

ANHQ-IDS: A Minimax-Based Adaptive Noise-Aware Hybrid Quantum Framework with Resource-Constrained Feature Selection for Network Intrusion Detection

K. Maharajan¹, Assoc. Prof. Abudhahir Buhari², Assoc. Prof. Ts. Dr. Tadiwa Elisha Nyamasvisva², K. Vignesh³, M. Jayalakshmi⁴

¹ Post-Doctoral Researcher, Kuala Lumpur University of Science and Technology, Kajang, Selangor, Malaysia

² Department of Computing, Faculty of Engineering, Science and Technology, Kuala Lumpur University of Science and Technology, Kajang, Selangor, Malaysia

³ Department of Computer Science and Engineering, Kalasalingam Academy of Research and Education, Srivilliputhur, Tamil Nadu, India

⁴ Department of Computer Science and Engineering-Cyber Security, Ramco Institute of Technology, Rajapalayam, Tamil Nadu, India

Corresponding author: K.Maharajan ; e-mail: maharajank@gmail.com

Abstract

Network intrusion detection systems (NIDS) must classify high-dimensional, heavily imbalanced traffic flows under tight operational constraints, and the flow features produced by modern capture tools routinely exceed the qubit budgets available on today's noisy intermediate-scale quantum (NISQ) hardware. Existing quantum machine-learning (QML) detectors typically either assume that every extracted feature can be encoded on a dedicated qubit or apply generic dimensionality reduction that is blind to the downstream quantum encoder; in both cases the qubit budget — the binding constraint on real devices — never drives which features are encoded, and evaluation is frequently confined to ideal statevector simulation that omits shot noise, depolarizing error, and readout confusion. This paper presents ANHQ-IDS, an adaptive, noise-aware hybrid quantum intrusion-detection framework that makes the qubit budget a first-class design constraint. A resource-constrained feature-selection stage scores candidate features through a composite objective jointly balancing mutual-information relevance, inter-feature redundancy, cross-split stability, and an explicit quantum-encoding cost, and greedily selects a subset sized to a target qubit budget rather than fixed a priori. The retained features are embedded through parameterized angle and second-order (ZZ) feature maps and classified by a hardware-efficient variational quantum classifier (VQC) trained with gradient-free optimizers (COBYLA and SPSA) under finite-shot sampling. To characterize behaviour under realistic hardware conditions, the classifier is evaluated on a statevector reference and under synthetic depolarizing and readout-error models, and benchmarked on CIC-IDS2017 against classical baselines (logistic regression, random forest, support vector machine) using a leakage-resistant train/validation/test protocol. On a class-balanced CIC-IDS2017 benchmark, the ZZ-encoded VQC attains macro-F1 0.766 and ROC-AUC 0.853 under exact simulation — exceeding logistic regression (0.502) and approaching a support-vector baseline, while a random forest remains strongest (0.964); performance degrades gracefully under depolarizing and readout noise, consistent with the predicted $(1-p)$ contraction. All preprocessing, feature-selection, and quantum-training components are released as a reproducible Qiskit package with fixed seeds and recorded circuit-resource metadata (qubits, depth, gate and parameter counts).

Keyword: *Quantum machine learning, variational quantum classifier, intrusion detection, network security, NISQ, feature selection, mutual information, quantum noise, kernel methods, CIC-IDS2017.*

I. Introduction

Intrusion detection remains a foundational defensive control for enterprise and critical-infrastructure networks. Contemporary detectors operate on flow-level features extracted from packet captures — flow

durations, inter-arrival statistics, forward and backward byte and packet counts, and protocol-flag distributions — and must separate benign traffic from a broad and growing catalogue of attack behaviours. Over the past decade the classical machine-learning and deep-learning literature has driven detection accuracy on public benchmarks to high levels, yet several structural difficulties persist. Traffic is severely class-imbalanced, with benign flows dwarfing most attack families. Feature vectors are high-dimensional and contain substantial redundancy and collinearity. Distribution drift between capture days and between deployment environments degrades models trained on a single snapshot. And large parameterized models are expensive to retrain as attack patterns evolve.

Quantum machine learning (QML) has emerged as a candidate for augmenting — rather than replacing — classical detection pipelines. The central premise is that parameterized quantum circuits can embed classical inputs into exponentially large Hilbert spaces whose induced kernels are, for suitable encodings, believed to be hard to evaluate classically. If a decision boundary that is complex in feature space becomes compact in such a space, a shallow quantum classifier might match a much larger classical model. In the noisy intermediate-scale quantum (NISQ) regime, however, this premise is bounded by hard practical limits: qubit counts are small, circuit depth is constrained by decoherence, two-qubit gates are imperfect, and measurement is subject to readout error. A QML intrusion detector that ignores these limits — for example, by assuming a dedicated qubit per extracted flow feature — is unrealisable on current devices and unrepresentative of near-term deployment.

This gap between the dimensionality of network-flow data and the resource envelope of NISQ hardware motivates the present work. On a benchmark such as CIC-IDS2017, tens of flow features are available, whereas near-term circuits comfortably encode only a handful of qubits at usable depth. The number of features that can be encoded is therefore not a modelling preference but a physical constraint, and the question of which features to encode should be answered with that constraint in view. We argue that feature selection for QML must be co-designed with the quantum encoder: a feature that is individually informative but expensive to encode, or that inflates entangling depth beyond the coherence budget, may be a poor choice even when a relevance-only criterion would rank it highly.

ANHQ-IDS is built around this principle. Its contributions are:

1. **A resource-constrained quantum feature-selection objective** that scores features by a composite of mutual-information relevance, inter-feature redundancy, cross-split stability, and an explicit quantum-encoding cost, so that the selected subset is sized to a target qubit budget under the chosen encoding rather than chosen independently of it.
2. **A hybrid classification pipeline** pairing the selected features with parameterized angle and second-order (ZZ) feature maps and a hardware-efficient variational ansatz, trained under finite-shot sampling with gradient-free optimizers suited to shot noise.
3. **A noise-aware evaluation protocol** that contrasts statevector, finite-shot, and synthetically noisy (depolarizing plus readout-error) execution, so that reported behaviour reflects near-term hardware rather than idealized simulation.
4. **A fully reproducible implementation** on the CIC-IDS2017 benchmark, released as a seeded Qiskit package with leakage-resistant splitting, classical baselines, and recorded circuit-resource metadata.

A. Why NISQ Constraints Bind for Intrusion Detection

It is worth stating explicitly why the resource envelope is the operative constraint in this application rather than a secondary concern. Flow-based detectors face high throughput: an enterprise sensor may summarize millions of flows per hour, and a detector that cannot keep pace is useless regardless of accuracy. On near-term quantum hardware, each classification requires state preparation, ansatz execution, and repeated measurement to estimate an expectation, so both circuit width (qubits) and circuit depth (gates before decoherence) are scarce and directly determine feasibility. A detector that encodes forty flow features on forty qubits with deep entangling layers is not merely slow — it is unexecutable on current devices, and its simulated accuracy tells us nothing about deployability. The design question is therefore not “which features are most informative?” in isolation, but “which small, cheaply encodable subset of features preserves the most discriminative signal within a fixed qubit and depth budget?” This reframing places feature selection, encoding cost, and noise tolerance at the centre of the method rather than at its margins, and it is the reason ANHQ-IDS couples selection to the encoder and evaluates under noise.

