IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received 29 April 2026, accepted 11 May 2026, date of publication 14 May 2026, date of current version 26 May 2026.

Digital Object Identifier 10.1109/ACCESS.2026.3693627

RESEARCH ARTICLE

Hybrid Quantum Machine Learning for Intrusion Detection: A Comparative Study of QNN and QSVM Models

RAID ANIS KERKATOU[1,2], HACENE BELHADEF[3], AICHA EUTAMENE[4], AND SVETLANA PETROVA STEFANOVA[5]

[1]University of Mujahid Abdelhafid Boussouf, Mila 43000, Algeria
[2]MISC Laboratory, University of Constantine 2-Abdelhamid Mehri, Constantine 25016, Algeria
[3]NTIC Faculty, University of Constantine 2-Abdelhamid Mehri, Constantine 25016, Algeria
[4]University Fréres Mentouri Constantine 1, Constantine 25000, Algeria
[5]Faculty of Electrical Engineering, Electronics and Automatics, University of Ruse "Angel Kanchev," 7017 Ruse, Bulgaria

Corresponding author: Raid Anis Kerkatou (kerkatouraidanis@centre-univ-mila.dz)

This work was supported by the Scientific Research Fund of the University of Ruse "Angel Kanchev."

ABSTRACT Intrusion Detection Systems (IDSs) are a cornerstone of modern cybersecurity; however, their performance is increasingly constrained by high-dimensional network traffic and the scalability limitations of classical machine learning approaches. Quantum Machine Learning (QML) has recently emerged as a promising paradigm by leveraging quantum feature spaces and hybrid quantum–classical optimization. In this work, we present a unified and reproducible benchmarking framework for QML-based intrusion detection under realistic Noisy Intermediate-Scale Quantum (NISQ) constraints. The proposed framework integrates hybrid Quantum Neural Networks (QNNs) and Quantum Support Vector Machines (QSVMs) within a consistent experimental pipeline, enabling a systematic comparison of quantum models, embedding strategies, and dimensionality reduction techniques. Experiments are conducted on the NSL-KDD benchmark dataset for both binary (normal versus attack) and multiclass intrusion detection tasks. Multiple quantum embedding strategies, including angle, amplitude, and Instantaneous Quantum Polynomial (IQP) embeddings, are evaluated in conjunction with PCA-based dimensionality reduction to analyze trade-offs between expressivity, circuit complexity, and computational cost. The proposed models are benchmarked against strong classical baselines under identical preprocessing and evaluation conditions. The results demonstrate that QNN- and QSVM-based IDS models achieve competitive and, in several scenarios, superior performance compared to classical approaches, particularly in multiclass settings characterized by nonlinear and imbalanced data distributions. In addition, we provide an in-depth analysis of scalability limitations, training complexity, and NISQ feasibility, highlighting practical considerations for real-world deployment. To ensure full reproducibility, all implementation details, datasets, and experimental configurations are publicly available at: https://github.com/raidanis/QML-IDS-NSLKDD. These findings position hybrid QML as a promising complementary paradigm for next-generation intrusion detection systems, particularly in low-dimensional and high-complexity feature spaces.

INDEX TERMS Cybersecurity, intrusion detection systems, quantum machine learning, quantum neural networks, quantum support vector machines, network security.

I. INTRODUCTION

Intrusion and anomaly detection play a critical role in protecting modern digital infrastructures, ranging from enterprise

The associate editor coordinating the review of this manuscript and approving it for publication was Ayman El-Baz.

networks to large-scale critical systems. Classical machine learning (ML) techniques, including Principal Component Analysis (PCA), autoencoders, and kernel-based classifiers, have been widely adopted to detect malicious behaviors in network traffic data. While effective in controlled environments, these approaches often struggle with high-dimensional, noisy, and highly imbalanced tabular data, limiting their ability to capture subtle or evolving attack patterns and increasing computational overhead in large-scale deployments.

Recent advances in Quantum Machine Learning (QML) offer a promising alternative by leveraging quantum-mechanical principles such as superposition and entanglement to construct highly expressive feature representations. Through quantum data embedding and quantum kernel mappings, QML models project classical data into high-dimensional Hilbert spaces, enabling the learning of complex nonlinear decision boundaries. Hybrid quantum–classical models based on variational quantum circuits provide a practical pathway for Noisy Intermediate-Scale Quantum (NISQ) devices by combining quantum feature extraction with classical optimization.

A growing body of literature has explored QML for anomaly and intrusion detection. Corli et al. [1] provide a comprehensive survey of QML-based anomaly detection methods. Oh and Park [2] introduced the Quantum Support Vector Data Description (QSVDD), demonstrating the effectiveness of quantum kernel-based anomaly detection. Rasyidi et al. [3] proposed hybrid quantum-classical autoencoders (HQC-AE) for intrusion detection, showing improved generalization under certain configurations. Rath and Date [4] emphasized the critical role of quantum embedding strategies in enhancing classification performance.

Despite these advances, several key limitations remain. First, most prior studies focus on a single quantum model, without providing a unified comparative evaluation across multiple QML paradigms. Second, the interaction between classical dimensionality reduction techniques and quantum feature embeddings remains insufficiently explored. Third, reproducibility is often limited due to incomplete disclosure of implementation details and experimental configurations. Finally, the scalability and practical feasibility of QML-based intrusion detection systems under NISQ constraints remain open challenges.

To address these gaps, this work proposes a unified and reproducible hybrid Quantum Machine Learning framework for intrusion detection. The framework integrates Quantum Neural Networks (QNNs) and Quantum Support Vector Machines (QSVMs) within a consistent experimental pipeline, enabling systematic comparison across models, embedding strategies, and dimensionality reduction configurations.

PCA-based dimensionality reduction is employed as a principled mechanism to map high-dimensional network traffic data into a lower-dimensional space compatible with quantum hardware constraints. Multiple quantum embedding strategies, including angle encoding, amplitude encoding, and Instantaneous Quantum Polynomial (IQP) encoding, are evaluated to analyze their impact on model expressivity, trainability, and computational cost.

The framework is evaluated on the NSL-KDD benchmark dataset for both binary and multiclass intrusion detection tasks. While widely used, this dataset has known limitations in representing modern cyber threats; therefore, its use in this work is intended to provide a controlled and reproducible benchmarking environment, and its limitations are explicitly discussed.

The main contributions of this paper are summarized as follows:

- A unified and reproducible QML-based intrusion detection framework integrating QNN and QSVM models under a consistent experimental protocol.
- A systematic comparative analysis of multiple quantum embedding strategies and PCA configurations, highlighting trade-offs between expressivity, scalability, and computational cost.
- A comprehensive benchmarking study against classical baselines, including SVM and Random Forest models, under identical preprocessing and evaluation conditions.
- An in-depth analysis of scalability, training complexity, and NISQ-era constraints, providing practical insights into the feasibility of QML-based IDS.
- Public release of all code, datasets, and experimental configurations to ensure full reproducibility of results.

The remainder of this paper is organized as follows. Section II presents the fundamentals of Quantum Machine Learning. Section III reviews related work on classical and quantum-based intrusion detection. Section IV describes the proposed hybrid quantum machine learning framework, including data preprocessing, quantum embeddings, and model architectures. Section V presents and discusses the experimental results for both binary and multiclass intrusion detection tasks. The limitations and NISQ related constraints of the proposed approach are analyzed in Section VI. Finally, Section VII concludes the paper and outlines directions for future research.

II. QUANTUM MACHINE LEARNING FUNDAMENTALS

This section briefly introduces the quantum computing concepts required to understand the proposed intrusion detection framework. The discussion is intentionally concise and focuses on principles directly relevant to quantum machine learning models.

A. QUBITS AND QUANTUM STATES

A quantum bit (qubit) is the fundamental unit of quantum information. A qubit can be represented as a normalized superposition of computational basis states

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \tag{1}$$

where $\alpha, \beta \in \mathbb{C}$ and $|\alpha|^2 + |\beta|^2 = 1$. Unlike classical bits, qubits can encode information in superposition, allowing a

system of n qubits to represent states in a 2^n-dimensional Hilbert space. This exponential representational capacity forms the basis for potential computational advantages in quantum machine learning.

B. QUANTUM GATES AND MEASUREMENT

Quantum gates are unitary operations applied to qubits and serve as the building blocks of quantum circuits. They are analogous to logical operations in classical computing. Common single qubit gates include the Pauli operators (X, Y, and Z) and the Hadamard gate H, which prepares superposition states. Multi qubit gates such as the controlled-NOT (CNOT) gate introduce entanglement between qubits. Measurement collapses a quantum state into a classical outcome, with probabilities determined according to the Born rule.

C. QUANTUM DATA EMBEDDING

In quantum machine learning, classical data must be encoded into quantum states through embedding circuits. Angle embedding maps classical features to rotation angles of parameterized quantum gates, while amplitude embedding encodes classical vectors directly into the amplitudes of quantum states. More expressive mappings, such as Instantaneous Quantum Polynomial (IQP) embeddings, exploit entanglement and higher order correlations to construct richer quantum feature spaces. These embeddings play a critical role in the performance of quantum models such as Quantum Support Vector Machines (QSVMs) and Quantum Neural Networks (QNNs).

D. VARIATIONAL QUANTUM CIRCUITS AND HYBRID LEARNING

Variational Quantum Circuits (VQCs) form the computational backbone of many near term quantum machine learning models [5]. A VQC typically consists of three components: a data embedding layer, a parameterized variational layer, and a measurement stage, as illustrated in Fig. 1.

In the embedding stage, classical input features $\mathbf{x}$ are mapped to a quantum state through a data dependent unitary transformation $U(\mathbf{x})$. This process encodes classical information into the quantum Hilbert space and defines the geometry of the resulting feature map. The choice of embedding strategy directly influences the expressivity and performance of the model.

The variational layers apply trainable unitary operations $U(\theta)$ composed of parameterized single qubit rotations and entangling gates. These layers introduce nonlinear transformations and correlations among qubits, enabling the circuit to learn complex decision boundaries. The circuit parameters are optimized using classical gradient-based or gradient free optimization algorithms within a hybrid quantum-classical training loop.

Finally, quantum measurements project the circuit output into classical observables, producing expectation values that define the model output. This measurement process connects quantum computation with classical optimization and enables practical learning within the constraints of current Noisy Intermediate Scale Quantum (NISQ) hardware.

