International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(22s) 2026 pp. 179-190
https://doi.org/10.70917/ijcisim-2026-5424
Article

A Resource-Aware Hybrid Quantum-Classical Framework for Network Intrusion Detection with Selective Quantum Processing

Rahul R. Bhoge[1], Ranjit R. Keole[2], Pravin P. Karde[3]

[1]Department of Computer Science and Engineering, HVPM's College of Engineering and Technology, Amravati 444605, India
[2]Department of Computer Science and Engineering, HVPM's College of Engineering and Technology, Amravati 444605, India
[3]Department of Information Technology, Government Polytechnic Amravati 444603, India
[1]bhoge.rahul@gmail.com, 2ranjitkeole@gmail.com, 3ppkarde@gmail.com

Abstract: Quantum machine learning (QML) is increasingly investigated for network intrusion detection, but the practical value of current quantum models remains uncertain because many studies rely on small simulated data sets, idealized execution, weak classical baselines, or isolated accuracy comparisons. This study develops and evaluates a resource-aware hybrid quantum-classical intrusion detection framework that combines fixed fidelity quantum kernels, repaired Quantum Kernel Alignment (QKA), a Variational Quantum Classifier (VQC), LightGBM, prototype-based kernel reduction, and probabilistic stacking. Experiments use the UNSW-NB15 benchmark under a leakage-aware preprocessing protocol and repeated evaluation across five random seeds. The experimental program covers 4- and 6-qubit representations, multiple training sizes, prototype budgets, and finite-shot configurations, yielding 260 experimental configurations/runs. The strongest quantum-assisted model, a 6-qubit QKA-VQC-LightGBM stack at 100 samples per class, obtains mean F1=0.8480, MCC=0.7092, ROC-AUC=0.9398, and PR-AUC=0.9302. Repaired QKA improves F1 from 0.8043 for the corresponding fixed kernel to 0.8089, while 50-prototype QKA attains F1=0.8257 with a more favorable kernel-workload trade-off than larger prototype budgets. Finite-shot experiments between 1,024 and 4,096 shots remain stable in simulation. However, full-data LightGBM remains stronger overall (F1=0.8957, MCC=0.8068, ROC-AUC=0.9755, PR-AUC=0.9704), and corrected statistical tests do not establish general quantum superiority. These findings motivate a Selective Quantum Intrusion Detection System (S-QIDS), in which classical inference handles high-confidence traffic and quantum processing is reserved for uncertain, novel, or high-risk observations. The study therefore reframes near-term quantum intrusion detection from unconditional quantum advantage toward measurable quantum utility, quantum-classical complementarity, and resource-aware quantum intervention.

Keywords: quantum machine learning, intrusion detection, quantum kernel alignment, variational quantum classifier, LightGBM, hybrid learning, selective quantum processing, S-QIDS, UNSW-NB15, NISQ, quantum utility

1. Introduction

Network intrusion detection has moved from static signature matching toward data-driven anomaly and classification methods because contemporary networks combine cloud services, Internet of Things devices, edge systems, encrypted communication, and continuously changing attack behavior.

Benchmark data sets such as UNSW-NB15 were created to provide more realistic normal and malicious traffic than legacy KDD-derived corpora and remain widely used for reproducible IDS evaluation [1,2]. Classical machine-learning methods, including support vector machines, neural networks, gradient boosting, and ensembles, can provide strong predictive performance on such tabular security data [3-5]. Consequently, a useful quantum IDS must be evaluated against competitive classical baselines rather than against deliberately weak comparators.

Quantum machine learning offers a different mechanism for constructing representations and decision functions. Quantum feature maps embed classical observations into quantum state spaces; fidelity kernels use overlaps between encoded states as similarity measures; and parameterized quantum circuits can be trained as variational classifiers [6-12]. These approaches are attractive for intrusion detection because network traffic often contains nonlinear interactions among a large number of correlated attributes. Nevertheless, a high-dimensional Hilbert space does not by itself imply a predictive or computational advantage. Huang et al. showed that access to training data can enable classical models to match quantum models even when the underlying quantum feature map is difficult to simulate directly, making fair classical comparison essential [13].