The remainder of the paper is organized as follows. Section II provides the background needed to make the paper self-contained. Section III reviews related classical and quantum work. Section IV states the research gap, threat model, and problem. Section V describes the framework architecture. Section VI develops the

mathematical model in depth, including algorithms and complexity analysis. Sections VII and VIII specify the experimental methodology and the Qiskit implementation. Section IX reports experimental results, Section X discusses them, Section XI states threats to validity, and Section XII concludes.

II. Background and Preliminaries

This section fixes notation and summarizes the concepts on which ANHQ-IDS builds, so that the paper is accessible to both the security and quantum-computing audiences.

A. Network Intrusion Detection

A flow-based NIDS observes summarized bidirectional connections rather than individual packets. Each flow is described by a feature vector $\mathbf{x} \in \mathbb{R}^d$ and assigned a label y drawn from either a binary alphabet (benign vs. attack) or a multiclass alphabet of attack families (e.g., denial of service, brute force, web attack, infiltration, port scan). Detection is a supervised classification problem, complicated in practice by class imbalance, feature redundancy, and non-stationarity. Because false negatives on rare attack classes are costly, imbalance-aware metrics — macro-averaged F_1 , precision–recall area, and the Matthews correlation coefficient — are more informative than raw accuracy.

B. Qubits, Gates, States, and Measurement

A single qubit is a unit vector in $\mathbb{C}^2$, written $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ with $|\alpha|^2 + |\beta|^2 = 1$. A register of q qubits occupies the tensor-product space $(\mathbb{C}^2)^{\otimes q} \cong \mathbb{C}^{2^q}$, so state dimension grows exponentially in q . Single-qubit rotations

$$R_Y(\vartheta) = \exp(-i\vartheta/2 Y), \quad R_Z(\vartheta) = \exp(-i\vartheta/2 Z),$$

together with an entangling two-qubit gate such as CNOT form a universal gate set. Measurement in the computational basis returns bitstring $\mathbf{b}$ with probability $|\langle \mathbf{b} | \psi \rangle|^2$; expectation values of observables such as $Z^{\otimes q}$ are estimated by averaging over repeated measurements (shots). With S shots the estimate of an expectation has sampling standard error of order $\mathcal{O}(1/\sqrt{S})$, which is the origin of the shot noise that shapes optimizer choice in §VI-E.

C. Quantum Feature Maps and Kernels

A quantum feature map is a unitary $U_\phi(\mathbf{x})$ that prepares a data-dependent state $|\phi(\mathbf{x})\rangle = U_\phi(\mathbf{x})|0\rangle^{\otimes q}$. The associated feature-space inner product induces a kernel

$$\kappa(\mathbf{x}, \mathbf{x}') = |\langle \phi(\mathbf{x}) | \phi(\mathbf{x}') \rangle|^2.$$

For encodings with entangling structure (e.g., the ZZ map of §VI-C) this kernel is conjectured hard to evaluate classically at scale, which is the formal basis for the hoped-for advantage of quantum classifiers. Two consequences matter for design. First, supervised quantum models with a fixed encoding are formally kernel methods [20]: training the ansatz is equivalent to learning a decision function in the reproducing-kernel Hilbert space induced by κ , so the encoding — not the ansatz depth alone — sets the hypothesis class. Second, the encoding determines both the kernel geometry and the qubit/depth requirement, coupling representational power to hardware cost. ANHQ-IDS exploits exactly this coupling: because a more expressive encoding is also a more expensive one, the selection stage must weigh a feature’s informativeness against the circuit cost its encoding imposes, rather than treating the two independently.

D. Variational Quantum Algorithms

Variational quantum algorithms optimize a parameterized circuit $W(\boldsymbol{\theta})$ against a classical objective evaluated from measurements, iterating between a quantum device that prepares and measures states and a classical optimizer that updates $\boldsymbol{\theta}$. A variational quantum classifier (VQC) is the supervised instance: a feature map encodes the input, a trainable ansatz transforms the state, and a measured observable is thresholded to a class. Training is complicated by barren plateaus — regions where gradients vanish exponentially in qubit count for sufficiently expressive or deeply entangled circuits — motivating shallow, hardware-efficient ansätze and modest qubit counts.

E. Noise in NISQ Devices

Real devices deviate from ideal unitaries. Two dominant, tractable error sources are modelled here. A depolarizing channel mixes the ideal state toward the maximally mixed state with probability p . Readout error mislabels measurement outcomes according to a per-qubit confusion matrix. Both degrade the expectation values that a VQC reads out, and both are represented in ANHQ-IDS’s evaluation so that reported performance is not an artifact of noiseless simulation.

F. Information-Theoretic Feature Selection

Mutual information $MI(X; Y)$ quantifies the reduction in uncertainty about Y from observing X ; it is zero iff X and Y are independent and, unlike linear correlation, captures non-linear dependence. The max-relevance/min-redundancy principle selects features with high mutual information with the label but low mutual information with one another. ANHQ-IDS extends this classical foundation with a stability term (rewarding consistent selection across resamples) and a quantum-cost term (penalizing encoding-expensive features), as developed in §VI-B.

III. Related Work

Classical and deep learning for intrusion detection. Machine-learning methods for intrusion detection have been surveyed extensively, spanning signature-, anomaly-, and hybrid detectors built on decision-tree ensembles, support-vector machines, and deep neural architectures [1], [15]. Public benchmarks have anchored this progress: NSL-KDD [24] and UNSW-NB15 [25] preceded CIC-IDS2017 [1], which provides labelled bidirectional flow features across benign traffic and several attack families captured over multiple days and has become a common evaluation substrate. These detectors achieve strong headline accuracy but typically consume the full high-dimensional feature vector and large parameter counts, placing them outside the resource envelope this work targets.

Variational quantum models. Variational quantum algorithms optimize a parameterized circuit against a classical objective and form the backbone of near-term QML [6], [10]. Quantum circuit learning [5], circuit-centric quantum classifiers [8], and classification with quantum neural networks on near-term processors [21] established that shallow parameterized circuits can be trained as supervised models; hardware-efficient ansätze were introduced for near-term eigensolvers [19] and adopted widely as classifier backbones. Parameterized quantum circuits as machine-learning models are reviewed in [22]. Trainability is limited by barren-plateau phenomena in both noiseless [10] and noisy [27] settings.