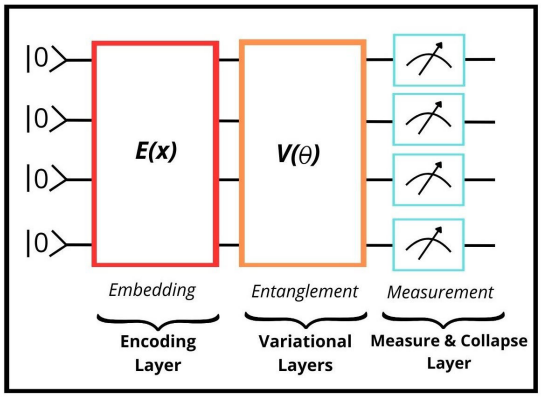

FIGURE 1. Structure of a variational quantum circuit used in this work. Classical features are encoded into quantum states via a data dependent embedding unitary $U(\mathbf{x})$, followed by parameterized variational layers $U(\theta)$ that introduce entanglement. Measurement collapses the quantum state to classical values used for hybrid optimization.

III. RELATED WORK

Intrusion detection and anomaly detection have been extensively studied using classical machine learning techniques. Traditional approaches typically rely on statistical models, shallow classifiers, or feature engineering combined with dimensionality reduction methods such as Principal Component Analysis (PCA). More recently, deep learning architectures, including autoencoders and recurrent neural networks, have been adopted to capture complex temporal and nonlinear patterns in network traffic data. Despite their success, these classical methods often suffer from scalability limitations, sensitivity to high dimensional noise, and reduced generalization when attack patterns evolve or data distributions shift.

Quantum Machine Learning (QML) has recently emerged as a promising paradigm capable of constructing highly expressive feature spaces through quantum data encoding and kernel-based learning. Several studies have explored QML for anomaly detection in generic tabular and signal processing contexts. Corli et al. [1] provide a comprehensive overview of QML-based anomaly detection techniques, categorizing them into supervised, unsupervised, and reinforcement learning paradigms. Their survey highlights the theoretical potential of quantum models while emphasizing the lack of standardized evaluation protocols and large-scale experimental validation.

Among supervised quantum approaches, kernel-based methods have received particular attention. Oh and Park [2] introduced Quantum Support Vector Data Description (QSVDD) for anomaly detection, extending the classical SVDD framework into the quantum domain. Their approach employs quantum feature maps to embed classical data into

a high dimensional Hilbert space, where a quantum kernel is used to learn a hypersphere that encloses normal data points. Samples falling outside the learned boundary are classified as anomalies. The method leverages quantum kernel estimation implemented through parameterized quantum circuits, enabling potentially richer feature representations than classical kernels. Experimental evaluations on benchmark anomaly detection datasets demonstrated that QSVDD can achieve competitive or improved detection performance compared to classical SVDD, highlighting the potential of quantum-enhanced kernel methods for anomaly detection tasks.

Hybrid quantum-classical architectures have also been explored as a strategy to balance expressivity and trainability. Rasyidi et al. [3] proposed a hybrid quantum–classical autoencoder (HQC-AE) framework for unsupervised network intrusion detection. Their approach integrates parameterized quantum circuits within the autoencoder architecture to enhance feature compression and anomaly detection. The model is trained only on normal network traffic and detects intrusions based on reconstruction error. The authors conducted a large scale experimental study exploring several quantum design choices, including quantum layer placement, measurement strategies, variational and non-variational formulations, and latent space regularization. The framework was evaluated on benchmark intrusion detection datasets such as UNSW-NB15, NSL-KDD, and CIC-IDS2017. Experimental results showed that well configured HQC autoencoders can match or sometimes outperform classical autoencoder-based approaches and demonstrate stronger generalization in zero-day attack scenarios, although the performance is sensitive to architectural design and quantum noise.

The choice of data embedding plays a crucial role in the effectiveness of QML models. Rath and Date [4] conducted a comparative analysis of several classical to quantum mapping techniques, including basis encoding, angle encoding, and amplitude encoding, to evaluate their impact on machine learning performance. Their study examined the integration of quantum data embeddings with various classical models such as Logistic Regression, K-Nearest Neighbors, Support Vector Machines, and ensemble methods. The results showed that quantum embedding strategies can improve classification accuracy and F1 scores by enabling richer feature representations, although they may introduce additional computational overhead. However, their work focuses primarily on general classification tasks and does not investigate intrusion detection scenarios or evaluate the impact of these embeddings across multiple quantum machine learning architectures.

Overall, existing studies demonstrate the potential of QML for anomaly and intrusion detection, but several important limitations remain. Most prior works focus on a single quantum model or embedding strategy, restrict experiments to binary classification tasks, or apply dimensionality reduction exclusively within classical pipelines. Furthermore, comprehensive comparisons between quantum and classical models under identical preprocessing and evaluation settings remain limited, particularly for tabular cybersecurity datasets.

In contrast, the present work provides a unified and systematic evaluation of hybrid Quantum Neural Networks and Quantum Support Vector Machines for intrusion detection. By explicitly integrating PCA-based dimensionality reduction with multiple quantum embedding strategies and benchmarking the proposed models against strong classical baselines, this study offers a comprehensive assessment of the practical benefits and limitations of QML-based intrusion detection systems under realistic NISQ constraints.

IV. METHODOLOGY

In this section, we present the methodological framework adopted in this study, which integrates classical machine learning models with hybrid quantum-classical neural networks (QNNs) and Quantum Support Vector Machines (QSVMs). The methodology is structured into four main components: dataset preprocessing, dimensionality reduction, model design, and evaluation.

A. DATASET DESCRIPTION

In this study, we utilized the widely used NSL-KDD benchmark dataset, a refined version of the original KDD'99 dataset, which is extensively employed for evaluating intrusion detection systems (IDS) [6]. The NSL-KDD dataset addresses several limitations of the KDD'99 dataset, such as the presence of redundant records and biased class distributions. It therefore provides a more balanced and representative benchmark for evaluating machine learning methods in cybersecurity applications.

To ensure computational feasibility given the simulation cost of quantum circuits, we used a representative subset corresponding to 20% of the full NSL-KDD dataset. This subset preserves the key statistical properties of the original distribution while significantly reducing the computational overhead required for training and evaluating quantum models.

B. DATA PREPROCESSING

Before applying dimensionality reduction and embedding techniques, classical preprocessing steps were performed to ensure data quality and consistency. The preprocessing pipeline consisted of three main stages:

- **Data Cleaning:** Duplicate and irrelevant records were removed to avoid training bias. Missing values, although rare in the NSL-KDD dataset, were handled using mean or mode imputation depending on the feature type.
- **Data Normalization:** Continuous attributes were normalized using Min-Max scaling to map feature values to the range $[0, 1]$. This prevents features with larger numeric ranges from dominating the learning process and ensures stable training for both classical and quantum models.

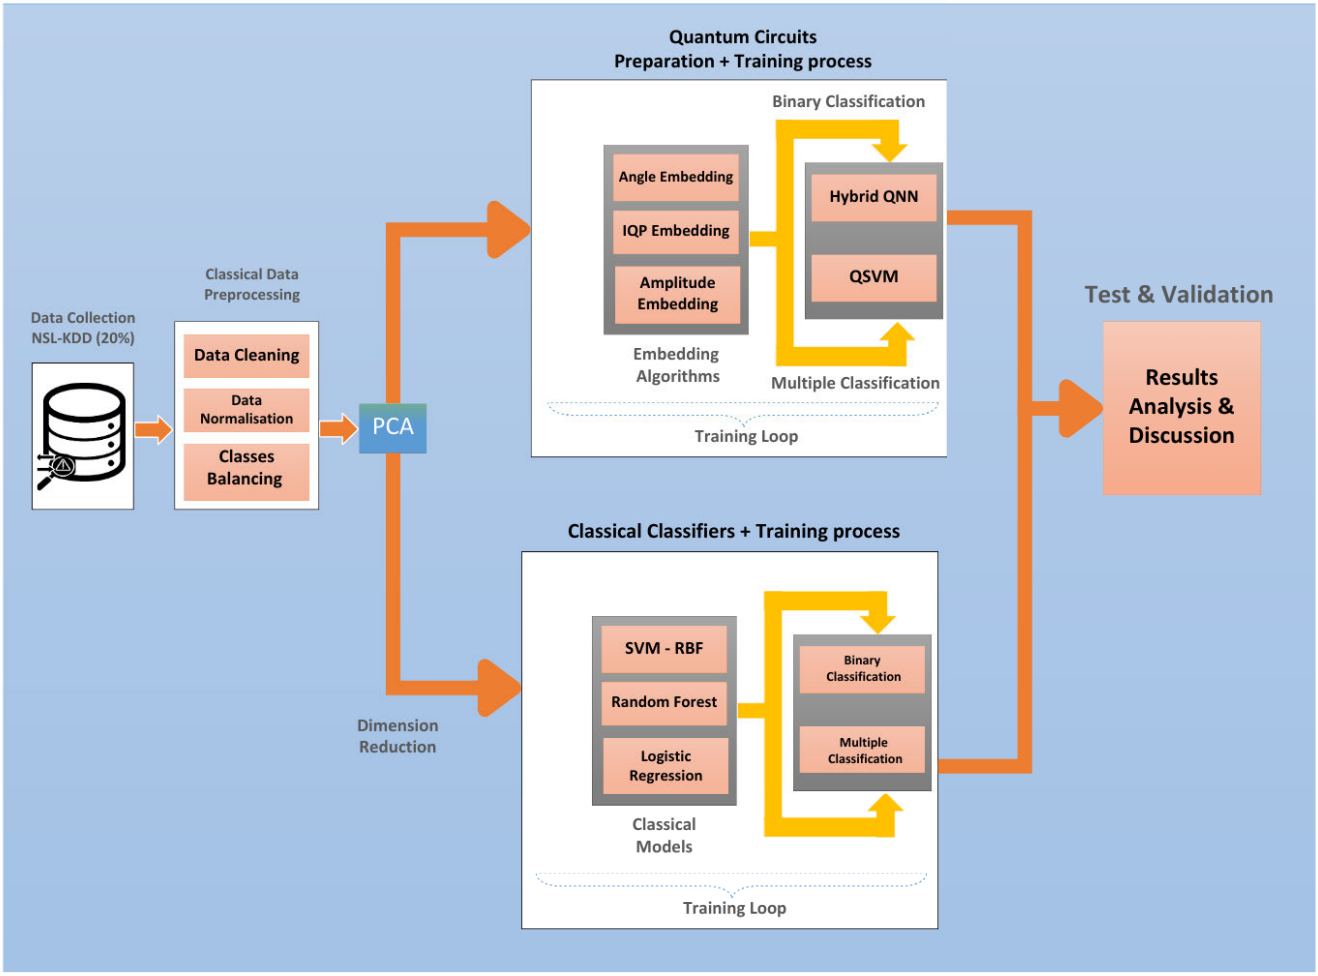

FIGURE 2. Methodology workflow of the proposed hybrid quantum intrusion detection framework.

