References

- [1] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM Journal on Computing*, vol. 26, no. 5, pp. 1484–1509, 1997.
- [2] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proc. 28th ACM Symposium on Theory of Computing (STOC)*, pp. 212–219, 1996.
- [3] L. Chen et al., "Report on post-quantum cryptography," National Institute of Standards and Technology, NISTIR 8105, 2016.
- [4] D. J. Bernstein and T. Lange, "Post-quantum cryptography," *Nature*, vol. 549, pp. 188–194, 2017.
- [5] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agüera y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. AISTATS*, pp. 1273–1282, 2017.
- [6] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," *arXiv:1610.05492*, 2016.
- [7] J. Biamonte, P. Wittek, N. Pancotti, N. Wiebe, P. Rebentrost, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, pp. 195–202, 2017.
- [8] V. Havlíček et al., "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, pp. 209–212, 2019.
- [9] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, "Parameterized quantum circuits as machine learning models," *Quantum Science and Technology*, vol. 4, no. 4, 043001, 2019.
- [10] M. Tavallaei, E. Bagheri, W. Lu, and A. A. Ghorbani, "A detailed analysis of the KDD CUP 99 data set," in *Proc. IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA)*, pp. 1–6, 2009.

- [11] J. Preskill, "Quantum computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, 2018.
- [12] V. Bergholm et al., "PennyLane: Automatic differentiation of hybrid quantum-classical computations," *arXiv:1811.04968*, 2018.
- [13] A. L. Buczak and E. Guven, "A survey of data mining and machine learning methods for cyber security intrusion detection," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 2, pp. 1153–1176, 2016.
- [14] C.-F. Tsai, Y.-F. Hsu, C.-Y. Lin, and W.-Y. Lin, "Intrusion detection by machine learning: A review," *Expert Systems with Applications*, vol. 36, no. 10, pp. 11994–12000, 2009.
- [15] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology*, vol. 10, no. 2, pp. 1–19, 2019.
- [16] K. Bonawitz et al., "Practical secure aggregation for privacy-preserving machine learning," in *Proc. ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pp. 1175–1191, 2017.
- [17] M. Schuld and N. Killoran, "Quantum machine learning in feature Hilbert spaces," *Physical Review Letters*, vol. 122, no. 4, 040504, 2019.
- [18] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, "Quantum circuit learning," *Physical Review A*, vol. 98, no. 3, 032309, 2018.
- [19] M. Schuld, I. Sinayskiy, and F. Petruccione, "An introduction to quantum machine learning," *Contemporary Physics*, vol. 56, no. 2, pp. 172–185, 2015.
- [20] E. Farhi and H. Neven, "Classification with quantum neural networks on near term processors," *arXiv:1802.06002*, 2018.
- [21] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. International Conference on Learning Representations (ICLR)*, 2015.
- [22] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, "Barren plateaus in quantum neural network training landscapes," *Nature Communications*, vol. 9, 4812, 2018.
- [23] S. Y.-C. Chen, C.-H. H. Yang, J. Qi, P.-Y. Chen, X. Ma, and H.-S. Goan, "Variational quantum circuits for deep reinforcement learning," *IEEE Access*, vol. 8, pp. 141007–141024, 2020.