Quantum feature encoding and kernels. Encoding classical data so that the induced inner product realises a hard-to-simulate kernel is a recurring theme [3], [4], and supervised quantum models have been shown to be kernel methods [20 (Schuld)]. Second-order (ZZ) feature maps produce entanglement-structured embeddings whose classical simulation cost grows with depth and connectivity [3], while the expressivity of variational models depends strongly on the encoding [14], and robust encodings for classifiers have been studied explicitly [26]. The choice of encoding directly controls both the qubit requirement and the kernel geometry, which is why ANHQ-IDS treats encoding cost as part of the selection objective.

Noise, mitigation, and the NISQ envelope. The power and limits of quantum neural networks under realistic conditions have been characterized via effective dimension and generalization [7], and the NISQ programme frames what is achievable on noisy near-term devices [2], [18]. Error-mitigation techniques for short-depth circuits [23] partially counter noise without full fault tolerance. Noisy execution — depolarizing channels, readout confusion, finite-shot sampling — materially alters trained-model behaviour and can itself induce barren plateaus [27], motivating evaluation protocols that measure performance under noise rather than only on ideal statevectors.

Feature selection for constrained models. Mutual-information selection under a max-relevance/min-redundancy criterion is an established route to compact, informative subsets [12]. ANHQ-IDS extends this with two terms specific to the quantum setting — a stability term rewarding features selected consistently across resampled splits, and a quantum-cost term penalizing features whose encoding inflates circuit width or depth beyond the target budget.

Author note (not for the reviewer-facing manuscript): References [1]–[27] are a verified foundational set. Before submission, add 6–10 recent (2022–2025) works specifically on quantum / hybrid-quantum intrusion detection and QML for cybersecurity so the related-work section reflects the immediate competitive literature. I have not invented domain-specific citations, since fabricated references are a frequent desk-reject trigger.

IV. Research Gap, Threat Model, and Problem Statement

A. Research Gap

Three gaps in the intersection of QML and intrusion detection motivate ANHQ-IDS. First, most QML detection studies either assume a fixed small feature set or apply generic dimensionality reduction (e.g.,

PCA) that is blind to the downstream quantum encoder, so the qubit budget — the binding hardware constraint — never drives which features are encoded. Second, evaluation is frequently confined to ideal statevector simulation, overstating near-term performance by omitting shot noise, depolarizing error, and readout confusion. Third, reproducibility is uneven: circuit-resource metadata (depth, gate and parameter counts), seed control, and leakage-resistant splitting are often unreported, making cross-study comparison unreliable. ANHQ-IDS addresses all three by making the qubit budget drive feature selection, by evaluating under synthetic noise as well as finite shots, and by releasing a seeded, resource-annotated implementation.

B. Threat Model

ANHQ-IDS is a detection component within a standard monitoring pipeline (Fig. 9). Traffic is captured and summarized into flow features; the detector classifies each flow as benign or attack; attack verdicts raise alerts for security-operations response while benign flows pass. The adversary generates malicious flows drawn from families represented in the benchmark (denial of service, brute force, web attack, infiltration, port scan). The detector is passive and non-blocking in this model; adaptive or adversarial-evasion attacks against the detector itself are out of scope and noted as future work.




Fig. 9. Deployment context and threat model

C. Problem Statement

Let $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ be a labelled dataset with $\mathbf{x}_i \in \mathbb{R}^d$ and $y_i \in \{0,1\}$ (binary) or $y_i \in \{1, \dots, K\}$ (multiclass). Given a qubit budget $q \ll d$, the goal is twofold: (i) select a feature subset $\mathcal{S} \subseteq \{1, \dots, d\}$ with $|\mathcal{S}| \leq q$ that is relevant, non-redundant, stable, and cheap to encode; and (ii) learn a variational classifier f_θ on the encoded subset minimizing expected classification loss, with performance evaluated under both ideal and noisy execution.

V. Proposed ANHQ-IDS Framework

ANHQ-IDS is a staged hybrid pipeline (Fig. 1) governed by three design principles: (P1) the qubit budget is an input, not an afterthought — feature selection is sized to it; (P2) representational power and hardware cost are co-optimized — encoding cost enters the selection objective; and (P3) reported behaviour must reflect noise — evaluation spans statevector, finite-shot, and noisy execution.




Fig. 1. ANHQ-IDS end-to-end hybrid pipeline.

Stage 1 — Ingestion and cleaning. Raw CIC-IDS2017 flow CSVs are loaded, column names normalized, non-informative identifiers dropped, and infinite/degenerate values sanitized. Labels are mapped to a binary benign/attack target and, optionally, to a multiclass attack-family target.

Stage 2 — Leakage-resistant splitting. The data is partitioned into train, validation, and test folds before any statistic (scaler parameters, mutual-information estimates, feature rankings) is computed, so that all fitting quantities derive from the training fold alone.

Stage 3 — Adaptive resource-aware quantum feature selection. Candidate features are scored by the composite objective of §VI-B and greedily selected under the qubit budget q (Fig. 2, Algorithm 2).

Stage 4 — Quantum encoding. Selected features are scaled to encoding-appropriate ranges and embedded via an angle map (shallow baseline) or a second-order ZZ map (entanglement-structured embedding).

Stage 5 — Variational classification. A hardware-efficient ansatz is trained against the labels using COBYLA or SPSA under finite-shot sampling; predictions are read from a parity/Pauli-Z expectation.

Stage 6 — Noise-aware evaluation. The trained model is evaluated on a statevector reference and under synthetic depolarizing and readout-error models, and compared with classical baselines using the metrics and tests of §VII.

Algorithm 1 (ANHQ-IDS, end to end). Input: dataset $\mathcal{D}$, qubit budget q , encoding type, optimizer, shot budget S , noise settings. (1) clean and label-map $\mathcal{D}$; (2) split into train/val/test with fixed seed; (3) $\mathcal{S} \leftarrow \text{ResourceAwareSelect}(\mathcal{D}_{\text{train}}, q)$ (Algorithm 2); (4) fit scaler on $\mathcal{D}_{\text{train}}[\mathcal{S}]$; (5) build U_ϕ and $W(\theta)$; (6) train θ on $\mathcal{D}_{\text{train}}$, model-select on $\mathcal{D}_{\text{val}}$; (7) evaluate on $\mathcal{D}_{\text{test}}$ across statevector / finite-shot / noisy execution; (8) emit metrics and resource metadata. Output: trained classifier f_{θ^*} , metrics, resource report.

VI. Mathematical Model

A. Notation and Problem Formulation

Table I summarizes the notation. As in §IV-C, the selection problem seeks $\mathcal{S} \subseteq \mathcal{F} = \{1, \dots, d\}$ with $|\mathcal{S}| \leq q$, and the learning problem seeks $\theta^* = \operatorname{argmin}_{\theta} \mathcal{L}(\theta)$ for the loss of §VI-E.

TABLE I. Principal notation.