Near-term QML also faces practical limitations. Variational circuits may exhibit optimizer sensitivity and barren plateaus [14-16], while kernel methods require a number of pairwise evaluations that grows approximately quadratically with the training-set size. Current NISQ devices add further constraints through noise, finite connectivity, circuit-depth limitations, readout errors, and limited qubit availability. Recent QML-IDS studies therefore increasingly evaluate accuracy-efficiency trade-offs, circuit complexity, and noisy execution rather than reporting only classification accuracy [17-21,31,33-35].

Our earlier Springer review synthesized quantum supervised learning approaches for intrusion detection and identified hybrid quantum-classical learning, scalability, robustness, and hardware feasibility as important research directions [22]. Subsequent practical work compared VQC and QKA in a unified Qiskit 2.5 pipeline using the UNSW-NB15 data set and established a reproducible basis for comparing variational and trainable-kernel approaches [23]. The present study advances that work in three ways. First, it evaluates fixed kernels, repaired QKA, VQC, LightGBM, prototype-QKA, finite-shot behavior, and probabilistic stacking under repeated experimental conditions [37]. Second, it explicitly evaluates quantum performance relative to a strong full-data classical baseline and uses statistical testing to avoid interpreting numerical improvements as proof of quantum advantage. Third, it derives a deployment-oriented Selective Quantum Intrusion Detection System (S-QIDS), in which quantum computation is used only for observations where its additional representation capacity may justify its resource cost.

The central hypothesis is therefore not that quantum models universally replace classical IDS. Instead, the study asks whether quantum components can provide measurable utility within a hybrid security pipeline, particularly in low-data or ambiguous regions of the decision space. This distinction between quantum advantage, quantum utility, and quantum-classical complementarity forms the conceptual basis of the paper.

1.1 Contributions

A unified hybrid IDS pipeline integrating fixed fidelity quantum kernels, repaired QKA, VQC, LightGBM, and probabilistic stacking.

A repeated experimental protocol spanning five random seeds, 4- and 6-qubit configurations, multiple data budgets, prototype budgets, and finite-shot conditions, resulting in 260 experimental configurations/runs.

A repaired QKA study that compares task-adapted kernels directly with the corresponding fixed quantum kernel.

A prototype-QKA analysis that quantifies predictive performance against approximate kernel workload and runtime.

Finite-shot evaluation at 1,024, 2,048, and 4,096 shots to distinguish sampling stability from true hardware-noise robustness.

Comparison against matched and full-data LightGBM baselines with statistical testing, explicitly avoiding unsupported claims of universal quantum advantage.

The S-QIDS architecture, which treats quantum processing as a selective escalation resource rather than an always-on replacement for mature classical IDS models.

2. Related Work and Research Gap

2.1 Classical IDS and the need for strong baselines

UNSW-NB15 contains contemporary attack families and was designed to address several weaknesses of KDD99-era data sets [1,2]. Classical IDS research on UNSW-NB15 has demonstrated that nonlinear SVMs, tree ensembles, and gradient-boosting methods can achieve strong detection performance, although false alarms and class imbalance remain important concerns [3-5]. For quantum IDS research, this matters because conclusions about quantum benefit are meaningful only when classical baselines are appropriately optimized for the same task and information budget.

2.2 Quantum feature spaces, kernels, and trainable embeddings

Quantum feature-space learning can be formulated by encoding a classical input x through a unitary $U\varphi(x)$ and defining a similarity measure from state overlaps [6-8]. Havlíček et al. introduced quantum-enhanced feature spaces for supervised learning [6], while Schuld and Killoran formalized QML in feature Hilbert spaces [7]. Quantum kernel methods retain classical convex optimization after quantum similarity estimation, making them attractive for NISQ settings. Trainable quantum embedding kernels extend fixed feature maps by optimizing circuit parameters for the target task [11,12]. Such adaptation is closely related to kernel-target alignment and forms the basis of the repaired QKA implementation studied here.

2.3 Variational quantum classifiers and trainability

VQCs learn a parameterized decision function directly from measurements of a trainable circuit [9,10]. Their flexibility comes with a nonconvex optimization problem that can be sensitive to initialization, circuit depth, noise, and parameterization. Barren plateaus can cause gradients to vanish exponentially in unfavorable settings [14-16,34,35]. Accordingly, deeper or more expressive circuits cannot be assumed to improve intrusion detection; trainability and resource cost must be measured empirically.