- **Class Balancing:** The NSL-KDD dataset exhibits inherent class imbalance between normal traffic and attack classes. In the binary classification scenario, all attack categories were merged into a single class labeled *Attack*, contrasted against the *Normal* class.
  In the multiclass classification scenario, the original attack types were grouped into four major families: Denial of Service (DoS), Probe, Remote to Local (R2L), and User to Root (U2R). To mitigate the imbalance problem, we applied the Synthetic Minority Oversampling Technique (SMOTE), which increases the representation of minority classes such as R2L and U2R while maintaining comparability with the *Normal* class. This balancing step reduces the risk of the Accuracy Paradox, where a classifier achieves high overall accuracy by favoring the majority class while failing to detect minority attack patterns [7].

These preprocessing steps ensure that the dataset is properly conditioned for subsequent PCA-based dimensionality reduction and embedding within quantum machine learning pipelines.

C. DIMENSIONALITY REDUCTION USING PCA

Principal Component Analysis (PCA) has been widely applied in intrusion detection research to address the high dimensionality of datasets such as NSL-KDD. Prior studies have used PCA to identify compact feature subsets while preserving discriminative information [8], or combined it with classical classifiers to maintain performance under dimensionality constraints. The theoretical foundations of PCA and its variance-based component selection are well documented in [9] and [10], Comparable preprocessing methodologies have been demonstrated to improve intrusion detection system performance by optimizing feature representation and lowering computational complexity [11].

The NSL-KDD dataset presents two main challenges:

- It is highly imbalanced across attack categories.
- It contains numerous heterogeneous features, many of which are redundant or weakly informative.

PCA mitigates these issues by projecting the original feature space into a lower dimensional subspace that concentrates the variance within the first few principal

TABLE 1. Mapping of NSL-KDD attack classes into four families.

Attack Family	Original Classes
DoS (Denial of Service)	back, land, neptune, pod, smurf, teardrop
Probe	ipsweep, nmap, portsweep, satan
R2L (Remote to Local)	ftp_write, guess_passwd, imap, multihop, phf, spy, warezclient, warezmaster
U2R (User to Root)	buffer_overflow, loadmodule, rootkit

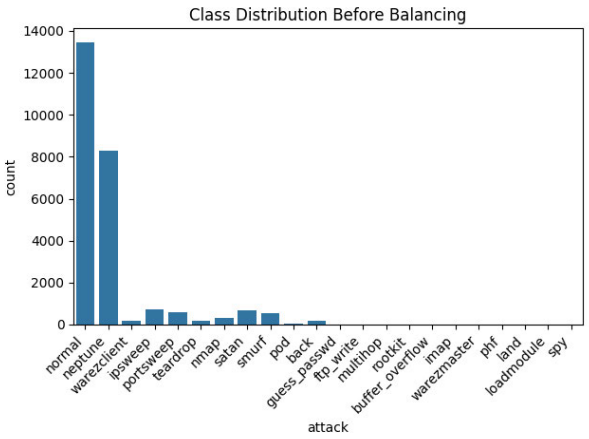

(a) Before Balancing

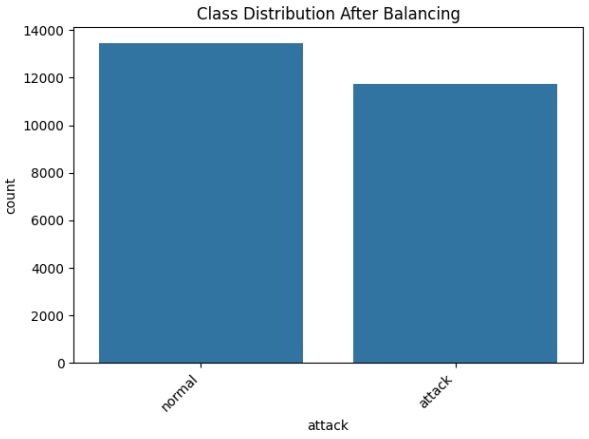

(b) After Balancing

FIGURE 3. Class distribution of the NSL-KDD dataset before and after balancing for binary classification.

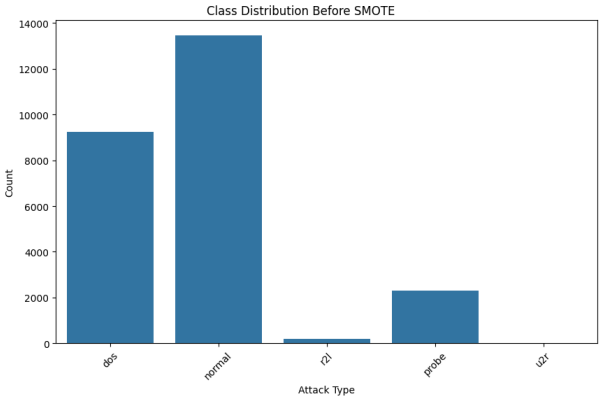

(a) Before SMOTE

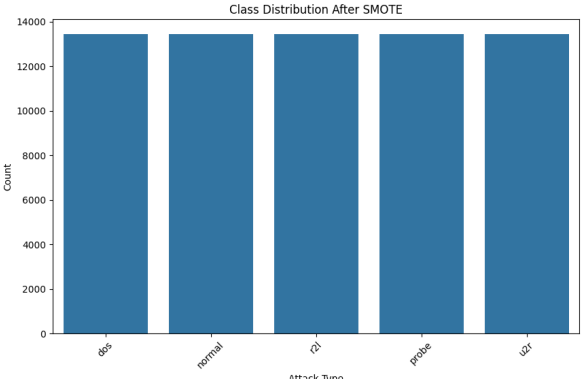

(b) After SMOTE

FIGURE 4. Class distribution of the NSL-KDD dataset before and after applying SMOTE.

components. This reduces computational complexity and limits the influence of noisy or redundant features.

In this work, PCA was applied to both binary and multi-class classification tasks. The cumulative explained variance was used to guide the selection of retained dimensions. Specifically, we considered $k \in \{2, 4, 6, 8, 10\}$ for Hybrid QNNs and classical baselines, while QSVM experiments were restricted to $k \in \{2, 4, 6, 8\}$ to remain compatible with the qubit and circuit depth limitations of current NISQ devices.

Formally, the PCA transformation is defined as

$$Z = XW_k \tag{2}$$

where X denotes the standardized data matrix, W_k represents the matrix composed of the top-k eigenvectors of the covariance matrix, and Z is the projected dataset in the reduced feature space.

For binary classification, Figure 5 illustrates the explained variance ratio and cumulative variance distribution across principal components.

For multiclass classification, the corresponding PCA analysis is shown in Figure 6.

The covariance matrix is defined as:

$$\Sigma = \frac{1}{n} X^T X \tag{3}$$

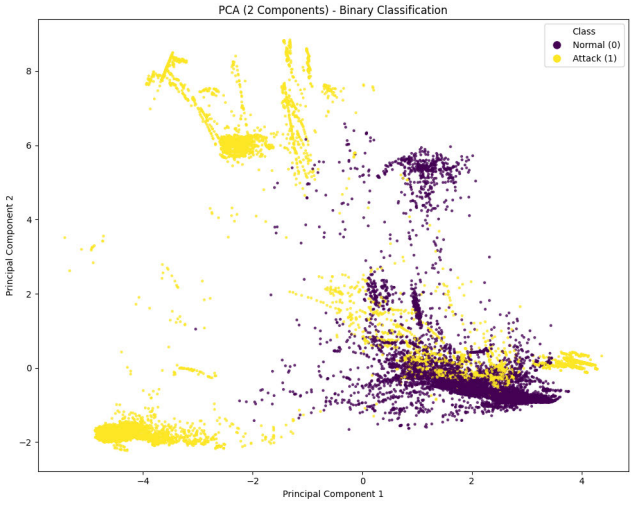

(a) Explained variance ratio for binary PCA.

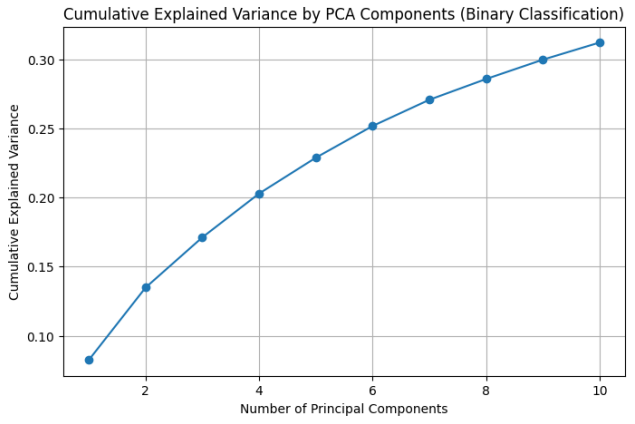

(b) Cumulative explained variance for binary PCA.

FIGURE 5. PCA results for binary classification.

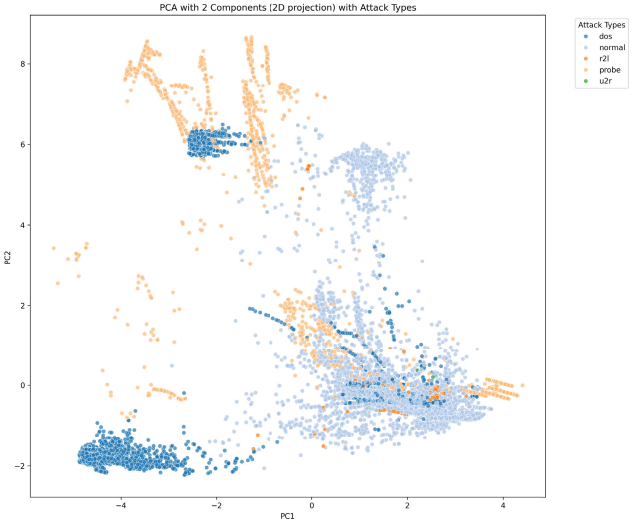

(a) Explained variance ratio for multiclass PCA.