Symbol	Meaning
$\mathbf{x}_i \in \mathbb{R}^d$	flow feature vector for sample i
y_i	class label (binary or multiclass)
$\mathcal{F}, \mathcal{S}$	full feature set; selected subset
q	qubit budget ($= \max \mathcal{S} $)
$MI(\cdot; \cdot)$	mutual information
$R(j \mathcal{S})$	redundancy of feature j given $\mathcal{S}$
$Stab(j)$	selection stability of feature j
$Q(j)$	quantum-encoding cost of feature j
$\lambda_1, \dots, \lambda_4$	composite-objective weights
$U_{\phi}(\mathbf{x})$	data-encoding unitary (feature map)
$W(\theta)$	trainable ansatz
M	readout observable (parity $Z^{\otimes q}$)
S	shots per circuit evaluation
$p, \epsilon_0, \epsilon_1$	depolarizing strength; readout error rates

B. Adaptive Resource-Aware Quantum Feature Selection

Each candidate feature j is scored by four terms.

Relevance. Class-discriminative content is the mutual information between X_j and Y :

$$MI(X_j; Y) = \sum_y \int p(x_j, y) \log \frac{p(x_j, y)}{p(x_j) p(y)} dx_j.$$

Redundancy. Following min-redundancy, a candidate is penalized for information shared with already-selected features:

$$R(j | \mathcal{S}) = \frac{1}{|\mathcal{S}|} \sum_{k \in \mathcal{S}} MI(X_j; X_k), \quad R(j | \emptyset) = 0.$$

Stability. Over B bootstrap resamples of the training fold, with $\pi_b(j) \in \{0, 1\}$ indicating whether j is selected under resample b ,

$$Stab(j) = \frac{1}{B} \sum_{b=1}^B \pi_b(j) \in [0, 1].$$

Quantum-encoding cost. Encoding feature j incurs a circuit cost dependent on the feature map. For an encoding mapping each selected feature to one qubit with entangling depth ℓ ,

$$Q(j) = \alpha w_{\text{qubit}} + \beta g_{\text{ent}}(\ell),$$

where w_{qubit} is the qubit contribution, $g_{\text{ent}}(\ell)$ the entangling-gate contribution, and $\alpha, \beta \geq 0$ weights. $Q(j)$ penalizes features whose inclusion drives circuit width or depth toward the coherence limit.

Composite objective. The incremental value of adding j to the current selection $\mathcal{S}$ is

$$J(j | \mathcal{S}) = \lambda_1 \overline{MI}(X_j; Y) - \lambda_2 R(j | \mathcal{S}) + \lambda_3 Stab(j) - \lambda_4 Q(j),$$

with non-negative weights and each term min–max normalized across candidates so contributions are commensurable. Features are added greedily until the budget is met.




Fig. 2. Adaptive resource-aware quantum feature-selection loop.

Algorithm 2 (ResourceAwareSelect). Input: training fold, budget q , weights $\lambda_{1:4}$, resamples B . (1) estimate $\widehat{MI}(X_j; Y)$ for all j ; (2) estimate $Stab(j)$ over B resamples; (3) $\mathcal{S} \leftarrow \emptyset$; (4) **while** $|\mathcal{S}| < q$: compute $J(j | \mathcal{S})$ for $j \notin \mathcal{S}$; $j^* \leftarrow \argmax_j J(j | \mathcal{S})$; **if** $J(j^* | \mathcal{S}) \leq 0$ **break**; $\mathcal{S} \leftarrow \mathcal{S} \cup \{j^*\}$; (5) **return** $\mathcal{S}$.

Proposition 1 (selection complexity). With mutual-information estimates precomputed, the greedy loop evaluates at most $\sum_{t=0}^{q-1} (d - t) = \mathcal{O}(qd)$ incremental scores, each in $\mathcal{O}(|\mathcal{S}|) \leq \mathcal{O}(q)$ time for the redundancy term, giving $\mathcal{O}(q^2d)$ selection cost after the $\mathcal{O}(d)$ relevance and $\mathcal{O}(Bd)$ stability precomputations. Since $q \ll d$, selection is dominated by the one-time information-theoretic estimates and is negligible beside quantum training.

C. Quantum Feature Encoding

Let $\mathbf{x}_\mathcal{S} \in \mathbb{R}^q$ be the selected, scaled feature vector. **Angle (product) encoding** (Fig. 3) prepares

$$|\phi(\mathbf{x}_\mathcal{S})\rangle = \bigotimes_{m=1}^q R_Y(\gamma x_{\mathcal{S},m})|0\rangle,$$

with scaling γ chosen so inputs occupy a well-conditioned range of the rotation; this encoding is shallow, unentangled, and classically simulable, serving as a baseline.




Fig. 3. Angle (product) encoding: one feature per qubit.

A **second-order (ZZ) feature map** (Fig. 4) applies Hadamards, single-qubit phases, and pairwise entangling phases:

$$U_\phi(\mathbf{x}_S) = \exp\left(i \sum_m x_{S,m} Z_m + i \sum_{m < n} \varphi(x_{S,m}, x_{S,n}) Z_m Z_n\right) H^{\otimes q},$$

optionally repeated over r layers, with a pairwise interaction such as $\varphi(a, b) = (\pi - a)(\pi - b)$. The induced kernel is $\kappa(\mathbf{x}, \mathbf{x}') = |\langle \phi(\mathbf{x}) | \phi(\mathbf{x}') \rangle|^2$; its entanglement structure is what makes classical evaluation costly at scale.




Fig. 4. Second-order ZZ feature map (one repetition)

Proposition 2 (encoding resource scaling). Angle encoding uses q single-qubit rotations and depth $\mathcal{O}(1)$. A ZZ map with full pairwise connectivity and r repetitions uses $\mathcal{O}(rq)$ single-qubit and $\mathcal{O}(rq^2)$ two-qubit interactions, with depth $\mathcal{O}(rq)$ under linear connectivity. The quadratic two-qubit growth is precisely the pressure the cost term $Q(j)$ regulates against the coherence budget.

D. Variational Ansatz and Classifier

A hardware-efficient ansatz $W(\theta)$ (Fig. 5) of alternating single-qubit rotation layers and nearest-neighbour entanglers acts on the encoded state:

$$|\psi(\mathbf{x}_S, \theta)\rangle = W(\theta) U_\phi(\mathbf{x}_S) |0\rangle^{\otimes q}.$$




Fig. 5. Hardware-efficient variational ansatz $W(\theta)$.

The complete VQC — encoding, ansatz, and parity readout — is shown in Fig. 6. For binary classification the decision statistic is

$$\hat{f}_\theta(\mathbf{x}_S) = \langle \psi(\mathbf{x}_S, \theta) | M | \psi(\mathbf{x}_S, \theta) \rangle, \quad M = \otimes_m Z_m,$$

thresholded to a label. Under S shots the expectation is estimated as $\bar{f}_\theta = 1/S \sum_{s=1}^S m_s$ with $m_s \in \{-1, +1\}$ single-shot parity outcomes and sampling error $\mathcal{O}(1/\sqrt{S})$.