2.4 QML for intrusion detection

Security intrusion detection using QML has progressed from proof-of-concept experiments toward comparative and noise-aware studies. Kalinin and Krundyshev evaluated QML techniques for security intrusion detection [17]. Abreu et al. proposed QML-IDS as a hybrid quantum-machine-learning intrusion detection framework [18]. More recent studies have evaluated variational circuits, hybrid QNNs, and quantum kernels on modern IDS or IoT data sets [19-21]. Cirillo et al. benchmarked

Pegasos QSVC, VQC, and a hybrid quantum-classical neural network under ideal and noisy conditions and explicitly analyzed execution-time and circuit-complexity trade-offs [19]. Bharathi et al. additionally investigated an IoT intrusion-detection framework with quantum-driven learning and hardware-oriented experiments [20]. Recent work has also reported practical quantum-kernel NIDS experiments on hardware and outlier-oriented quantum intrusion detection, broadening evidence beyond idealized simulator-only studies [31,32].

A 2026 survey of QML-IDS literature reports that hybrid quantum-classical approaches now dominate the field because they combine quantum representation mechanisms with the stability and scalability of classical learning. The same survey identifies simulator dependence, limited real-device validation, low reporting frequency for computational metrics, and uncertain generalization as recurring limitations [21,31,32,36]. These findings motivate resource-aware benchmarking rather than accuracy-only comparison.

2.5 Research gap

Three gaps remain particularly relevant. First, many QML-IDS studies evaluate a single quantum model or compare studies that use different preprocessing pipelines, data sizes, or metrics. Second, quantum-kernel scalability and variational-trainability problems are often discussed theoretically but not measured alongside intrusion-detection

performance. Third, most architectures assume that every observation should be processed by the quantum model, although current quantum resources are scarce and classical models can already classify many observations confidently. The present study addresses the first two gaps experimentally and proposes S-QIDS as a concrete design response to the third. Recent kernel studies further show that practical deployment must consider both hardware behavior and the possibility of kernel-value concentration as the quantum representation scales [31,33].

3. Materials and Methods

3.1 Data set and task definition

The study uses the publicly available UNSW-NB15 benchmark for binary classification of normal versus malicious network traffic [1,2]. The data contain numerical and categorical flow attributes together with attack labels. Binary classification is selected to enable controlled comparison of quantum and classical learners while preserving the operational IDS question of whether a flow should be escalated as malicious.

3.2 Leakage-aware preprocessing

Preprocessing follows a training-first protocol. Irrelevant identifiers are removed; categorical variables are encoded; numerical values are normalized; and feature reduction is learned using only training data before transformation of held-out observations. Low-dimensional representations are then rescaled to ranges compatible with quantum rotation gates. This design prevents test-set information from influencing feature engineering and makes classical and quantum branches comparable.

Table 1: Summarizes the common experimental pipeline.

Stage	Operation	Purpose
1	Load UNSW-NB15 training/testing partitions	Common benchmark source
2	Remove identifiers and encode categorical features	Prevent non-informative leakage; create numerical input
3	Normalize and reduce dimensionality	Fit low-qubit quantum representation
4	Scale latent features to rotation ranges	Quantum feature-map compatibility
5	Train fixed QK, repaired QKA, VQC, and LightGBM	Comparable learning branches
6	Construct probabilistic hybrid stack	Test quantum-classical complementarity
7	Repeat across seeds/resources/shots	Assess stability and resource trade-offs
8	Evaluate predictive and statistical metrics	Avoid single-metric conclusions

3.3 Quantum feature map and fidelity kernel

For an n-dimensional transformed input x, the feature-map circuit prepares a state

$$|\varphi(x)\rangle = U\varphi(x)|0\rangle\otimes n \quad (1)$$

A fidelity-based quantum kernel is then defined as

$$K(x_i,x_j) = |\langle\varphi(x_i)|\varphi(x_j)\rangle|^2 \quad (2)$$

The kernel matrix is passed to a classical kernel classifier. This formulation is hybrid by construction: the quantum component generates similarity information, while the final decision boundary is optimized classically [6,7,11]. The fixed kernel is retained as an essential control because QKA should be judged by improvement relative to the same untrained quantum representation.

3.4 Repaired Quantum Kernel Alignment

QKA introduces trainable circuit parameters θ into the feature map:

$K\theta(x_i,x_j) = |\langle\varphi\theta(x_i)|\varphi\theta(x_j)\rangle|^2$ (3)

The kernel parameters are optimized to improve agreement between the learned kernel matrix and the target-label structure. A normalized target-alignment objective can be written as

$A(K,Y) = \langle K,Y\rangle F / (||K||F\ ||Y||F)$ (4)

where $\langle\cdot,\cdot\rangle$ F denotes the Frobenius inner product and Y is an ideal label kernel. An increase in alignment and downstream F1 provides evidence that task-specific kernel adaptation functions as intended; it does not by itself prove a computational quantum advantage [11-13].

