Evaluation of Quantum Machine Learning Models for Network Anomaly Detection

Swathi Chandrasekhar and Shiva Raj Pokhrel
IoT & SE Research Lab, School of IT, Deakin University, Geelong, Australia
Email: swathi.chandrasekhar@deakin.edu.au, shiva.pokhrel@deakin.edu.au

***Abstract*—The proliferation of Internet-connected devices has intensified the risk of large-scale network intrusions, necessitating more advanced detection methodologies for the post quantum era. This work presents, to the best of our knowledge, a comprehensive investigation of quantum machine learning (QML) for anomaly detection on the BoT-IoT dataset, a widely adopted benchmark for IoT security. Leveraging feature-engineered representations, we implement and evaluate four representative quantum classifiers: Quantum Support Vector Classification (QSVC), Quantum Neural Networks (QNN), and Variational Quantum Classifiers (VQC).**

***Index Terms*—Quantum Machine Learning (QML), Anomaly Detection.**

I. Introduction

The proliferation of digital technologies and the Internet of Things (IoT) has generated vast, high-dimensional datasets in domains such as consumer electronics, finance, and critical infrastructure. As these systems become more interconnected, the challenge of detecting anomalies—ranging from device failures to financial fraud—has grown increasingly urgent. While classical machine learning (ML) has achieved notable success in anomaly detection, it often struggles with scalability, robustness in high-dimensional spaces, and generalization to novel threats [1]. Quantum computing, leveraging superposition and entanglement, offers a fundamentally new computational paradigm that can surpass classical limits [2]–[4]. Within this paradigm, quantum machine learning (QML) [5], [6] has shown promise for complex, high-dimensional, and noisy data [1], [3], [7], with emerging evidence suggesting its potential to efficiently identify rare and subtle anomalies.

Hybrid quantum-classical models have demonstrated early success in IoT and consumer electronics by capturing intricate, non-linear correlations [1]. In finance, quantum kernel methods have outperformed classical RBF kernels as feature dimensionality increases, especially in unsupervised settings [3]. Variational quantum circuits and quantum support vector machines further mitigate overfitting and enhance generalization in heterogeneous and noisy data [3].

Despite these advances, practical deployment of QML-based anomaly detection faces challenges, including the computational cost of large quantum kernel matrices and hardware limitations such as qubit count and gate fidelity [7]–[9]. Ongoing research into efficient kernel evaluation, sparse matrix techniques, and hardware-specific optimizations is gradually improving feasibility [3].

In this work, we systematically investigate the potential and limitations of QML for anomaly detection in real-world domains. By benchmarking quantum and classical methods on representative datasets, we delineate the conditions for quantum advantage and highlight key algorithmic and hardware bottlenecks. Our findings offer new insights into the expressivity, robustness, and scalability of quantum-enhanced anomaly detection, guiding future progress toward practical quantum advantage in security-critical, data-intensive environments.

A. Key Contributions

- We propose a unified QML framework for anomaly detection, instantiated and evaluated on the BoT-IoT benchmark dataset [10], providing the first systematic study of quantum enhanced methods in this security critical domain.
- We rigorously assess the influence of dimensionality reduction techniques including Behavioral and Statistical Analysis (BSA), principal component analysis (PCA), and t-distributed stochastic neighbor embedding (t-SNE) on the performance and robustness of quantum classifiers.
- We implement and benchmark four state-of-the-art QML models Quantum Support Vector Classifier (QSVC), Variational Quantum Classifier (VQC), Quantum Neural Network (QNN) for anomaly detection in IoT network traffic.
- We provide a comparative performance evaluation of these quantum approaches against classical baselines, identifying the conditions under which QML models yield advantages, and analyze their scalability, generalization, and deployment feasibility for real-world IoT security applications.

II. Methodology

We propose a novel framework for implementing and evaluating quantum machine learning (QML) algorithms for anomaly detection in IoT network traffic, using the BoT-IoT dataset. The framework comprises two main stages: (i) dimensionality reduction and feature selection, and (ii) quantum-based classification.