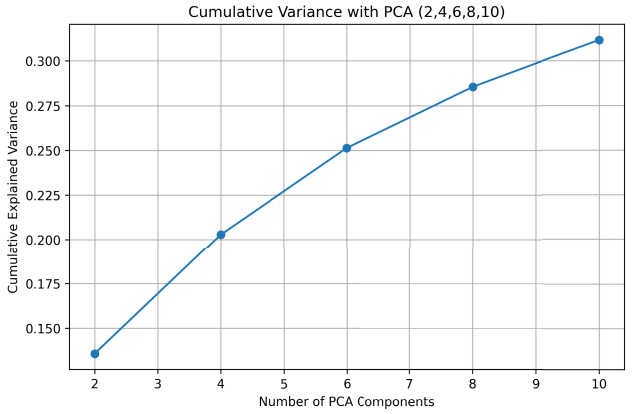

(b) Cumulative explained variance for multiclass PCA.

FIGURE 6. PCA results for multiclass classification.

where X is the centered data matrix. The projection matrix W_k is constructed from the eigenvectors corresponding to the top-k eigenvalues, preserving the maximum variance.

The retained variance ratio is defined as:

$$\text{Explained Variance Ratio} = \frac{\sum_{i=1}^{k} \lambda_i}{\sum_{i=1}^{d} \lambda_i} \tag{4}$$

D. QUANTUM DATA EMBEDDING

After dimensionality reduction, the resulting classical feature vectors must be encoded into quantum states before being processed by a quantum circuit. This step, commonly referred to as *quantum data embedding* or *quantum feature mapping*, transforms classical information into a representation suitable for quantum computation. By mapping classical data into high dimensional Hilbert spaces, quantum embeddings enable nonlinear feature transformations that can improve the expressive capability of quantum machine learning models.

Let the classical feature vector be defined as

$$\mathbf{x} = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^n. \tag{5}$$

The embedding procedure converts this vector into a quantum state

$$|\psi(\mathbf{x})\rangle \in \mathcal{H}^{2^n}, \tag{6}$$

where $\mathcal{H}^{2^n}$ denotes the 2^n-dimensional Hilbert space associated with an n-qubit quantum system. Different embedding strategies define different mappings between classical data and quantum states, resulting in varying circuit complexity and representational capacity.

In this work, three commonly used embedding techniques are investigated: angle embedding, amplitude embedding, and Instantaneous Quantum Polynomial (IQP) embedding.

- **Angle Embedding**
  Angle embedding encodes classical features into the rotation angles of parameterized single-qubit gates. Each feature controls the rotation applied to a qubit initialized in the ground state $|0\rangle$. A common implementation uses rotations around the Y axis, resulting in the

following quantum state representation:

$$|\psi(\mathbf{x})\rangle = \bigotimes_{i=1}^{n} R_Y(x_i)|0\rangle. \tag{7}$$

where $R_Y(x_i)$ denotes the rotation operator

$$R_Y(\theta) = \begin{bmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{bmatrix}. \tag{8}$$

This encoding strategy produces shallow circuits with low gate complexity, making it suitable for noisy intermediate scale quantum (NISQ) devices [12]. However, because each feature independently controls a single-qubit rotation, the ability of angle embedding to capture complex feature relationships may be limited.

- **Amplitude Embedding**
  Amplitude embedding maps a normalized classical feature vector directly into the amplitudes of a quantum state. Given a normalized vector

$$\mathbf{x} = (x_0, x_1, \ldots, x_{2^n-1}), \quad \sum_i |x_i|^2 = 1, \tag{9}$$

the resulting quantum state is defined as

$$|\psi(\mathbf{x})\rangle = \sum_{i=0}^{2^n-1} x_i|i\rangle. \tag{10}$$

This method allows 2^n classical features to be encoded using only n qubits, providing an exponentially compact representation of classical data. However, preparing such states generally requires complex state preparation circuits with deeper gate sequences, which may introduce noise and training difficulties on current quantum hardware [13], [14].

- **IQP Embedding**
  Instantaneous Quantum Polynomial (IQP) embedding constructs expressive feature maps using commuting diagonal operators that introduce nonlinear phase relationships between classical features. The circuit typically consists of an initial layer of Hadamard gates followed by data dependent phase operations.
  The encoded quantum state can be written as

$$|\psi(\mathbf{x})\rangle = U_{\text{IQP}}(\mathbf{x})H^{\otimes n}|0\rangle^{\otimes n}, \tag{11}$$

where the feature dependent unitary operator is defined as

$$U_{\text{IQP}}(\mathbf{x}) = \exp\left(i\sum_i x_i Z_i + i\sum_{i<j} x_i x_j Z_i Z_j\right). \tag{12}$$

Here, Z_i represents the Pauli-Z operator acting on qubit i, while Z_iZ_j introduces entanglement between pairs of qubits. These interactions allow IQP embedding to capture nonlinear correlations between classical features in the quantum Hilbert space, thereby improving the expressivity of quantum kernel methods [15], [16].

Each embedding strategy presents different trade-offs in terms of circuit depth, expressivity, and hardware feasibility. Evaluating their impact is therefore essential for understanding how quantum feature maps influence the performance of quantum machine learning models for intrusion detection tasks.

E. QUANTUM NEURAL NETWORK ARCHITECTURE

The Hybrid Quantum Neural Network (HQNN) implemented in this work follows a variational quantum circuit paradigm that integrates classical neural processing with quantum feature transformation.

Input network traffic features are first reduced using PCA and processed by a classical fully connected layer (FC1), which serves as a learnable preprocessing stage before quantum encoding.

The output of FC1 is embedded into a parameterized quantum circuit using one of the proposed embedding strategies. The resulting quantum state is then transformed through layers of trainable single qubit rotations and entangling gates forming a variational ansatz parameterized by θ.

The variational quantum circuit (ansatz) used in this work follows a hardware-efficient design composed of alternating layers of parameterized single-qubit rotations and entangling operations.

In each layer, rotation gates of the form $R_Y(\theta)$ and $R_Z(\theta)$ are applied independently to every qubit. These are followed by entangling operations implemented the Controlled-NOT (CNOT) gates arranged in a linear topology, where each qubit is entangled with its nearest neighbor. This structure enables efficient information propagation while maintaining shallow circuit depth.

For a system of n qubits and L layers, the total number of trainable parameters scales as $O(n \times L)$. In this work, we fix the number of layers to $L = 2$, as illustrated in Figures 7–9, to balance expressivity and trainability under NISQ constraints.

This architecture is inspired by hardware-efficient ansätze commonly used in variational quantum algorithms, enabling compatibility with near-term quantum devices.

This explicit specification ensures that the proposed HQNN architecture can be directly reproduced across different quantum machine learning frameworks.

TABLE 2. Variational quantum circuit (Ansatz) configuration.

Component	Specification
Ansatz Type	Hardware-efficient (Strongly Entangling Layers)
Number of Qubits	Equal to PCA components (n)
Number of Layers (L)	2
Single-Qubit Gates	$R_Y(\theta)$, $R_Z(\theta)$
Entanglement Strategy	Linear nearest-neighbor CNOT chain
Parameterization	Independent parameters per qubit per layer
Measurement	Pauli-Z expectation values
Output	Expectation values passed to classical FC layer

Quantum measurements produce expectation values that are forwarded to a classical post-processing layer (FC2), which outputs the final binary or multiclass prediction.

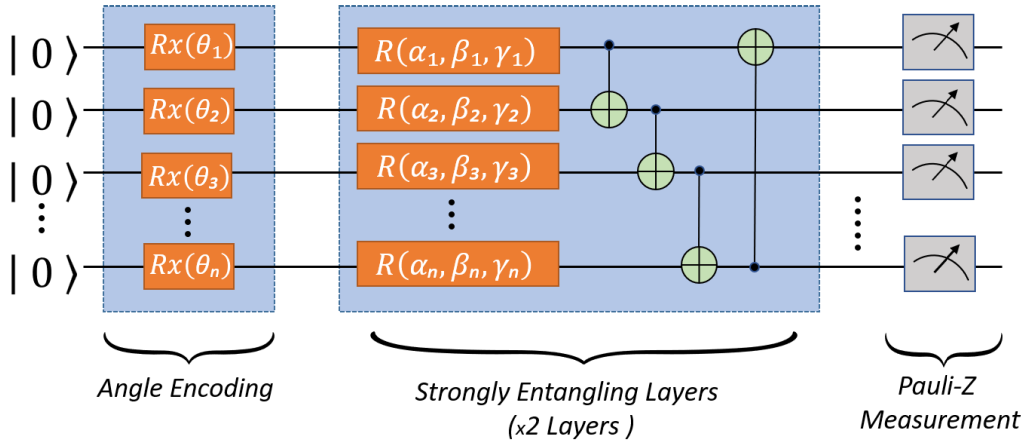

FIGURE 7. Quantum circuit representation of angle embedding where classical features control the rotation angles of R_Y gates applied to qubits initialized in the state $|0\rangle$.

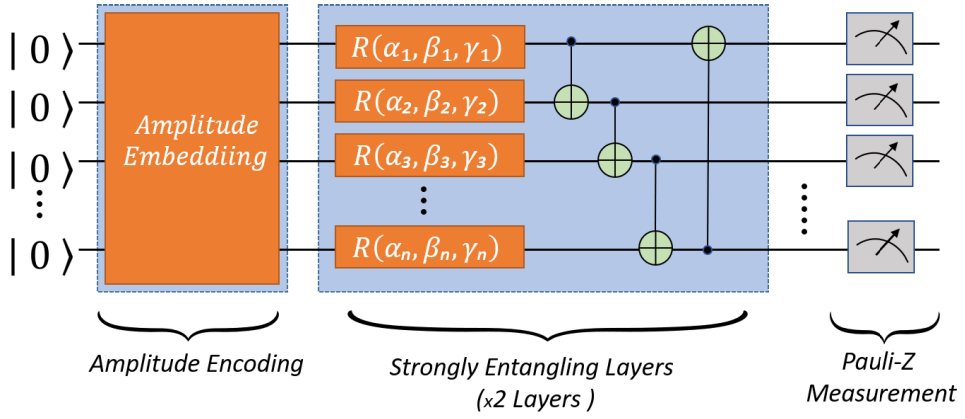

FIGURE 8. Conceptual quantum circuit for amplitude embedding, where a classical vector x is encoded into the amplitudes of a quantum state using a state preparation unitary $U(\mathbf{x})$.