3.5 Variational Quantum Classifier

The VQC branch encodes the transformed input and applies a trainable variational ansatz:

$|\psi(x,\theta)\rangle = U\theta\ U\varphi(x)|0\rangle\otimes n$ (5)

Measurements are converted to class probabilities or decision values, and a classical optimizer updates θ. The VQC is evaluated as a compact nonlinear quantum classifier rather than being assumed to provide faster or more accurate inference than classical alternatives. This interpretation is consistent with known trainability limitations of variational circuits [9,14-16].

3.6 Classical LightGBM baseline

LightGBM is used because gradient boosting is a strong baseline for structured tabular security data. Two roles are distinguished. Matched LightGBM is trained using the same restricted information budget as the quantum experiment, enabling a more direct representation comparison. Full-data

LightGBM uses the complete classical training information and therefore represents the more realistic operational baseline that a deployable quantum-assisted IDS must justify exceeding or complementing.

3.7 Probabilistic hybrid stacking

Let pQ(x), pV(x), and pC(x) denote probabilities from QKA, VQC, and LightGBM. A classical meta-learner g combines these outputs:

$pH(x) = g(pQ(x), pV(x), pC(x))$ (6)

The stack tests whether quantum and classical learners make partially different errors. Improvement over matched LightGBM is interpreted as complementarity only when supported by repeated trials and statistical analysis.

3.8 Prototype-QKA

Exact quantum-kernel construction requires approximately $O(N^2)$ pairwise similarities for N training points. To reduce the workload, prototype-QKA constructs the quantum training representation from a smaller set of representative observations. Prototype budgets of 20, 50, and 100 per class are evaluated. The objective is not merely to minimize runtime but to identify the point at which additional quantum evaluations cease to provide useful predictive gain.

3.9 Finite-shot evaluation

The repaired 6-qubit QKA is evaluated with 1,024, 2,048, and 4,096 shots. Finite-shot experiments quantify sampling variability but do not reproduce all hardware effects. In particular, shot-based simulation does not fully model decoherence, gate infidelity, readout errors, connectivity restrictions, or device queue latency. Therefore, finite-shot stability is reported separately from hardware robustness.

3.10 Experimental protocol and metrics

Experiments are repeated using random seeds 11, 22, 33, 44, and 55. The study includes 4- and 6-qubit representations, multiple samples per class, fixed and trainable kernels, VQC, matched/full LightGBM, hybrid stacking, prototype-QKA, and finite-shot configurations. Across the experimental program, 260 configurations/runs are analyzed. Primary metrics are F1-score, Matthews Correlation Coefficient (MCC), ROC-AUC, and PR-AUC; accuracy, precision, recall, confusion matrices, runtime, and approximate kernel-pair counts are considered where available. For imbalanced IDS data, F1, MCC, and PR-AUC are emphasized because accuracy alone can conceal poor minority-class detection [21].

4. Proposed Selective Quantum Intrusion Detection System (S-QIDS)

4.1 Motivation

The experimental results reported in Sect. 5 show that full-data LightGBM is stronger overall, while quantum-assisted stacking can improve over matched classical learning in selected low-data conditions. This pattern suggests that an always-on quantum classifier is unlikely to be the most efficient near-term deployment strategy. S-QIDS therefore treats quantum computation as a specialized escalation resource.