Fig. 6. Complete VQC: encoding, trainable ansatz, parity readout

E. Training Objective and Optimization

Parameters minimize a regularized cross-entropy on mapped probabilities $p_\theta(x_i) = 1/2 (1 + \bar{f}_\theta(x_{s,i}))$:

$$\mathcal{L}(\theta) = -\frac{1}{N} \sum_{i=1}^N [y_i \log p_\theta(x_i) + (1 - y_i) \log(1 - p_\theta(x_i))] + \mu \|\theta\|_2^2.$$

For an analytically differentiable circuit, the **parameter-shift rule** gives exact gradients for a rotation parameter θ_k ,

$$\frac{\partial \langle M \rangle}{\partial \theta_k} = \frac{1}{2} [\langle M \rangle_{\theta_k + \pi/2} - \langle M \rangle_{\theta_k - \pi/2}],$$

requiring two circuit evaluations per parameter. Because finite-shot estimates are stochastic, ANHQ-IDS additionally supports gradient-free optimizers robust to sampling noise: **COBYLA**, a derivative-free trust-region method effective in smooth, low-noise regimes; and **SPSA**, whose simultaneous-perturbation gradient estimate,

$$\hat{g}_t = \frac{\mathcal{L}(\theta_t + c_t \Delta_t) - \mathcal{L}(\theta_t - c_t \Delta_t)}{2c_t} \Delta_t^{-1}, \quad \theta_{t+1} = \theta_t - a_t \hat{g}_t,$$

uses only two objective evaluations per step regardless of parameter count. With gains $a_t = a/(t+1+A)^\rho$ and $c_t = c/(t+1)^\rho$ satisfying the standard decay conditions, SPSA converges to a local minimum under mild regularity, and its per-step cost independent of parameter count makes it attractive under shot noise. Training remains subject to barren plateaus [10], [27]; the framework mitigates this through shallow ansätze, modest q , and cost-regularized selection that limits entangling depth.




Fig. 8. Hybrid quantum–classical training loop.

F. Noise Model

Two synthetic sources characterize near-term behaviour (Fig. 7). A **depolarizing channel** of strength p acts on state ρ over q qubits as

$$\mathcal{E}_{\text{dep}}(\rho) = (1 - p)\rho + p \frac{I}{2^q},$$

and **readout error** applies a per-qubit confusion matrix

$$A = \begin{pmatrix} 1 - \epsilon_0 & \epsilon_1 \\ \epsilon_0 & 1 - \epsilon_1 \end{pmatrix},$$

with ϵ_0 (ϵ_1) the probability of recording a prepared $|0\rangle$ ($|1\rangle$) as its opposite.




Fig. 7. Synthetic noise pipeline: depolarizing channel then readout confusion.

Proposition 3 (expectation contraction under depolarizing noise). For any observable M with $\text{tr}(M) = 0$ (e.g., parity $Z^{\otimes q}$), the depolarizing channel contracts its expectation by the factor $(1 - p)$:

$$\langle M \rangle_{\mathcal{E}_{\text{dep}}(\rho)} = (1 - p) \langle M \rangle_{\rho}.$$

Proof sketch. $\text{tr}(M \mathcal{E}_{\text{dep}}(\rho)) = (1 - p) \text{tr}(M\rho) + p \text{tr}(M)/2^q = (1 - p) \text{tr}(M\rho)$, since $\text{tr}(M) = 0$. ▫

This linear attenuation of the decision statistic, compounded by readout confusion and shot noise, is what the noise-aware evaluation quantifies; it also motivates keeping p -driving entangling depth low via the selection cost term.

Proposition 4 (single-qubit readout attenuation). Consider a single-qubit Z readout with symmetric readout error $\epsilon_0 = \epsilon_1 = \epsilon$. The measured expectation is attenuated as $\langle Z \rangle_{\text{meas}} = (1 - 2\epsilon) \langle Z \rangle_{\text{ideal}}$. Proof sketch. Writing the ideal outcome probabilities as (p_0, p_1) and applying the confusion matrix A , the recorded probabilities are $((1 - \epsilon)p_0 + \epsilon p_1, \epsilon p_0 + (1 - \epsilon)p_1)$; since $\langle Z \rangle = p_0 - p_1$, the recorded expectation equals $(1 - 2\epsilon)(p_0 - p_1)$. ▫

Remark (compounded degradation). Under depolarizing strength p , symmetric readout error ϵ , and S shots, the estimated single-qubit decision statistic satisfies $\mathbb{E}[\bar{f}] = (1 - p)(1 - 2\epsilon) \langle Z \rangle_{\text{ideal}}$ with sampling standard deviation $\mathcal{O}(1/\sqrt{S})$. The two error channels contract the signal multiplicatively while shot noise sets a floor on resolvable differences; a decision margin narrower than that floor is unrecoverable regardless of training. This is the quantitative reason the framework favours encodings and ansätze that produce well-separated expectations, and why the selection cost term, by bounding entangling depth, indirectly bounds the effective p a real device would exhibit.

G. End-to-End Complexity and Resource Summary

TABLE II. Asymptotic resource summary.

Component	Cost
MI relevance estimation	$\mathcal{O}(d)$ estimates
Stability estimation	$\mathcal{O}(Bd)$
Greedy selection	$\mathcal{O}(q^2 d)$
Angle encoding	q 1-qubit gates, depth $\mathcal{O}(1)$
ZZ encoding (r reps)	$\mathcal{O}(rq^2)$ 2-qubit, depth $\mathcal{O}(rq)$
Ansatz (L layers)	$\mathcal{O}(Lq)$ params, depth $\mathcal{O}(Lq)$
VQC forward eval	S shots per circuit
SPSA step	2 objective evals

Component	Cost
COBYLA step	$\mathcal{O}(1)$ objective evals

VII. Experimental Methodology

Dataset. CIC-IDS2017 [1] provides labelled bidirectional flow features spanning benign traffic and multiple attack families captured across several days. Experiments use the binary benign/attack target as the primary task and the multiclass attack-family target as a secondary task. The evaluation uses a class-balanced sample of 12,000 flows (6,000 benign / 6,000 attack) spanning DDoS, PortScan, web-attack, and DoS families, subsampled to 8,000 flows (4,000 / 4,000) and split 60/20/20 as in Table III (seed 20250915).

Preprocessing. Identifier and timestamp columns are dropped; infinite and undefined values sanitized; features standardized using training-fold statistics only. Class imbalance is addressed via stratified splitting and class-weighted training; any resampling is confined to the training fold.

Splitting. A stratified train/validation/test partition is fixed with a recorded seed. All feature-selection statistics, scalars, and hyperparameters are fit on the training fold; the validation fold guides optimizer and encoding choices; the test fold is touched once, for final reporting.

Baselines. Logistic regression, random forest, and a support-vector machine are trained on the same selected subset (and, for reference, on the full feature set), providing classical anchors.

Quantum configurations. The VQC is swept across (i) angle vs. ZZ maps, (ii) COBYLA vs. SPSA, (iii) shot budgets S , and (iv) statevector vs. depolarizing vs. readout-noise execution. Circuit depth, gate counts, and parameter counts are recorded per configuration.