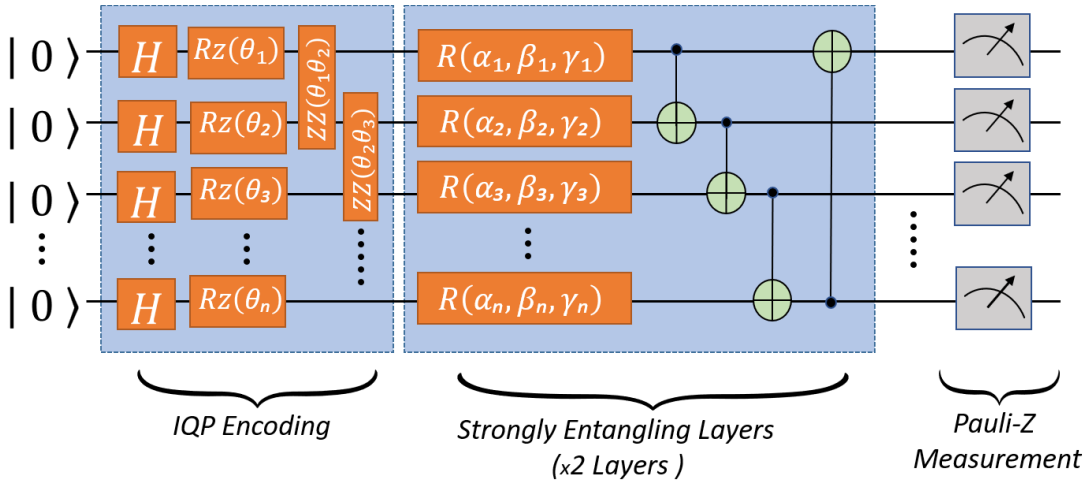

FIGURE 9. Quantum circuit representation of the IQP embedding feature map consisting of Hadamard gates, data dependent phase rotations, and entangling interactions between qubits.

Specifically, measurements are performed using Pauli-Z operators on each qubit, and the resulting expectation values are used as inputs to the classical output layer.

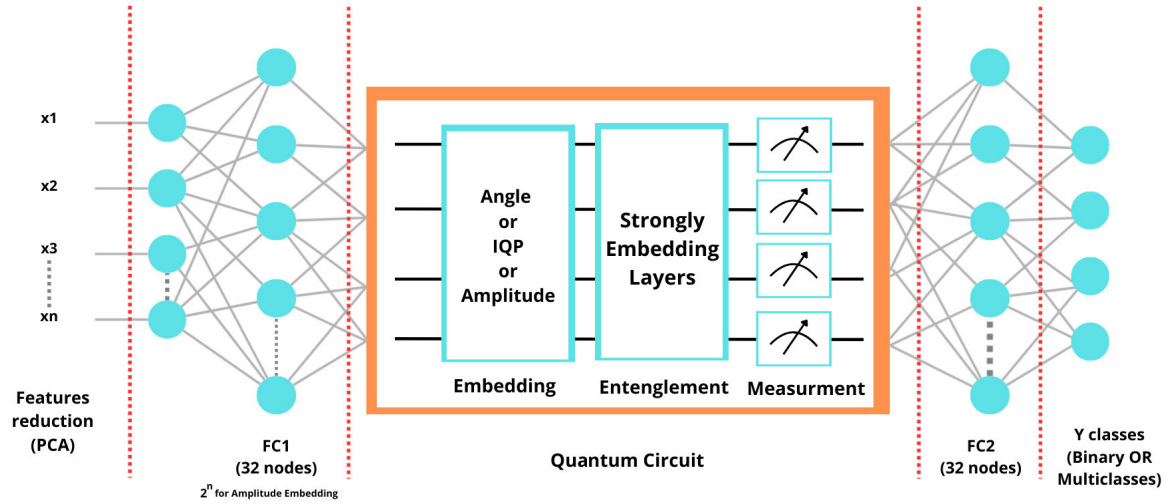

FIGURE 10. Hybrid quantum neural network (HQNN) architecture.

The complete circuit structures corresponding to the embedding strategies and variational layers are illustrated in Figures 7, 8, and 9, ensuring full reproducibility of the proposed quantum models.

Formally, the model output is defined as

$$y = f\left(\langle\psi(\mathbf{z},\theta)|\hat{O}|\psi(\mathbf{z},\theta)\rangle\right) \tag{13}$$

where $|\psi(\mathbf{z},\theta)\rangle$ denotes the quantum state after embedding and variational transformations, $\hat{O}$ is the measurement operator, and $f(\cdot)$ represents a classical decision function.

F. QUANTUM SUPPORT VECTOR MACHINE

In addition to the QNN architecture, this work also evaluates a Quantum Support Vector Machine (QSVM), which leverages quantum kernel methods for classification.

After classical preprocessing and PCA-based dimensionality reduction, feature vectors are encoded into quantum states using the selected embedding circuits.

Unlike QNNs, QSVMs do not involve trainable quantum parameters. Instead, the quantum circuit implicitly computes kernel values by evaluating overlaps between quantum states.

Given two samples $\mathbf{z}_i$ and $\mathbf{z}_j$, the quantum kernel is defined as

$$K(\mathbf{z}_i,\mathbf{z}_j) = |\langle\phi(\mathbf{z}_i)|\phi(\mathbf{z}_j)\rangle|^2 \tag{14}$$

where $|\phi(\mathbf{z})\rangle$ represents the quantum state generated by the embedding circuit.

In practice, the quantum kernel is estimated using a fidelity-based approach implemented via the adjoint method. Specifically, for two input vectors $\mathbf{z}_i$ and $\mathbf{z}_j$, the kernel is computed by applying the embedding circuit to $\mathbf{z}_i$ followed by the adjoint (inverse) embedding of $\mathbf{z}_j$, and measuring the probability of returning to the all-zero state:

$$K(\mathbf{z}_i,\mathbf{z}_j) = |\langle 0|U^\dagger(\mathbf{z}_j)U(\mathbf{z}_i)|0\rangle|^2 \tag{15}$$

This approach enables efficient estimation of state overlap without explicitly computing inner products in Hilbert space.

The full quantum kernel matrix $K \in \mathbb{R}^{N\times N}$ is constructed by evaluating pairwise kernel values between all training samples. Due to the quadratic complexity of kernel evaluation, a subset of approximately 1000 samples is used to ensure computational feasibility.

Kernel computation is performed using batched parallel evaluation to accelerate simulation, and all circuits are executed on the `lightning.qubit` simulator backend. The number of measurement shots is fixed to 1024 to balance estimation accuracy and computational cost.

The resulting kernel matrix is used within a classical Support Vector Machine (SVM) classifier with a precomputed kernel. The regularization parameter C is selected via grid search over the range $C \in \{0.1, 0.5, 1, 10, 100\}$. Model training and evaluation follow an 80/20 train-test split with stratified sampling.

G. IMPLEMENTATION DETAILS AND REPRODUCIBILITY

To ensure full reproducibility of the proposed framework, all experiments were implemented using the *PennyLane* quantum machine learning library integrated with *PyTorch* for hybrid optimization.

The quantum circuits were simulated using the `default.qubit` and `lightning.qubit` simulators. All experiments were conducted on a classical computing environment with fixed random seeds to ensure reproducibility.

The complete implementation, including preprocessing scripts, model architectures, hyperparameters, and evaluation pipelines, is publicly available at:

https://github.com/raidanis/QML-IDS-NSLKDD

Table 3 summarizes the key training parameters used for the HQNN and QSVM models.

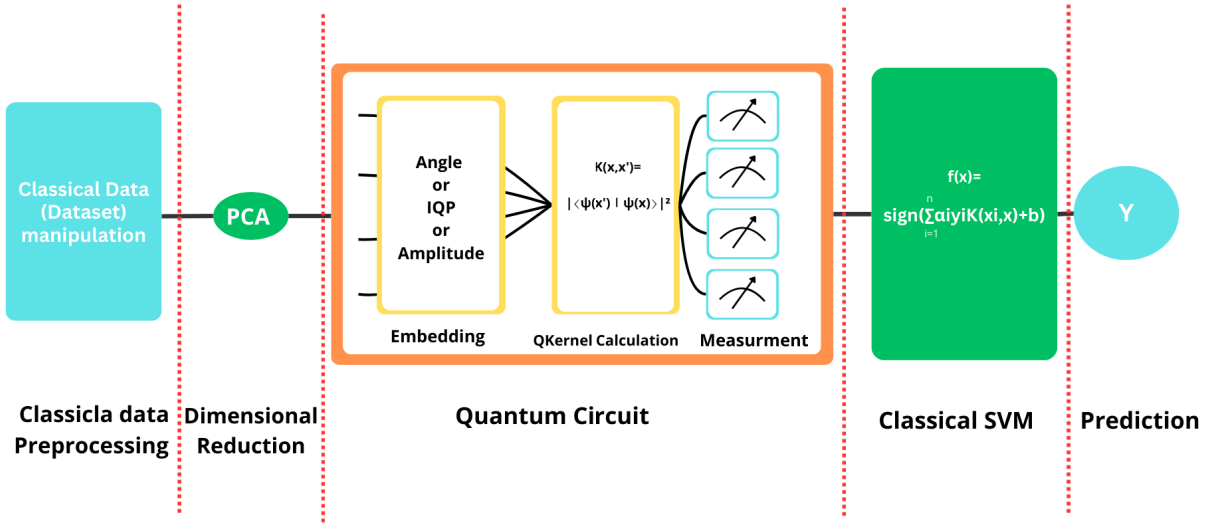

FIGURE 11. Hybrid quantum support vector machine (QSVM) architecture.

TABLE 3. QSVM configuration and quantum kernel estimation.

Component	Specification
Kernel Type	Quantum fidelity kernel
Embedding Methods	Angle, Amplitude, IQP
Kernel Computation	$U(\mathbf{x}_i) + U^{\dagger}(\mathbf{x}_j)$
Measurement	Probability of $\|0\rangle$ state
Simulator	`lightning.qubit`
Number of Shots	1024
Dataset Size	∼1000 samples (subsampled)
Kernel Matrix	Precomputed ($N \times N$)
Parallelization	Batch + multi-core execution
SVM Type	Precomputed kernel SVM
Regularization (C)	$\{0.1, 0.5, 1, 10, 100\}$
Train/Test Split	80/20 (stratified)

H. ALGORITHMIC OVERVIEW

Algorithm 1 presents the detailed training procedure of the proposed HQNN model.

Algorithm 1 HQNN Training Pipeline

Require: Dataset X, labels y
Ensure: Trained HQNN model
1: Apply preprocessing (cleaning, normalization, balancing)
2: Apply PCA to obtain reduced features Z
3: Initialize classical and quantum parameters
4: **for** each epoch **do**
5: Encode Z into quantum states
6: Apply variational quantum circuit
7: Measure expectation values
8: Compute loss
9: Update parameters
10: **end for**
11: **return** trained HQNN model

I. BASELINE MODELS AND EVALUATION PROTOCOL

To evaluate the effectiveness of the proposed QML framework, the HQNN and QSVM models are compared against several classical baselines, including Support Vector Machines with radial basis function kernels, Random Forests, and classical neural networks.