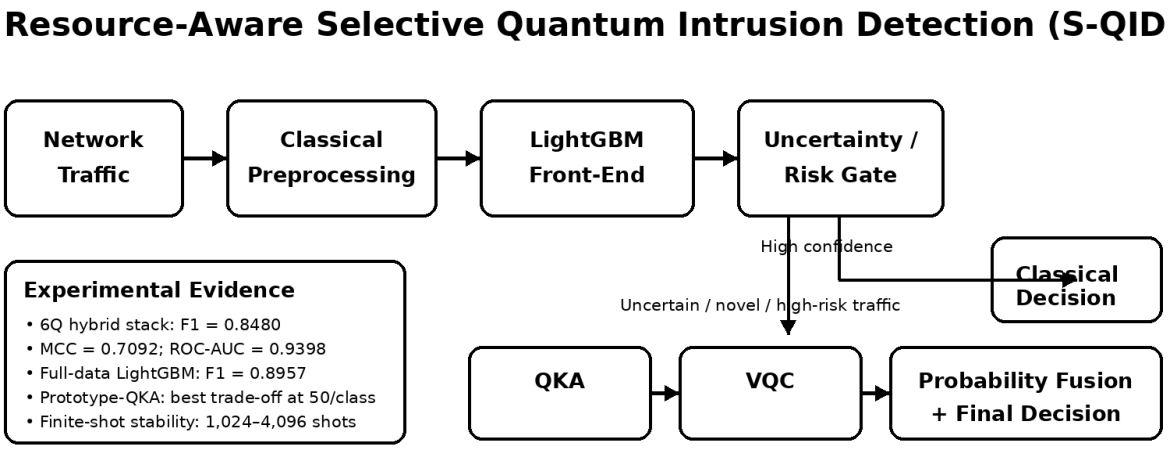

Figure 1: Proposed Resource-Aware Selective Quantum Intrusion Detection (S-QIDS)

4.2 Architecture

Table 2: Proposed S-QIDS selective quantum-processing architecture.

Layer	Operation	Decision
1. Traffic ingestion	Extract and normalize flow features	Create classical input
2. Classical IDS	LightGBM produces probability and uncertainty	Accept high-confidence decisions
3. Risk/novelty gate	Evaluate uncertainty, novelty, and security-risk score	Route only selected flows
4. Quantum branch	QKA and/or VQC on compressed representation	Generate complementary quantum evidence
5. Fusion	Stack classical and quantum probabilities	Final IDS decision

6. Feedback	Log escalated flows and outcomes	Update thresholds and models

Define a binary routing variable R(x) for observation x:

$$R(x) = 1\{ U(x)>\tau U \text{ OR } N(x)>\tau N \text{ OR } S(x)>\tau S \}, \quad (7)$$

where U(x) is classical predictive uncertainty, N(x) is a novelty or out-of-distribution score, and S(x) is a security-risk score. R(x)=1 forwards the observation to the quantum branch; otherwise the classical decision is retained. If ρ denotes the fraction of traffic routed to quantum processing, then the selective quantum workload is

$$Q_{selective} = \rho\, Q_{full} \quad (8)$$

This expression is a resource model, not a claimed speedup. The practical research question becomes: what routing fraction ρ yields the best trade-off among detection quality, false alarms, latency, and quantum circuit executions?

4.3 Proposed S-QIDS evaluation

A future implementation should evaluate routing budgets such as $\rho \in \{0.05, 0.10, 0.20, 0.30, 0.50, 1.00\}$. At each budget, F1, MCC, PR-AUC, false-positive rate, attack detection rate, calibration, number of quantum calls, circuit executions, shots, and end-to-end latency should be reported. This converts quantum usage itself into an explicit experimental variable and directly addresses the deployment limitations identified by recent QML-IDS surveys [19-21].

5. Results

5.1 Overall predictive performance

The strongest quantum-assisted configuration is the 6-qubit QKA-VQC-LightGBM stack at 100 samples per class. It achieves mean F1=0.8480, MCC=0.7092, ROC-AUC=0.9398, and PR-AUC=0.9302. Full-data LightGBM remains stronger overall, with F1=0.8957, MCC=0.8068, ROC-AUC=0.9755, and PR-AUC=0.9704. The central result is therefore competitive hybrid quantum utility rather than general quantum superiority.

Table 3: Selected overall performance comparison.

Model / configuration	**F1**	**MCC**	**ROC-AUC**	**PR-AUC**
6Q QKA-VQC-LightGBM stack (100/class)	0.8480	0.7092	0.9398	0.9302
Full-data LightGBM	0.8957	0.8068	0.9755	0.9704

5.2 Repaired QKA and 4Q/6Q ablation

At 6Q/100, repaired QKA improves mean F1 from 0.8043 for the fixed kernel to 0.8089 and MCC from approximately 0.6471 to 0.6538. The improvement is modest and is interpreted as evidence of successful trainable adaptation rather than evidence of quantum advantage. Six-dimensional representations provide small but consistent improvements for the major models at the 100-sample/class setting.