Feature Selection and Dimensionality Reduction: The dataset is first preprocessed to identify the most relevant features. Feature importance is quantified using behavioral and statistical metrics, including point-biserial correlation, chi-square tests, and ANOVA F-tests, with aggregated scores used

979-8-3315-7896-1/26/$31.00 ©2026 IEEE

to rank and select top features. For linear dimensionality reduction, we apply Principal Component Analysis (PCA) [11], which projects data onto a lower-dimensional subspace by solving the eigenvalue problem: $\mathbf{C}\mathbf{v}_i = \lambda_i \mathbf{v}_i$, where $\mathbf{C}$ is the covariance matrix, $\mathbf{v}_i$ are eigenvectors, and λ_i are eigenvalues.

To capture non-linear structure, we employ t-distributed Stochastic Neighbor Embedding (t-SNE), which preserves local similarities by minimizing the Kullback–Leibler divergence between high- and low-dimensional probability distributions: $\mathrm{KL}(P\|Q) = \sum_{i\neq j} p_{ij} \log \frac{p_{ij}}{q_{ij}}$, where p_{ij} and q_{ij} denote pairwise similarities in the original and embedded spaces, respectively.

Quantum Classification: The reduced feature sets are then used as input to several quantum classifiers. Each model exploits different quantum mechanisms to enhance discrimination power in high-dimensional and noisy regimes.

A. Quantum Support Vector Classifier (QSVC)

QSVC extends classical support vector machines by employing quantum kernels that map data into a high-dimensional Hilbert space [5], [8], [12]. The kernel function between two input vectors x and x' is expressed as $K(x, x'; \theta) = |\langle \psi(x, \theta) \mid \psi(x', \theta)\rangle|^2$, where $\psi(x, \theta)$ is the quantum state generated by a parameterized feature map. This formulation enables QSVC to capture complex correlations inaccessible to classical kernels, particularly in high-dimensional IoT traffic data.

Algorithm 1 Quantum Support Vector Classifier (QSVC)

```
procedure QSVC_TRAIN($\{(\vec{x}_i, y_i)\}_{i=1}^{N}$, feature_map, optimizer)
    Apply dimensionality reduction to $\vec{x}_i$
    Normalize features
    Initialize parameters $\theta$
    for each pair $(\vec{x}_i, \vec{x}_j)$ do
        Prepare quantum states $|\psi(\vec{x}_i, \theta)\rangle$, $|\psi(\vec{x}_j, \theta)\rangle$
    end for
    Optimize $\theta$ to minimize SVC loss
    Train QSVC with optimized kernel $K$ and labels $y_i$
    return trained QSVC model
end procedure
```

B. Variational Quantum Classifier (VQC)

VQC is a hybrid quantum-classical algorithm that uses a parameterized quantum circuit and a classical optimizer to learn a decision boundary [4], [12]. The model minimizes:

$$\min_{\theta} \mathcal{L}(\theta) = \sum_{i=1}^{N} \ell\left(f_\theta(x_i), y_i\right)$$

where $f_\theta(x_i)$ is the predicted probability and ℓ is a loss function.

Algorithm 2 Variational Quantum Classifier (VQC)

```
procedure VQC_TRAIN($\{(\vec{x}_i, y_i)\}_{i=1}^{N}$, feature_map, ansatz, optimizer)
    Apply dimensionality reduction to $\vec{x}_i$
    Normalize features and Initialize parameters $\theta$
    for each training iteration do
        for each $\vec{x}_i$ do
            Prepare quantum state with feature_map,ansatz
            Compute predicted probabilities $f_\theta(\vec{x}_i)$
        end for
        Compute loss $\mathcal{L}(\theta)$
    end for
    return trained VQC model
end procedure
```

C. Quantum Neural Network (QNN)

QNNs are parameterized quantum circuits that approximate complex functions, analogous to classical neural networks [2], [4], [8], [12]. The QNN output is the expectation value:

$$f_\theta(x) = \langle 0| U^\dagger(x, \theta) \hat{O} U(x, \theta) |0\rangle$$

where $U(x, \theta)$ is the parameterized circuit and $\hat{O}$ is the measurement observable.