All models are trained and evaluated using identical preprocessing steps, PCA configurations, and dataset splits to ensure a fair comparison. The dataset is divided into training and testing sets using an 80/20 split, and hyperparameters for classical models are selected using cross validation on the training data.

Model performance is evaluated using standard classification metrics including accuracy, precision, recall, and F1-score. This unified evaluation protocol allows a transparent comparison between classical and quantum approaches for both binary and multiclass intrusion detection tasks.

J. NISQ CONSTRAINTS AND COMPUTATIONAL CONSIDERATIONS

The proposed framework operates under NISQ-era limitations, including restricted qubit counts, circuit noise, and limited coherence times. To address these constraints, shallow circuits and low-dimensional embeddings were employed.

Quantum simulations exhibit significantly higher computational cost compared to classical models, particularly for QSVM kernel estimation, which scales quadratically with the number of samples. This motivates the use of dimensionality reduction and dataset subsampling.

These constraints highlight the importance of balancing expressivity and computational feasibility when designing QML-based intrusion detection systems.

V. RESULTS AND DISCUSSION

In this section, we present the performance evaluation of the proposed QML-based intrusion detection models. Results are

reported for both binary and multiclass classification tasks and are compared against classical baselines. Evaluation metrics include accuracy, F1-score, and confusion matrices. In addition, we analyze the impact of quantum embedding strategies and PCA-based dimensionality reduction.

A. SCOPE OF EXPERIMENTAL COMPARISON

It is important to clarify the scope of the experimental evaluation conducted in this study. The primary objective of this work is not to provide an exhaustive benchmark across all existing quantum intrusion detection systems, but rather to perform a controlled and systematic analysis of hybrid quantum machine learning design choices within a unified framework.

Specifically, this study focuses on: (i) comparing two major QML paradigms (variational Quantum Neural Networks and quantum kernel-based QSVM), (ii) analyzing the impact of different quantum embedding strategies (angle, amplitude, IQP), and (iii) evaluating the role of PCA-based dimensionality reduction under NISQ constraints.

To ensure methodological consistency and fair comparison, all models are implemented within the same pipeline, using identical preprocessing, feature spaces, and evaluation protocols.

While recent approaches such as hybrid quantum autoencoders and QSVDD-based methods have demonstrated promising results, they rely on fundamentally different learning paradigms (unsupervised or reconstruction-based objectives), making direct comparison non-trivial without introducing additional experimental variables.

Therefore, incorporating such models within the current framework would require a substantial redesign of the experimental protocol, which falls outside the scope of this study. Instead, this work aims to provide a focused and reproducible evaluation of supervised hybrid QML architectures for intrusion detection.

A comprehensive benchmark including heterogeneous quantum IDS paradigms is identified as an important direction for future research.

B. BINARY INTRUSION DETECTION RESULTS

This section reports the performance of the proposed hybrid Quantum Machine Learning models for binary intrusion detection (Normal vs. Attack) on the NSL-KDD dataset. To maintain clarity and conciseness, we present results only for the best performing configurations within each model family: Hybrid Quantum Neural Networks (H-QNN), Quantum Support Vector Machines (QSVM), and classical baselines. All models are evaluated under identical preprocessing pipelines, PCA configurations, and train-test splits.

1) OVERALL PERFORMANCE COMPARISON

Table 4 summarizes the binary classification performance of the selected models. Among classical approaches, the Random Forest classifier achieves the strongest baseline performance, benefiting from ensemble learning and robustness to feature noise. Kernel-based SVMs and classical neural networks exhibit competitive accuracy but show greater sensitivity to the choice of PCA dimensionality.

The H-QNN-based intrusion detection system achieves performance comparable to, and in certain configurations slightly exceeding, the best classical baselines. Notably, the H-QNN demonstrates strong generalization despite operating on a reduced feature space, indicating that the quantum embedding and variational circuit effectively capture discriminative patterns relevant to intrusion detection. This observation suggests that hybrid quantum-classical models can compensate for aggressive dimensionality reduction through expressive quantum feature representations.

The QSVM model also achieves competitive results, particularly when combined with expressive embedding strategies such as amplitude encoding. While its absolute accuracy remains close to that of the best performing H-QNN configuration, the QSVM exhibits stable behavior across PCA settings, highlighting the robustness of quantum kernel-based methods in binary intrusion detection tasks.

2) CONFUSION MATRIX EVALUATION

Figure 12 presents the confusion matrices for the best performing models. Although overall accuracy values appear similar, the distribution of errors reveals meaningful differences. The H-QNN with amplitude embedding achieves the lowest false negative rate, indicating improved sensitivity to attack traffic. In intrusion detection scenarios, minimizing false negatives is critical because undetected attacks pose significant security risks.

The Random Forest baseline, while competitive in overall accuracy, exhibits a slightly higher false positive rate, potentially leading to increased false alarms in practical deployments. The QSVM demonstrates balanced classification behavior, further confirming the robustness of quantum kernel-based approaches when operating on reduced dimensional feature representations.

3) IMPACT OF QUANTUM EMBEDDING

The choice of quantum embedding plays a critical role in binary classification performance. Angle embedding consistently yields stable results with shallow circuits and relatively low training overhead, making it particularly suitable for NISQ constrained settings. Amplitude embedding provides compact state representations but introduces increased circuit depth, which may affect trainability. IQP embedding offers enhanced expressivity through entanglement and nonlinear phase interactions, leading to improved performance in several configurations, particularly for QSVM models.

These findings reveal an important trade-off between circuit complexity and classification performance, emphasizing the importance of embedding selection when deploying QML-based intrusion detection systems on near term quantum hardware.

TABLE 4. Best binary classification results on NSL-KDD Dataset for H-QNN, QSVM, and classical models.

Model	Embedding	PCA Components	Accuracy	F1 Score	Training Time (s)
H-QNN	Angle	10	0.993	0.993	3745.980
	Amplitude	10	0.993	0.993	4620.45
	IQP	6	0.977	0.977	3947.880
QSVM	Angle	8	0.985	0.984	35000.00
	Amplitude	8	0.988	0.987	35200.00
	IQP	6	0.957	0.955	28780.00
SVM (RBF)	-	8	0.963	0.961	15.2
Random Forest	-	10	0.972	0.970	14.8
Logistic Regression	-	10	0.948	0.945	3.2

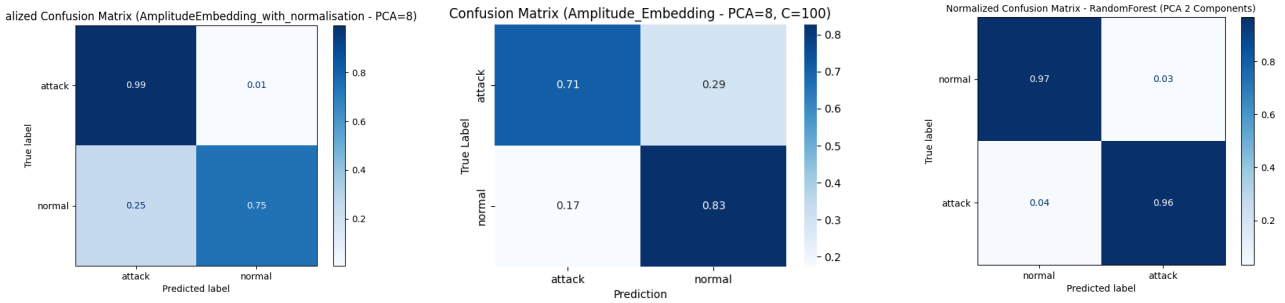

(a) H-QNN (Amplitude, PCA=8) (b) QSVM (Amplitude, PCA=8) (c) Random Forest

FIGURE 12. Confusion matrices of the best-performing quantum and classical models on the NSL-KDD binary classification task.

4) DISCUSSION

Overall, the binary intrusion detection results demonstrate that hybrid QML models can achieve detection performance comparable to strong classical baselines while operating in significantly reduced feature spaces. The H-QNN, in particular, benefits from variational quantum circuits that enable flexible nonlinear decision boundaries under NISQ constraints. Although classical ensemble methods such as Random Forest remain highly competitive, the observed performance of H-QNN and QSVM models provides empirical evidence that QML constitutes a viable alternative modeling paradigm for binary intrusion detection.

C. MULTICLASS INTRUSION DETECTION RESULTS

This section evaluates the performance of quantum and classical models on the multiclass intrusion detection task, where network traffic is categorized into five classes: Normal, DoS, Probe, R2L, and U2R. Compared to binary classification, this task presents a higher level of complexity due to severe class imbalance and overlapping feature distributions among attack categories.

1) OVERALL MULTICLASS PERFORMANCE

Table 5 summarizes the multiclass classification results in terms of accuracy and macro-averaged F1-score.

The Hybrid QNN demonstrates balanced performance across classes, achieving competitive macro-F1 scores despite operating in a reduced dimensional feature space. This behavior suggests that the variational quantum circuit combined with quantum embedding improves the model's ability to capture subtle decision boundaries associated with minority attack patterns.

2) CONFUSION MATRIX EVALUATION

Figure 13 illustrates the class wise confusion patterns. While overall accuracy differences appear moderate, the distribution of errors across minority classes reveals meaningful distinctions. The H-QNN with amplitude embedding demonstrates improved detection of R2L and U2R attacks, reducing false negatives compared to classical baselines.

In intrusion detection systems, minority attack classes are particularly challenging due to severe class imbalance. The improved recall observed for these categories suggests that quantum-enhanced feature representations may capture subtle structural patterns that are not fully exploited by classical ensemble methods.

The Random Forest model, although competitive in overall accuracy, exhibits increased misclassification between Probe and DoS categories, reflecting the difficulty of separating closely related traffic behaviors under skewed distributions.

3) EFFECT OF PCA AND QUANTUM EMBEDDING

The impact of dimensionality reduction is more pronounced in the multiclass scenario than in binary classification. Aggressive PCA reduction negatively affects classical classifiers by removing discriminative information required to separate closely related attack types. Quantum models, particularly QNNs, exhibit greater resilience to dimensionality reduction, maintaining stable performance even with a limited number of principal components.

TABLE 5. Best multiclass classification results on NSL-KDD dataset for QML and classical models.