Table 4: Four- versus six-qubit comparison at 100 samples per class.

Model	**4Q F1**	**6Q F1**
Fixed quantum kernel	0.8028	0.8043
Repaired QKA	0.8082	0.8089

Matched LightGBM	0.8238	0.8306
Hybrid stack	0.8416	0.8480

5.3 Hybrid complementarity and statistical interpretation

At 6Q/100, mean F1 rises from 0.8089 for repaired QKA and 0.8306 for matched LightGBM to 0.8480 for the stack. The improvement over matched LightGBM is 0.0174 with paired effect size $dz \approx 0.93$. However, the raw Wilcoxon p-value is 0.0625 and does not remain significant after correction for multiple comparisons. The ensemble therefore demonstrates promising complementary behavior without statistically established superiority.

5.4 Prototype-QKA and resource trade-off

Prototype-QKA reveals a non-monotonic relationship between quantum workload and predictive performance. The 50-prototype configuration achieves the most favorable observed trade-off. Increasing the prototype budget from 50 to 100 per class more than doubles approximate kernel-pair evaluations while reducing F1. This result cautions against equating larger quantum workloads with better models.

Table 5: Six-qubit prototype-QKA at a 250-sample/class source pool.

Prototypes/class	F1	MCC	PR-AUC	Runtime (s)	Approx. kernel pairs
20	0.7018	0.4308	0.7904	5.44	20,820
50	0.8257	0.6919	0.9218	13.65	55,050
100	0.8040	0.6620	0.9209	32.50	120,100

5.5 Finite-shot stability

No systematic degradation is observed between 1,024 and 4,096 shots. The values fluctuate within a narrow range, supporting finite-sampling stability in simulation. This result should not be generalized to physical hardware because it does not include all device-specific noise mechanisms.

Table 6: Finite-shot performance of repaired 6Q QKA.

Shots	F1	MCC	PR-AUC
1,024	0.8112	0.6584	0.9077
2,048	0.8092	0.6544	0.9054
4,096	0.8132	0.6621	0.9056

6. Discussion

6.1 From quantum advantage to quantum utility

The results support a cautious but meaningful interpretation. Repaired QKA improves over the corresponding fixed kernel, showing that task-specific quantum representation learning is functioning. The hybrid stack improves over matched LightGBM in the 6Q/100 setting, indicating that quantum and classical learners can capture partially different decision structure. Yet the strongest full-data LightGBM remains better overall, and multiple-comparison-corrected statistics do not establish quantum superiority. These three facts are mutually consistent and motivate the concept of quantum utility rather than universal quantum advantage.

This interpretation also aligns with the theoretical literature. A quantum embedding may be computationally difficult to reproduce exactly while still providing little useful predictive separation once classical models are trained on data [13]. Likewise, kernel and variational models can suffer from poor inductive bias or trainability even when

their underlying circuits explore very large state spaces [13-16,33-35]. Consequently, the relevant empirical question is whether the quantum component adds useful information under fair resource accounting.

6.2 Why the classical result strengthens the paper

The superiority of full-data LightGBM should not be treated as a failed quantum experiment. It establishes the practical baseline that any near-term quantum-assisted system must justify. The result directly motivates S-QIDS: if the classical model is highly confident for most flows, quantum evaluation should be reserved for ambiguous, novel, or high-risk traffic. Such routing transforms the classical model from a competitor into the first stage of a resource-aware hybrid architecture.

6.3 Scalability implications of prototype-QKA

The prototype results show that quantum-kernel scaling is not only a theoretical concern. Moving from 50 to 100 prototypes per class increases approximate pairwise evaluations from 55,050 to 120,100 and runtime from 13.65 s to 32.50 s, while F1 decreases. The 50-prototype result therefore provides empirical evidence that resource reduction can improve the performance-cost balance. Future kernel research should optimize predictive information per quantum evaluation rather than maximizing sample count indiscriminately. This resource-aware interpretation is also consistent with recent theory showing that quantum-kernel values can become exponentially concentrated, increasing the measurement burden needed to resolve informative similarities [33].

6.4 Finite shots, noise, and hardware realism

Stable finite-shot simulation is encouraging but insufficient to claim hardware robustness. Cirillo et al. demonstrated that intrusion-detection QML performance can vary substantially with entanglement structure, CNOT count, backend noise, and circuit complexity [19,31,32]. Physical-hardware studies therefore need to report device properties, transpiled depth, two-qubit-gate counts, shots, queue-independent execution time, and readout mitigation. These measurements are planned as a hardware-feasibility extension rather than being implied by the current simulation results.