Algorithm 3 Quantum Neural Network (QNN)

```
procedure QNN_TRAIN($\{(\vec{x}_i, y_i)\}_{i=1}^{N}$, feature_map, ansatz, optimizer)
    Apply dimensionality reduction to $\vec{x}_i$
    Normalize features and Initialize parameters $\theta$
    for each training iteration do
        for each $\vec{x}_i$ do
            Prepare quantum state with feature_map,ansatz
            Compute output $f_\theta(\vec{x}_i)$
        end for
        Compute loss $\mathcal{L}(\theta)$
    end for
    return trained QNN model
end procedure
```

To instantiate our framework, we use the BoT-IoT dataset preprocessed following [10]. From the original traffic traces, we retain flow-level statistical and behavioral features (e.g., packet counts, byte rates, inter-arrival times, and flag-based indicators), yielding 20 features after BSA-based ranking. The resulting dataset is randomly partitioned into 70% training, 15% validation, and 15% test sets, stratified by attack/benign labels to preserve class proportions. All features are z-score normalized prior to dimensionality reduction (PCA or t-SNE) and subsequent quantum encoding. Quantum circuits are simulated using Qiskit Aer on a Mac-OS, ensuring that performance comparisons reflect realistic near-term execution costs for QML-based anomaly detection.

TABLE I: Performance of Classical ML and QML Algorithms for Anomaly Detection on the BoT-IoT Dataset Using Different Dimensionality Reduction Techniques.

Algorithm	**Accuracy**	**Precision**	**F1-score**	**Recall**	**Time (s)**	**Dimensionality reduction Technique**
SVC (classical ML)	0.87	0.71	0.83	0.88	268	BSA
QSVC	0.85	0.72	0.85	0.78	736	BSA
QNN	0.85	0.85	0.72	0.92	164	BSA
VQC	0.95	0.96	0.95	0.95	266	BSA
SVC (classical ML)	0.97	0.91	0.93	0.88	141	PCA
QSVC	0.94	0.97	0.95	0.94	2152	PCA
QNN	0.74	0.71	0.83	0.79	472	PCA
VQC	0.52	0.64	0.49	0.58	2230	PCA
SVC (classical ML)	0.77	0.81	0.76	0.78	285	t-SNE
QSVC	0.89	0.80	0.91	0.94	872	t-SNE
QNN	0.62	0.67	0.62	0.73	130	t-SNE
VQC	0.74	0.73	0.74	0.85	132	t-SNE

Abbreviations: QSVC = Quantum Support Vector Classifier; QNN = Quantum Neural Network; VQC = Variational Quantum Classifier; BSA = Behavioral and Statistical Analysis; PCA = Principal Component Analysis; t-SNE = t-distributed Stochastic Neighbor Embedding.

III. Results and Experimentation

Table I presents the performance of three QML algorithms (QSVC, QNN, and VQC) and classical SVC on the BoT-IoT dataset with three dimensionality reduction techniques: BSA, PCA, and t-SNE. Metrics include accuracy, precision, recall, F1-score, and computational time. VQC with PCA achieves the highest accuracy and F1-score, while QSVC offers strong precision with lower computational cost. Training loss visualizations (Figures 1–4) indicate that VQC converges faster and more stably across dimensionality reduction methods, suggesting its promise for scalable IoT anomaly detection.

Classical SVC generally matches or outperforms QML models in both accuracy and efficiency. Nonetheless, QML algorithms already show competitive results, and as quantum hardware and algorithms mature, they are expected to surpass classical methods, particularly for complex, high-dimensional data.

A. Behavioral and Statistical Analysis (BSA)

Figure 1 illustrates the convergence trajectories of QNN, VQC, and QSVC using BSA features. VQC exhibits rapid loss minimization, approaching near-zero within 30 iterations, consistent with its top performance (accuracy = 0.95, F1 = 0.95). QNN converges more gradually over 50 iterations, reflecting its robust recall (0.92), while QSVC displays pronounced oscillations between 0.6–1.0, indicating unstable training despite moderate accuracy (0.85). These results suggest that BSA-derived representations are more effectively leveraged by parameterized quantum circuits (VQC, QNN) than by kernel-based methods.

B. Principal Component Analysis (PCA)

Figure 2 shows the convergence behavior of QSVC, QNN, and VQC on PCA-reduced features. QSVC exhibits smooth and consistent convergence, achieving the highest accuracy (0.94) and F1-score (0.95), highlighting its suitability for linearly compressed representations. QNN converges slowly, suggesting limited expressivity on PCA-transformed features, while VQC experiences sharp loss spikes and unstable training, resulting in poor accuracy (0.52). Notably, both QSVC and VQC incur high computational costs (>2000 s), indicating that while PCA can enhance QSVC performance, it also substantially increases runtime.