Model	Embedding	PCA Components	Accuracy	F1 Score	Training Time (s)
H-QNN	Angle	6	0.9240	0.9300	3007.89
	Amplitude	10	0.9476	0.9466	8182.48
	IQP	6	0.8104	0.8206	4608.77
QSVM	Angle	8	0.920	0.950	35213.49
	Amplitude	8	0.960	0.920	39112.32
	IQP	6	0.933	0.930	28784.01
SVM (RBF)	-	10	0.895	0.890	16.0
Random Forest	-	10	0.920	0.918	14.2
Logistic Regression	-	10	0.875	0.870	3.5

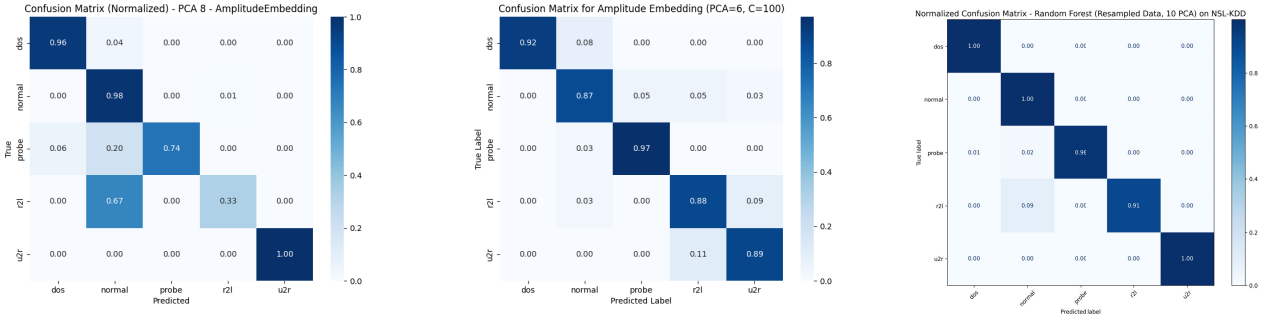

(a) H-QNN (Amplitude, PCA=8) (b) QSVM (Amplitude, PCA=8) (c) Random Forest

FIGURE 13. Confusion matrices for the best-performing models in multiclass intrusion detection on NSL-KDD.

Among the embedding strategies, IQP embedding consistently improves multiclass discrimination by introducing entanglement induced nonlinear correlations between features. Angle embedding remains computationally efficient and stable, while amplitude embedding provides compact encodings at the cost of increased circuit depth. These observations highlight a trade-off between expressivity and trainability that is central to the design of quantum-enhanced IDS.

4) DISCUSSION

The multiclass results provide stronger evidence of the potential advantages of QML in intrusion detection compared to binary classification. While classical ensemble methods remain effective in terms of raw accuracy, quantum models demonstrate improved robustness in handling class imbalance and complex attack distributions. In particular, the Hybrid QNN shows promise as a scalable and expressive model for fine grained intrusion categorization under NISQ constraints.

These findings suggest that QML-based IDS architectures are especially well suited for multiclass intrusion detection tasks, where the ability to model high order feature interactions and nonlinear class boundaries is critical.

D. COMPREHENSIVE DISCUSSION AND INTERPRETATION

Taken together, the binary and multiclass experiments provide a coherent view of the capabilities and current limitations of hybrid quantum machine learning for intrusion detection.

First, in terms of predictive performance, both H-QNN and QSVM models achieve results comparable to strong classical baselines across the evaluated metrics. In the binary classification setting, the quantum models obtain accuracy and F1-scores that are competitive with Random Forest and RBF-SVM classifiers. Examination of the confusion matrices indicates that the H-QNN maintains a balanced distribution between correctly detected attacks and correctly identified normal traffic, suggesting that the model does not overly favor either class during optimization. This balance is particularly important in intrusion detection scenarios, where both missed attacks and excessive false alarms can degrade operational usability.

Second, the multiclass experiments provide further insight into the behavior of quantum-enhanced models. While classical ensemble methods remain highly competitive in terms of overall accuracy, the Hybrid QNN exhibits improved stability in recognizing minority attack categories such as R2L and U2R, which are known to be difficult classes in the NSL-KDD dataset. The improved recognition of these classes is reflected in the confusion matrices, where fewer samples are misclassified into dominant categories compared to some classical baselines. This behavior suggests that variational quantum circuits, through entanglement and nonlinear feature transformations in Hilbert space, may capture higher-order correlations among features that are not easily exploited by shallow classical models operating on compressed inputs.

Third, dimensionality reduction plays a central role in the proposed pipeline. Because near-term quantum circuits are constrained by the number of available qubits, PCA is used

to project the original feature space into a lower dimensional representation. In the experiments, classical models exhibit a noticeable performance degradation when the dimensionality is aggressively reduced. In contrast, the quantum models maintain relatively stable performance across different PCA configurations. This observation suggests that quantum feature embeddings may help mitigate part of the information loss introduced by classical dimensionality reduction by mapping compressed feature vectors into richer representations within a high dimensional Hilbert space.

It is also important to acknowledge certain dataset-related limitations. The experiments in this study are conducted using the NSL-KDD dataset, which remains one of the most widely adopted benchmarks for intrusion detection research. However, because it is derived from the earlier KDD'99 dataset, it may not fully reflect the characteristics of modern network traffic and contemporary cyber threats. Consequently, while the results provide valuable insights into the behavior of hybrid quantum-classical models, further evaluation on more recent intrusion detection datasets would be necessary to confirm the generalizability of the proposed approach in real-world environments.

From a computational perspective, however, the quantum approaches introduce notable overhead. QSVM models require the evaluation of quantum kernel matrices whose complexity grows quadratically with the number of training samples, leading to longer training times compared with classical classifiers. Similarly, the training of variational quantum circuits in the Hybrid QNN involves iterative parameter optimization and repeated circuit evaluations on simulators. While the circuits remain compatible with the depth limitations of current NISQ devices, scalability to very large datasets is currently constrained by simulation costs and limited qubit resources.

It is therefore important to clarify that the results presented in this study do not claim a definitive quantum advantage in the strict computational complexity sense. Instead, the findings provide empirical evidence that hybrid quantum-classical architectures can achieve competitive performance for intrusion detection tasks when carefully designed with suitable embeddings and circuit structures. In particular, the results indicate that quantum models can remain robust under dimensionality constraints and class imbalance, two conditions commonly encountered in cybersecurity datasets.

Overall, the experimental results demonstrate that QNN and QSVM models constitute viable alternatives to classical intrusion detection approaches in scenarios where feature dimensionality must be reduced and complex feature interactions remain important for classification. The observed stability across PCA configurations and the improved recognition of minority attack categories motivate further investigation of quantum-enhanced IDS architectures on larger datasets and, ultimately, on real quantum hardware.

Consequently, hybrid quantum machine learning should currently be viewed not as a replacement for classical intrusion detection systems, but as a complementary modeling paradigm whose potential advantages emerge in high-complexity classification problems characterized by nonlinear feature interactions and constrained representational spaces.

1) PRACTICAL IMPLICATIONS FOR INTRUSION DETECTION SYSTEMS

From a practical cybersecurity perspective, the results suggest that hybrid quantum–classical models may offer advantages in scenarios where feature dimensionality must be constrained while preserving complex feature interactions. In operational intrusion detection systems, dimensionality reduction is often required to ensure real-time processing and manageable model complexity. The experimental results indicate that quantum embeddings combined with variational circuits can maintain competitive predictive performance even under aggressive feature compression.

This property may be particularly valuable in resource-constrained environments such as edge networks, IoT infrastructures, or distributed monitoring systems where computational efficiency and compact feature representations are critical. While current NISQ hardware limitations prevent immediate deployment of large-scale quantum IDS, the proposed hybrid framework provides a conceptual pathway for integrating quantum-enhanced feature representations into future cybersecurity pipelines as quantum hardware continues to mature.

In addition, it is important to note that several recent quantum intrusion detection approaches, such as hybrid quantum autoencoders and quantum data description models, operate under different learning paradigms and optimization objectives. While these methods have demonstrated promising results, a direct comparison would require unified evaluation protocols and consistent preprocessing pipelines. Establishing such a standardized benchmarking framework remains an open challenge in quantum cybersecurity research.

Moreover, the absence of standardized benchmarks and reproducible implementations for many recent QML-based IDS approaches further complicates fair comparison. In contrast, this work emphasizes reproducibility by providing a fully open-source implementation and controlled experimental design.

VI. LIMITATIONS AND NISQ CONSTRAINTS

Despite the encouraging performance achieved by the proposed QML-based IDS, several limitations remain, primarily related to the constraints of current Noisy Intermediate Scale Quantum (NISQ) devices [17]. First, the dimensionality of classical intrusion detection datasets significantly exceeds the number of qubits available on present-day quantum hardware. This necessitates aggressive dimensionality reduction, such as PCA, which may discard fine grained feature interactions relevant to certain attack classes.

Second, quantum circuits employed in this work are shallow by design to mitigate decoherence and gate noise [18], [19]. While this improves trainability and execution feasibility, it inherently limits the expressive

power of the quantum models, particularly for complex multi-class intrusion patterns. Deeper variational circuits or more expressive embeddings could potentially yield better performance but are currently impractical on NISQ devices.

Third, the experiments are conducted using high fidelity quantum simulators rather than real quantum hardware. Although simulators enable controlled evaluation and scalability analysis, they do not fully capture hardware specific noise sources, such as crosstalk, readout errors, and device calibration drift [20]. As a result, performance on actual quantum processors may degrade without additional error mitigation strategies [21].

Finally, training hybrid quantum–classical models remains computationally expensive due to repeated circuit evaluations during optimization [22]. This limits large scale hyperparameter exploration and restricts the applicability of QML models in real-time intrusion detection scenarios under current hardware constraints.

From a scalability perspective, the computational complexity of QSVM kernel estimation grows quadratically with the number of training samples, while variational circuit training in QNNs requires iterative evaluations that scale with both circuit depth and dataset size. These constraints highlight the need for more efficient quantum kernel approximations, circuit compression techniques, and hybrid acceleration strategies to enable large-scale deployment.

Overall, these limitations reflect the current state of quantum technology rather than fundamental shortcomings of the proposed methodology.

In addition, the present study focuses exclusively on supervised hybrid quantum learning paradigms (QNN and QSVM) under a unified experimental framework. Alternative quantum intrusion detection approaches, such as hybrid quantum autoencoders and quantum data description methods, follow fundamentally different unsupervised or reconstruction-based learning strategies. As a result, direct comparison within the same pipeline would require significant modifications to the training objectives and evaluation protocols. Establishing standardized and reproducible benchmarking across heterogeneous QML paradigms remains an open challenge in the field.