6.5 S-QIDS as a deployment-oriented interpretation

S-QIDS follows directly from the empirical results. Full-data classical learning is presently stronger and cheaper, while quantum models can provide complementary information in constrained regions. A selective routing gate therefore allocates scarce quantum computation where marginal information gain is more likely to be valuable. The concept is compatible with cloud quantum execution because only a fraction of observations require external quantum calls, and it preserves classical fallback behavior when quantum resources are unavailable. Importantly, S-QIDS remains a proposed architecture in the present paper; its routing policy and end-to-end latency must be validated experimentally before deployment claims are made. More broadly, recent cybersecurity-oriented QML studies reinforce the need to treat hardware access, latency, model optimization, and hybrid orchestration as first-class deployment constraints [31,36].

7. Threats to Validity and Limitations

Table 7: Principal threats to validity.

Limitation	Implication
Single primary data set	The comprehensive experimental analysis uses UNSW-NB15. Cross-dataset validation is needed before broad generalization.
Simulation dependence	Finite-shot simulation does not reproduce all physical-device effects, including decoherence, topology, gate errors, readout errors, and queue latency.

Restricted qubit counts	Four- and six-qubit experiments provide controlled NISQ evidence but do not establish large-scale quantum scalability.
Classical preprocessing dependence	Quantum models rely on classical cleaning, encoding, scaling, and dimensionality reduction. Performance therefore cannot be attributed solely to quantum computation.
Kernel complexity	Prototype reduction decreases the workload but does not remove the fundamental pairwise scaling of explicit quantum-kernel estimation.
Statistical power	Five seeds are substantially better than a single run, but larger repeated studies would provide more power for small effect sizes.
Adversarial robustness	Robustness against adversarially manipulated network traffic has not yet been experimentally demonstrated.
S-QIDS status	Selective routing is proposed from the observed evidence but has not yet been executed as a complete online system.

8. Future Work

The immediate next step is to implement S-QIDS using stored or newly generated per-sample probabilities. Classical uncertainty thresholds should route approximately 5%, 10%, 20%, 30%, 50%, and 100% of ambiguous observations to the quantum branch. The resulting curves should report F1, MCC, PR-AUC, false-positive rate, quantum-call count, and latency as functions of the routing budget.

Adversarial robustness is the second priority. Network-valid perturbations should be generated under feature-domain constraints and used to compare LightGBM, QKA, VQC, and S-QIDS. Robustness degradation, attack success rate, and calibration under attack should be reported. Third, the complete methodology should be frozen and externally validated on a modern independent benchmark such as CIC-IDS2017, CSE-CIC-IDS2018, ToN-IoT, or a comparable data set. Fourth, representative final circuits should be evaluated on IBM Quantum hardware or realistic backend-derived noise models, with explicit reporting of transpiled depth, gate counts, shots, and execution time. Barren-plateau mitigation and optimization-aware circuit design should also be considered when extending VQC depth or qubit count [34,35].

These extensions will determine whether the observed quantum-classical complementarity persists under distribution shift, and whether selective quantum processing can offer a practically favorable resource-performance trade-off.

9. Conclusion

This study presents a resource-aware hybrid quantum-classical framework for network intrusion detection that integrates fixed fidelity quantum kernels, repaired QKA, VQC, LightGBM, prototype reduction, finite-shot evaluation, and probabilistic stacking. Across a repeated experimental program of 260 configurations/runs and five random seeds, the strongest quantum-assisted configuration, a 6-qubit QKA-VQC-LightGBM ensemble at 100 observations per class, achieves mean F1=0.8480 and MCC=0.7092. Repaired QKA provides modest improvement over the corresponding fixed kernel, prototype-QKA identifies a favorable 50-prototype performance-cost operating point, and finite-shot simulation remains stable across 1,024-4,096 shots.

However, full-data LightGBM remains superior overall and corrected pairwise tests do not establish universal quantum superiority. The most defensible conclusion is therefore that near-term quantum learning can act as a complementary, selectively applied representation mechanism rather than as a universal replacement for mature classical IDS algorithms. This conclusion motivates S-QIDS, which routes only uncertain, novel, or high-risk observations to quantum processing. The resulting research direction moves quantum cybersecurity from broad advantage claims toward reproducible

benchmarking, explicit resource accounting, selective intervention, and experimentally testable quantum utility.