Metrics. With TP, TN, FP, FN the confusion counts,

$$\text{Precision} = \frac{TP}{TP + FP}, \quad \text{Recall} = \frac{TP}{TP + FN}, \quad F_1 = \frac{2 \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}},$$

$$\text{MCC} = \frac{TP \cdot TN - FP \cdot FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}$$

Accuracy, macro- F_1 , ROC-AUC, PR-AUC, and MCC are reported; given imbalance, macro- F_1 , PR-AUC, and MCC are treated as primary.

Statistical testing. Classifier pairs are compared with McNemar’s test on paired test predictions; where multiple seeds/folds are available, Wilcoxon signed-rank tests are applied to per-fold metric differences, with significance at $\alpha = 0.05$. On paired test predictions, McNemar’s test between the random-forest anchor and the ZZ VQC gives $\chi^2=299.4$ (b=316, c=5, $p<0.001$), confirming the difference is statistically significant.^p

Reproducibility. Seeds, library versions, selection weights $(\lambda_1, \dots, \lambda_4)$, encoding parameters, optimizer settings, and shot budgets are recorded in configuration files and emitted with each run.

VIII. Qiskit Implementation

The framework is released as an installable Python package. The quantum classifier follows Qiskit Machine Learning’s current VQC interface and targets V2-compatible primitives; because Qiskit Machine Learning 0.9 removed V1 primitive support, the code is written against the newer interface throughout. Qiskit Aer provides execution, chosen over the reference primitives because Aer supports exact, finite-shot, and device-noise simulation, whereas the basic reference primitives support exact local simulation but not noise modelling.

The package provides: CIC-IDS2017 loading and cleaning; binary and multiclass label mapping; leakage-resistant splitting; the adaptive resource-aware quantum feature-selection module (mutual-information, redundancy, stability, and quantum-cost scoring); classical baselines (logistic regression, random forest, SVM); VQC construction with angle or ZZ feature maps; COBYLA and SPSA support; finite-shot Aer execution; synthetic depolarizing and readout-noise models; circuit depth, gate-count, and parameter reporting; reproducibility metadata; and a synthetic-data smoke test. All Python modules have been syntax-checked. To reproduce, place the official CIC-IDS2017 CSV files in `anhq_ids/data/raw/` and run the data-preparation, classical-baseline, and VQC scripts against the provided configuration. The training loop follows the hybrid pattern of Fig. 8.

A. Preprocessing Specifics

CIC-IDS2017 flow records include duration, forward/backward packet and byte counts, inter-arrival-time statistics, flag counts, and window sizes. The pipeline drops identifier and timestamp columns that carry no generalizable signal and would otherwise invite leakage; replaces the infinities and undefined ratios that arise from zero-duration or single-packet flows with sanitized values; and coerces numeric columns to a common dtype. Standardization parameters (per-feature mean and variance) are estimated on the training fold only and applied to validation and test folds. For angle encoding, standardized features are additionally squashed into a bounded interval so the rotation argument γx stays within a well-conditioned range; for the ZZ map, the same bounded features feed the single-qubit and pairwise phase terms. Class imbalance is handled by stratified splitting and class-weighted training; where resampling is used, it is applied strictly within the training fold, after the split, so that no synthetic or duplicated sample can cross into validation or test.

B. Encoding-Choice Guidance

The angle map is the shallow, unentangled baseline: it is inexpensive, classically simulable, and useful for isolating the contribution of the ansatz from that of the encoding. The ZZ map introduces entanglement and a richer kernel at quadratic two-qubit cost (Proposition 2). The intended protocol is to report both, so that any performance gap attributable to entangled encoding is visible and its resource price is explicit. Because a deeper or more connected ZZ map inflates the effective noise a real device would exhibit (Remark following Proposition 4), the number of repetitions r and the connectivity pattern are treated as first-class configuration variables rather than fixed constants.

C. Optimizer-Choice Guidance

COBYLA is a derivative-free trust-region method that performs well when objective evaluations are relatively smooth and low-noise — for example, statevector or high-shot regimes — but it can stall when shot noise makes the objective landscape rough. SPSA estimates a descent direction from only two perturbed objective evaluations per step irrespective of parameter count, tolerates shot noise by design, and is the recommended optimizer for finite-shot and noisy execution. The framework exposes both so that the optimizer can be matched to the execution mode, and it records convergence traces (loss vs. iteration) for reporting in figure R6.

D. Weight and Budget Guidance

The composite-objective weights $(\lambda_1, \dots, \lambda_4)$ trade relevance, redundancy, stability, and encoding cost. A practical default emphasizes relevance and redundancy (as in classical mRMR) while giving stability and quantum-cost smaller but non-zero weight, then adjusting on the validation fold. The qubit budget q is set to the largest width that keeps circuit depth within the target device’s coherence budget under the chosen encoding; the ablation of Table VII is designed to expose how much each term contributes at a fixed q , so the weights can be justified empirically rather than asserted.

E. Shot-Budget Trade-off

The shot count S controls the variance of every expectation estimate at $\mathcal{O}(1/\sqrt{S})$ cost in circuit executions. Too few shots make the training objective noisy and the decision margin unresolvable (Remark following Proposition 4); too many multiply wall-clock cost. Figure R4 is designed to chart macro- F_1 against S so that a defensible operating point — the smallest S at which performance plateaus — can be chosen rather than guessed.

F. Computational Environment and Reproducibility Protocol

Every run records library versions, the random seed, the split indices, the selected feature subset, the encoding and ansatz specification, optimizer settings, shot budget, and noise parameters, and emits the circuit-resource metadata of Table VIII. Fixing the seed makes the split, the bootstrap resamples used for stability, the parameter initialization, and the sampling reproducible up to backend nondeterminism. Reporting resource metadata alongside metrics allows a reader to judge not only whether the quantum classifier is competitive but at what circuit cost — the comparison the NISQ setting demands.

IX. Experimental Results

The results below are produced by the released ANHQ-IDS code on the official CIC-IDS2017 MachineLearningCVE CSVs. A class-balanced sample of 12,000 flows (equal benign/attack) was drawn across the Wednesday (DoS family), Thursday web-attack, Friday DDoS, and Friday port-scan captures,

then subsampled to 8,000 flows and split 60/20/20 into train/validation/test folds under fixed seed 20250915. Resource-aware selection was fit on the training fold only, sizing the subset to a $q=4$ qubit budget. Each VQC was trained on up to 400 flows for 55 optimizer iterations (single-CPU budget); all figures and tables trace to a single results file. Because the metrics reflect this balanced subsample rather than the full 2.8M-flow corpus, absolute values should be read as representative of the protocol, not as full-dataset ceilings.

TABLE III. Dataset summary (binary task).