VII. CONCLUSION AND FUTURE WORK

This paper presented a controlled and reproducible study of hybrid Quantum Machine Learning for intrusion detection, focusing on the integration of Quantum Neural Networks (QNN) and Quantum Support Vector Machines (QSVM) within a unified framework. Unlike prior works that typically evaluate a single quantum model or rely on heterogeneous experimental settings, this study systematically analyzes the interaction between quantum embeddings, dimensionality reduction, and model architectures under consistent preprocessing and evaluation protocols. This design-oriented perspective enables a clearer understanding of the practical trade-offs that govern QML performance in cybersecurity applications.

From a quantum perspective, the study highlights that the choice of data embedding and circuit architecture plays a critical role in determining classification performance under NISQ constraints. The findings suggest that QML models are not intended to replace classical IDS approaches in the near term, but rather to complement them in hybrid pipelines where dimensionality is controlled and interpretability is prioritized.

It is important to note that the evaluation is conducted on the NSL-KDD dataset, which, although widely used, may not fully capture modern network traffic characteristics. Therefore, the reported results should be interpreted as a controlled benchmark rather than a definitive assessment of real-world deployment performance.

Future work will extend this study in several important directions. First, a comprehensive comparative evaluation including heterogeneous quantum intrusion detection models, such as hybrid quantum autoencoders and quantum data description approaches, will be conducted under a unified benchmarking framework. Second, experiments will be extended to more recent and large-scale cybersecurity datasets to evaluate real-world applicability. Third, we will investigate quantum-native dimensionality reduction and feature encoding techniques to reduce reliance on classical preprocessing. Finally, the integration of noise-aware training, error mitigation strategies, and execution on real quantum hardware will be explored to bridge the gap between simulation and deployment.

As quantum hardware continues to mature, the proposed methodology provides a reproducible and extensible benchmark for assessing the realistic potential of Quantum Machine Learning in cybersecurity applications.

Overall, this work provides a reproducible and systematically controlled foundation for future research in quantum-enhanced intrusion detection, with particular emphasis on understanding the practical trade-offs of QML under realistic NISQ constraints.

REFERENCES

[1] S. Corli, L. Moro, D. Dragoni, M. Dispenza, and E. Prati, "Quantum machine learning algorithms for anomaly detection: A review," *Future Gener. Comput. Syst.*, vol. 166, May 2025, Art. no. 107632. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167739X2400596X

[2] H. Oh and D. K. Park, "Quantum support vector data description for anomaly detection," *Mach. Learning: Sci. Technol.*, vol. 5, no. 3, Aug. 2024, Art. no. 035052, doi: 10.1088/2632-2153/ad6be8.

[3] M. A. Rasyidi, O. Alhussein, S. Muhaidat, and E. Damiani, "Hybrid quantum-classical autoencoders for unsupervised network intrusion detection," 2025, *arxiv: 2512.05069*.

[4] M. Rath and H. Date, "Quantum data encoding: A comparative analysis of classical-to-quantum mapping techniques and their impact on machine learning accuracy," *EPJ Quantum Technol.*, vol. 11, no. 1, p. 72, Oct. 2024, doi: 10.1140/epjqt/s40507-024-00285-3.

[5] H. Belhadef, H. Benchiheb, and L. Lebdjiri, "Exploring the capabilities and limitations of VQC and QSVC for sentiment analysis on real-world and synthetic datasets," in *New Trends in Database and Information Systems*. Cham, Switzerland: Springer, 2023, pp. 415–424.

[6] M. Tavallaee, E. Bagheri, W. Lu, and A. A. Ghorbani, "A detailed analysis of the KDD CUP 99 data set," in *Proc. IEEE Symp. Comput. Intell. Secur. Defense Appl.*, Jul. 2009, pp. 1–6.

[7] M. Uddin, "Addressing accuracy paradox using enhanced weighted performance metric in machine learning," in *Proc. 6th HCT Inf. Technol. Trends (ITT)*, Ras Al Khaimah, UAE, 2019, pp. 319–324, doi: 10.1109/ITT48889.2019.9075071.
[8] P. B. Udas, R. K. Barik, S. Tripathy, S. K. Pradhan, and H. S. Behera, "Spider: A shallow pca-based intrusion detection and classification system," *Comput. Electr. Eng.*, vol. 100, Mar. 2022, Art. no. 107828.
[9] J. Shlens, "A tutorial on principal component analysis," 2014, *arXiv:1404.1100*.
[10] F. L. Gewers, H. V. Ribeiro, L. E. C. da Rocha, and L. F. Costa, "Principal component analysis: A natural approach to data exploration," 2018, *arXiv:1804.02502*.
[11] M. B. Lehocine and H. Belhadef, "Preprocessing-based approach for prompt intrusion detection in SDN networks," *J. Netw. Syst. Manage.*, vol. 32, no. 4, p. 49, Oct. 2024.
[12] E. Ovalle-Magallanes, D. E. Alvarado-Carrillo, J. G. Avina-Cervantes, I. Cruz-Aceves, and J. Ruiz-Pinales, "Quantum angle encoding with learnable rotation applied to quantum-classical convolutional neural networks," *Appl. Soft Comput.*, vol. 141, Jul. 2023, Art. no. 110307.
[13] D. Ranga, A. Rana, S. Prajapat, P. Kumar, K. Kumar, and A. V. Vasilakos, "Quantum machine learning: Exploring the role of data encoding techniques, challenges, and future directions," *Mathematics*, vol. 12, no. 21, p. 3318, Oct. 2024.
[14] J. Han, N. S. DiBrita, Y. Cho, H. Luo, and T. Patel, "Enqode: Fast amplitude embedding for quantum machine learning using classical data," 2025, *arXiv:2503.14473*.
[15] M. Schuld and N. Killoran, "Quantum machine learning in feature Hilbert spaces," *Phys. Rev. Lett.*, vol. 122, no. 4, Feb. 2019, Art. no. 040504.
[16] R. Shaydulin and S. M. Wild, "Importance of kernel bandwidth in quantum machine learning," *Phys. Rev. A, Gen. Phys.*, vol. 106, no. 4, Oct. 2022, Art. no. 042407.
[17] J. Preskill, "Quantum computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, Aug. 2018.
[18] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, pp. 195–202, Jan. 2017.
[19] M. Cerezo, A. Arrasmith, R. Babbush, S. Benjamin, S. Endo, K. Fujii, J. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. Coles, "Variational quantum algorithms," *Nature Rev. Phys.*, vol. 3, pp. 625–644, Mar. 2021.
[20] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, no. 7671, pp. 242–246, Sep. 2017.
[21] K. Temme, S. Bravyi, and J. M. Gambetta, "Error mitigation for short-depth quantum circuits," *Phys. Rev. Lett.*, vol. 119, no. 18, Nov. 2017, Art. no. 180509.
[22] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe, "Circuit-centric quantum classifiers," *Phys. Rev. A, Gen. Phys.*, vol. 101, no. 3, Mar. 2020, Art. no. 032430.

RAID ANIS KERKATOU received the B.S. degree in computer science and the M.Sc. degree in artificial intelligence and its application from the University of Mujahid Abdelhafid Boussouf, Mila, Algeria. He is currently pursuing the Ph.D. degree in artificial intelligence, with a research focus on quantum machine learning (QML) and cybersecurity.

His research interests include on the development of hybrid quantum-classical learning frameworks, including quantum neural networks (QNNs), quantum support vector machines (QSVMs), and variational quantum algorithms, for advanced intrusion detection systems and network security applications. He is particularly interested in leveraging quantum feature representations and optimization techniques to enhance detection performance in high-dimensional and complex data environments. In addition, he investigates quantum reinforcement learning (QRL) for adaptive and intelligent cyber-defense strategies. His work also extends to computer vision applications for disaster management, including semantic segmentation and object detection for post-disaster damage assessment using datasets, such as FloodNet. His broader research interests include quantum optimization, dimensionality reduction techniques, and the integration of emerging artificial intelligence paradigms into real-world, and data-driven cybersecurity systems.

HACENE BELHADEF received the Ph.D. degree in computer science from the Faculty of Engineering, University Fréres Mentouri Constantine 1, in 2010, and the HDR-Habilitation degree, in 2013.

He was a Lecturer. He was a Full Professor, in 2019. He is currently a Professor with the Faculty of New Information and Communication Technologies, University of Constantine 2-Abdelhamid Mehri, where he teaches undergraduate and graduate courses, supervises engineer, and master's and Ph.D. students. Previously, he was the Head of Pedagogy and a Postgraduate Studies, where he was responsible for curriculum planning, quality assurance of academic programs, and coordination of graduate education. More recently, he was the Head of the Digital Transformation Department, contributing to strategic initiatives related to digital infrastructure, academic innovation, and institutional modernization. His academic activities include course development, curriculum coordination, student mentoring, and active participation in institutional and funded research projects. Through his research, teaching, leadership, and service activities, he contributes to the advancement of higher education and scientific research at both national and international levels. In addition to his teaching and research responsibilities, he has held several academic and administrative leadership positions within the university. His research interests include quantum computing and quantum machine learning, as well as data mining, natural language processing, artificial intelligence, and machine learning. His research addresses both theoretical foundations and applied solutions in advanced computational systems, with a particular emphasis on emerging quantum technologies.

Dr. Belhadef is an Active Member of the international scientific community. He regularly serves on program committees of international conferences and acts as a reviewer for several peer-reviewed international journals.

AICHA EUTAMENE received the M.S. degree in computer science from University Fréres Mentouri Constantine 1, Algeria, in 2009, the Ph.D. degree in computer science from the University of Constantine 2-Abdelhamid Mehri, Algeria, in 2016, and the HDR-Habilitation degree, in February 2026.

She was a Lecturer. She is currently a Lecturer with the Faculty of Natural and Life Sciences, University Fréres Mentouri Constantine 1. Her research interests include natural language processing, machine learning, bioinformatics, quantum computing, and applications of artificial intelligence in computational biology.

SVETLANA PETROVA STEFANOVA received the M.Sc. degree in computer science and the Ph.D. degree from the University of Ruse "Angel Kanchev," Ruse, Bulgaria, in 1989 and 2003, respectively.

Since 1990, she has been a Teacher with the Department of Electrical and Computer Engineering, University of Ruse "Angel Kanchev." She has been a CEO of FAAC Bulgaria, Ruse, since 1995. She has been an Associate Professor, since 2011. She has many years of experience in the industry. Her research interests include machine learning, intelligent systems, and computer graphics systems. She has published numerous articles in these areas.

Prof. Stefanova is a member of ACM.