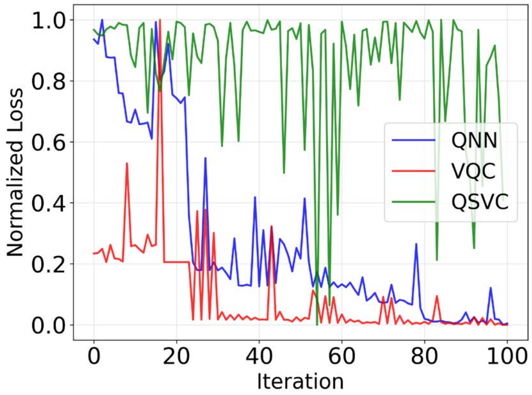

Fig. 1: Normalized training loss across iterations for QNN, VQC, and QSVC using BSA.

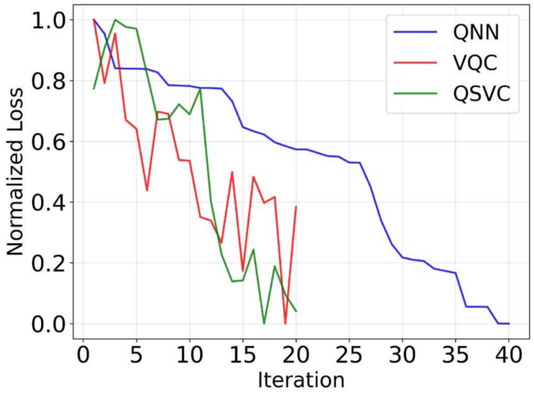

Fig. 2: Normalized training loss across iterations for QNN, VQC, and QSVC using PCA.

C. t-distributed Stochastic Neighbor Embedding (t-SNE)

Figure 3 illustrates the training dynamics of QSVC, QNN, and VQC on t-SNE features. QSVC demonstrates nearly linear, monotonic convergence from loss = 1.0 to 0.0, consistent with its strong recall (0.94) and robust F1-score (0.91). QNN and VQC exhibit oscillatory convergence, stabilizing after 100 iterations, with VQC ultimately achieving lower final loss and higher accuracy (0.74 vs. 0.62). Notably, t-SNE significantly improves computational efficiency, reducing runtime for QNN (130 s) and VQC (132 s) compared to PCA-based training, highlighting its effectiveness for fast, high-dimensional anomaly detection.

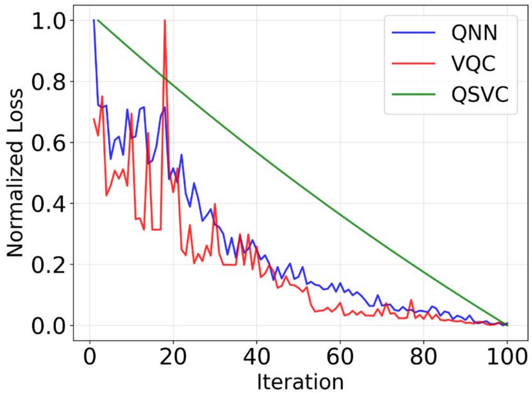

Fig. 3: Normalized training loss across iterations for QNN, VQC, and QSVC using t-SNE.

D. Loss Distribution Across Models and Dimensionality reduction Techniques

Figure 4 presents violin plots of normalized loss distributions for all algorithm–Dimensionality reduction combinations. VQC with BSA exhibits the narrowest, lowest-centered distribution, confirming both stable training and superior predictive performance. QSVC on PCA, despite high accuracy, shows the widest spread, reflecting substantial variability across iterations. t-SNE yields more balanced distributions for all algorithms, with QSVC maintaining consistently lower spread than QNN and VQC. Overall, these results indicate that BSA optimally stabilizes VQC, PCA favors QSVC, and t-SNE offers a compromise between computational efficiency and moderate reliability.