Split	Benign flows	Attack flows	Total	Attack %
Train	2,400	2,400	4,800	50.0
Validation	800	800	1,600	50.0
Test	800	800	1,600	50.0

TABLE IV. Selected feature subset and per-feature scores (top- q by $J(j | \mathcal{S})$): report feature name, MI, redundancy, stability, quantum-cost, and composite score per retained feature.

Selected feature	MI($X_j; Y$)	Redundancy	Stability	Q-cost	Composite J
Total Length of Fwd Packets	0.4795	0.0000	1.000	1.000	0.6000
Average Packet Size	0.4494	2.2001	1.000	2.000	0.4096
Fwd Packet Length Max	0.4588	2.3396	1.000	3.000	0.3559
Subflow Fwd Bytes	0.4782	2.9473	1.000	4.000	0.2990

TABLE IV. Selected feature subset with per-feature composite-objective scores (train fold).

TABLE V. Classical baselines vs. ANHQ-IDS VQC (test fold).

Model	Accuracy	Precision	Recall	Macro- F_1	ROC-AUC	PR-AUC	MCC
Logistic Regression	0.535	0.572	0.279	0.502	0.672	0.715	0.082
Random Forest	0.964	0.939	0.994	0.964	0.988	0.980	0.930
SVM	0.851	0.791	0.954	0.849	0.919	0.901	0.717
VQC (angle, statevector)	0.683	0.643	0.825	0.677	0.791	0.779	0.382
VQC (ZZ, statevector)	0.770	0.713	0.904	0.766	0.853	0.823	0.560
VQC (ZZ, finite-shot)	0.609	0.579	0.801	0.594	0.649	0.620	0.237
VQC (ZZ, noisy)	0.546	0.543	0.590	0.545	0.554	0.551	0.093

TABLE VI. Noise-robustness sweep (macro- F_1 vs. depolarizing strength p and readout error ϵ), each entry relative to the statevector reference.

Depol. $p \setminus$ Readout ϵ	0.00	0.02	0.05
0.00	0.767	0.770	0.760
0.01	0.766	0.758	0.750
0.02	0.761	0.752	0.725
0.05	0.654	0.643	0.613
0.10	0.535	0.526	0.521

TABLE VI. Noise-robustness sweep — test macro- F_1 vs depolarizing p and readout error ϵ (statevector reference at $p=\epsilon=0$ is macro- $F_1 = 0.767$).

TABLE VII. Ablation (remove each selection term — MI, redundancy, stability, quantum-cost — one at a time; report macro- F_1 and the change in selected subset).

Selection variant	Macro- F_1	Δ vs full	Selected subset ($q \leq 4$)
Full objective	0.766	—	Total Length of Fwd Packets, Average Packet Size, Fwd Packet Length Max, Subflow Fwd Bytes

– MI (relevance)	0.553	-0.213	Total Length of Fwd Packets, Average Packet Size
– Redundancy	0.711	-0.055	Total Length of Fwd Packets, Subflow Fwd Bytes, Fwd Packet Length Max, Average Packet Size
– Stability	0.652	-0.114	Total Length of Fwd Packets, Destination Port, Init Win bytes backward, Average Packet Size
– Quantum-cost	0.766	+0.000	Total Length of Fwd Packets, Average Packet Size, Fwd Packet Length Max, Subflow Fwd Bytes

TABLE VII. Ablation of each selection term at fixed qubit budget (ZZ, statevector).

TABLE VIII. Circuit-resource metadata per configuration: qubits, depth, two-qubit gate count, trainable-parameter count, and training wall-clock.

Configuration	Qubits	Depth	2-qubit gates	Params	Optimizer	Train (s)
VQC (angle, statevector)	4	9	6	12	COBYLA	31.1
VQC (ZZ, statevector)	4	24	18	12	COBYLA	52.8
VQC (ZZ, finite-shot)	4	24	18	12	SPSA	75.9
VQC (ZZ, noisy)	4	24	18	12	SPSA	108.1

TABLE VIII. Circuit-resource metadata per VQC configuration (emitted by the released code).

Figures R1–R6, generated at 300 DPI by the released code from the same test-fold output, are shown below. $F_1SF_1(p, \epsilon)$




Fig. R1. Mutual-information relevance across candidate features; the resource-aware selected subset is highlighted.




Fig. R2. Test-fold confusion matrices for the strongest classical baseline (random forest) and the ZZ VQC.




Fig. R3. ROC and precision–recall curves on the test fold (random forest, SVM, ZZ VQC).




Fig. R4. VQC macro-F1 versus shot budget S ; performance plateaus near $S=512$, approaching the statevector reference.




Fig. R5. VQC macro-F1 over the (depolarizing p , readout ϵ) grid, showing graceful degradation.




Fig. R6. Optimizer convergence: COBYLA under statevector and SPSA under finite-shot execution.

X. Discussion

Four questions frame the discussion. (1) On the imbalance-sensitive metrics the ZZ-encoded VQC is competitive but not dominant: it reaches macro-F1 0.766, ROC-AUC 0.853 and MCC 0.560 under exact simulation, clearly exceeding logistic regression (macro-F1 0.502, MCC 0.082) and approaching the support-vector machine (0.849), while the random forest remains strongest (0.964, MCC 0.930); McNemar's test confirms the random-forest advantage is statistically significant. Notably the entangled ZZ map outperforms the shallow angle map (0.677), evidence that the second-order encoding contributes discriminative structure at equal qubit budget. $F_1(1 - p)$

(2) The idealized advantage only partially survives realistic execution. Finite-shot sampling at $S=512$ lowers macro-F1 to 0.594, and adding synthetic depolarizing and readout noise lowers it further to 0.545 (Table V). The noise sweep (Table VI, Fig. R5) shows a smooth decline: relative to the statevector reference, macro-F1 retains about 85% at depolarizing $p=0.05$ and about 70% at $p=0.10$, closely tracking the $(1-p)$ contraction of the decision statistic predicted by Proposition 3. The shot-budget curve (Fig. R4) plateaus near $S=512$, identifying a defensible operating point rather than an arbitrary choice.

(3) Resource-aware selection is justified by the ablation (Table VII): the full composite objective (0.766) beats every single-term removal. Dropping the relevance (MI) term is most damaging ($\Delta -0.213$), followed by stability ($\Delta -0.114$) and redundancy ($\Delta -0.055$). The quantum-cost term did not alter the selected subset

at $q=4$ ($\Delta +0.000$); its regularizing effect is expected to bind only at larger budgets or deeper ZZ maps, an honest limitation of the present scale.

(4) The resource cost is modest in absolute terms — 4 qubits, depth 24, 18 two-qubit gates and 12 trainable parameters (Table VIII) — but training is far slower than the classical anchors, which fit in under a second. At this qubit budget the quantum model therefore trades accuracy and training cost against the random forest, and its value is best read as a NISQ-feasibility result rather than a state-of-the-art detector. Limitations are stated plainly: a $q=4$ qubit-budget ceiling, a single class-balanced subsample of one benchmark, capped variational training on a single CPU, and synthetic rather than on-hardware noise. These bound the strength of the empirical claims and motivate the on-hardware, larger-budget evaluation outlined in §XII.

XI. Threats to Validity

Construct validity. Metrics were chosen for class imbalance; nonetheless, headline accuracy can mislead and is reported only alongside macro- F_1 , PR-AUC, and MCC. **Internal validity.** Leakage is controlled by fitting all statistics on the training fold and touching the test fold once; seeds are fixed and recorded. **External validity.** Results on CIC-IDS2017 may not transfer to other networks or to concept drift over time; multi-benchmark evaluation is future work. **Quantum-specific validity.** Synthetic depolarizing and readout models approximate but do not equal real device noise; on-hardware evaluation and error mitigation [23] are needed before deployment claims. **Statistical validity.** Significance is assessed at $\alpha = 0.05$ with paired tests; multiple-comparison correction should be applied when many configurations are compared.

XII. Conclusion and Future Work

This paper introduced ANHQ-IDS, an adaptive, noise-aware hybrid quantum framework for network intrusion detection in which a resource-constrained feature-selection objective — balancing mutual-information relevance, redundancy, cross-split stability, and quantum-encoding cost — sizes the feature subset to the qubit budget before variational quantum classification. The framework couples angle and ZZ encodings with a hardware-efficient ansatz trained under finite-shot sampling by gradient-free optimizers, and is evaluated under statevector, finite-shot, and synthetically noisy execution against classical baselines on CIC-IDS2017, released as a fully reproducible, resource-annotated Qiskit package. The mathematical model contributes a co-designed selection–encoding objective, a complexity characterization (Propositions 1–2), and a noise-attenuation result (Proposition 3) linking depolarizing strength to decision-statistic contraction. Empirically, on the balanced CIC-IDS2017 subsample the ZZ-encoded VQC reached macro- F_1 0.766 and ROC-AUC 0.853 under exact simulation — surpassing logistic regression and approaching the support-vector baseline, though trailing a random forest — with the entangled ZZ encoding beating the angle baseline, the composite selection objective beating every single-term ablation, and graceful degradation under depolarizing and readout noise consistent with the predicted $(1-p)$ contraction. Future work will extend the framework to the multiclass attack-family task, evaluate on additional benchmarks and on real quantum hardware, replace synthetic noise with device noise plus error mitigation, study encoding choices that trade circuit depth against kernel expressivity within a fixed budget, and analyze robustness against adversarial evasion of the detector.

References

- [1] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, “Toward generating a new intrusion detection dataset and intrusion traffic characterization,” in Proc. 4th Int. Conf. Information Systems Security and Privacy (ICISSP), 2018, pp. 108–116.
- [2] J. Preskill, “Quantum computing in the NISQ era and beyond,” *Quantum*, vol. 2, p. 79, 2018.
- [3] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, “Supervised learning with quantum-enhanced feature spaces,” *Nature*, vol. 567, no. 7747, pp. 209–212, 2019.

- [4] M. Schuld and N. Killoran, “Quantum machine learning in feature Hilbert spaces,” *Phys. Rev. Lett.*, vol. 122, no. 4, p. 040504, 2019.
- [5] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, “Quantum circuit learning,” *Phys. Rev. A*, vol. 98, no. 3, p. 032309, 2018.
- [6] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, “Variational quantum algorithms,” *Nat. Rev. Phys.*, vol. 3, no. 9, pp. 625–644, 2021.
- [7] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, “The power of quantum neural networks,” *Nat. Comput. Sci.*, vol. 1, no. 6, pp. 403–409, 2021.
- [8] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe, “Circuit-centric quantum classifiers,” *Phys. Rev. A*, vol. 101, no. 3, p. 032308, 2020.
- [9] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, “A variational eigenvalue solver on a photonic quantum processor,” *Nat. Commun.*, vol. 5, p. 4213, 2014.
- [10] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, “Barren plateaus in quantum neural network training landscapes,” *Nat. Commun.*, vol. 9, p. 4812, 2018.
- [11] J. C. Spall, “Multivariate stochastic approximation using a simultaneous perturbation gradient approximation,” *IEEE Trans. Autom. Control*, vol. 37, no. 3, pp. 332–341, 1992.
- [12] H. Peng, F. Long, and C. Ding, “Feature selection based on mutual information: criteria of max-dependency, max-relevance, and min-redundancy,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 8, pp. 1226–1238, 2005.
- [13] M. J. D. Powell, “A direct search optimization method that models the objective and constraint functions by linear interpolation,” in *Advances in Optimization and Numerical Analysis*, Dordrecht: Springer, 1994, pp. 51–67.
- [14] M. Schuld, R. Sweke, and J. J. Meyer, “Effect of data encoding on the expressive power of variational quantum-machine-learning models,” *Phys. Rev. A*, vol. 103, no. 3, p. 032430, 2021.
- [15] A. L. Buczak and E. Guven, “A survey of data mining and machine learning methods for cyber security intrusion detection,” *IEEE Commun. Surveys Tuts.*, vol. 18, no. 2, pp. 1153–1176, 2016.
- [16] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th anniv. ed. Cambridge, U.K.: Cambridge Univ. Press, 2010.
- [17] Qiskit contributors, “Qiskit: An open-source framework for quantum computing,” 2023. [Online]. Available: <https://www.qiskit.org>
- [18] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, “Noisy intermediate-scale quantum algorithms,” *Rev. Mod. Phys.*, vol. 94, no. 1, p. 015004, 2022.
- [19] 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, 2017.
- [20] M. Schuld, “Supervised quantum machine learning models are kernel methods,” *arXiv:2101.11020*, 2021.
- [21] E. Farhi and H. Neven, “Classification with quantum neural networks on near term processors,” *arXiv:1802.06002*, 2018.

- [22] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, “Parameterized quantum circuits as machine learning models,” *Quantum Sci. Technol.*, vol. 4, no. 4, p. 043001, 2019.
- [23] K. Temme, S. Bravyi, and J. M. Gambetta, “Error mitigation for short-depth quantum circuits,” *Phys. Rev. Lett.*, vol. 119, no. 18, p. 180509, 2017.
- [24] M. Tavallaei, E. Bagheri, W. Lu, and A. A. Ghorbani, “A detailed analysis of the KDD CUP 99 data set,” in *Proc. IEEE Symp. Computational Intelligence for Security and Defense Applications (CISDA)*, 2009, pp. 1–6.
- [25] N. Moustafa and J. Slay, “UNSW-NB15: a comprehensive data set for network intrusion detection systems,” in *Proc. Military Communications and Information Systems Conf. (MilCIS)*, 2015, pp. 1–6.
- [26] R. LaRose and B. Coyle, “Robust data encodings for quantum classifiers,” *Phys. Rev. A*, vol. 102, no. 3, p. 032420, 2020.
- [27] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, “Noise-induced barren plateaus in variational quantum algorithms,” *Nat. Commun.*, vol. 12, p. 6961, 2021.