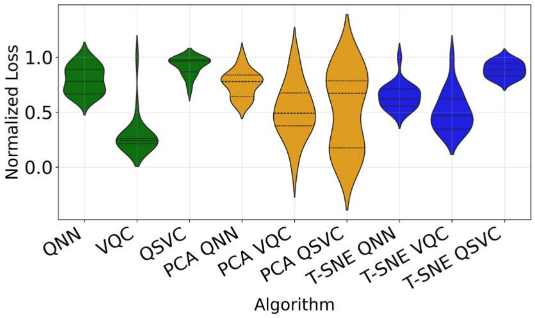

Fig. 4: Distribution of normalized losses for QNN, VQC, and QSVC under different Dimensionality reduction techniques (BSA, PCA, t-SNE).

E. Comparative Analysis

Integrating the quantitative outcomes (Table I) with convergence and distribution analyses (Figures 1–4), three conclusions emerge:

- **VQC + BSA** achieves the most favorable trade-off between stability, accuracy, and computational time.
- **QSVC + PCA** delivers the highest accuracy but incurs significant computational costs and variance in loss.
- **QSVC + t-SNE** provides strong recall with smooth convergence and efficient runtimes, making it attractive for resource-constrained deployments.

IV. Conclusion

This paper presented a comparative evaluation of QML algorithms for anomaly detection on the BoT-IoT dataset using different dimensionality reduction techniques. Our results show that VQC combined with BSA achieved the best overall trade-off between accuracy, stability, and runtime, while QSVC with PCA delivered the highest accuracy at the cost of significant computational time. QSVC with t-SNE provided efficient convergence with strong recall, making it suitable for resource-limited IoT environments. In practice, t-SNE driven feature compression reduces circuit depth and evaluation time, which is particularly beneficial when QML inference is offloaded from cloud backends to edge gateways or lightweight security appliances, thereby aligning our design with the stringent latency and resource constraints typical of IoT deployments. QNN demonstrated robustness in recall but exhibited greater variability across reductions.

References

[1] S. Bhowmik and H. Thapliyal, “Quantum machine learning for anomaly detection in consumer electronics,” in *2024 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*. IEEE, 2024, pp. 544–550.

[2] N. Liu and P. Rebentrost, “Quantum machine learning for quantum anomaly detection,” *Physical Review A*, vol. 97, no. 4, p. 042315, 2018.

[3] O. Kyriienko and E. B. Magnusson, “Unsupervised quantum machine learning for fraud detection,” *arXiv preprint arXiv:2208.01203*, 2022.

[4] S. Chandrasekhar, S. R. Pokhrel, and N. Singh, “Adapting quantum machine learning for energy dissociation of bonds,” *ChemRxiv*, 2025.

[5] S. Corli, L. Moro, D. Dragoni, M. Dispenza, and E. Prati, “Quantum machine learning algorithms for anomaly detection: A review,” *Future Generation Computer Systems*, vol. 166, p. 107632, 2025.

[6] S. Kumari, S. R. Pokhrel, S. Chandrasekhar, N. Singh, H. S. Dutta, A. Anwar, S. Rajasegarar, and R. Doss, “Modeling wavelet transformed quantum support vector for network intrusion detection,” 2025. [Online]. Available: https://arxiv.org/abs/2512.01365

[7] 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.

[8] N. Singh and S. R. Pokhrel, “Modeling feature maps for quantum machine learning,” *arXiv preprint arXiv:2501.08205*, 2025.

[9] J. Preskill, “Quantum computing in the nisq era and beyond,” *Quantum*, vol. 2, p. 79, 2018.

[10] J. M. Peterson, J. L. Leevy, and T. M. Khoshgoftaar, “A review and analysis of the bot-iot dataset,” in *2021 IEEE International Conference on Service-Oriented System Engineering (SOSE)*. IEEE, 2021, pp. 20–27.

[11] N. Singh and S. R. Pokhrel, “Modeling quantum machine learning for genomic data analysis,” *arXiv preprint arXiv:2501.08193*, 2025.

[12] S. Chandrasekhar, S. R. Pokhrel, S. Kumari, and N. Singh, “Modeling quantum autoencoder trainable kernel for iot anomaly detection,” 2025. [Online]. Available: https://arxiv.org/abs/2511.21932