Declarations

Conflict of interest: The authors declare that they have no conflict of interest.

Data availability: The study uses the publicly available UNSW-NB15 network intrusion detection data set.

https://www.kaggle.com/datasets/mrwellsdavid/unsw-nb15

Code availability:

https://github.com/bhogerahul-commits/Quantum-Supervised-Learning/tree/NA_HQKVE_IDS

REFERENCES

1. Moustafa N, Slay J (2015) UNSW-NB15: a comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set). In: Proceedings of the Military Communications and Information Systems Conference (MilCIS), Canberra, Australia, pp 1–6. https://doi.org/10.1109/MilCIS.2015.7348942
2. Moustafa N, Slay J (2016) The evaluation of network anomaly detection systems: statistical analysis of the UNSW-NB15 data set and the comparison with the KDD99 data set. Inf Secur J Glob Perspect 25(1–3):18–31. https://doi.org/10.1080/19393555.2015.1125974
3. Tama BA, Rhee KH (2019) An in-depth experimental study of anomaly detection using gradient boosted machine. Neural Comput Appl 31:955–965. https://doi.org/10.1007/s00521-017-3128-z
4. Shone N, Ngoc TN, Phai VD, Shi Q (2018) A deep learning approach to network intrusion detection. IEEE Trans Emerg Top Comput Intell 2(1):41–50. https://doi.org/10.1109/TETCI.2017.2772792
5. Tama BA, Comuzzi M, Rhee KH (2019) TSE-IDS: a two-stage classifier ensemble for intelligent anomaly-based intrusion detection system. IEEE Access 7:94497–94507.
6. Havlíček V, Córcoles AD, Temme K, Harrow AW, Kandala A, Chow JM, Gambetta JM (2019) Supervised learning with quantum-enhanced feature spaces. Nature 567:209–212. https://doi.org/10.1038/s41586-019-0980-2
7. Schuld M, Killoran N (2019) Quantum machine learning in feature Hilbert spaces. Phys Rev Lett 122:040504. https://doi.org/10.1103/PhysRevLett.122.040504
8. Schuld M, Petruccione F (2018) Supervised learning with quantum computers. Springer, Cham. https://doi.org/10.1007/978-3-319-96424-9
9. Farhi E, Neven H (2018) Classification with quantum neural networks on near term processors. arXiv:1802.06002
10. Benedetti M, Lloyd E, Sack S, Fiorentini M (2019) Parameterized quantum circuits as machine learning models. Quantum Sci Technol 4:043001. https://doi.org/10.1088/2058-9565/ab4eb5
11. Hubregtsen T, Wierichs D, Gil-Fuster E, Derks PJHS, Faehrmann PK, Meyer JJ (2022) Training quantum embedding kernels on near-term quantum computers. Phys Rev A 106:042431. https://doi.org/10.1103/PhysRevA.106.042431
12. Schuld M (2021) Supervised quantum machine learning models are kernel methods. arXiv:2101.11020
13. Huang HY, Broughton M, Mohseni M, Babbush R, Boixo S, Neven H, McClean JR (2021) Power of data in quantum machine learning. Nat Commun 12:2631. https://doi.org/10.1038/s41467-021-22539-9
14. McClean JR, Boixo S, Smelyanskiy VN, Babbush R, Neven H (2018) Barren plateaus in quantum neural network training landscapes. Nat Commun 9:4812. https://doi.org/10.1038/s41467-018-07090-4
15. Holmes Z, Sharma K, Cerezo M, Coles PJ (2022) Connecting ansatz expressibility to gradient magnitudes and barren plateaus. PRX Quantum 3:010313. https://doi.org/10.1103/PRXQuantum.3.010313
16. Cerezo M, Verdon G, Huang HY, Cincio L, Coles PJ (2022) Challenges and opportunities in quantum machine learning. Nat Comput Sci 2:567–576. https://doi.org/10.1038/s43588-022-00311-3
17. Kalinin M, Krundyshev V (2023) Security intrusion detection using quantum machine learning techniques. J Comput Virol Hack Tech 19:125–136. https://doi.org/10.1007/s11416-022-00435-0
18. Abreu D, Rothenberg CE, Abelem A (2024) QML-IDS: Quantum Machine Learning Intrusion Detection System. arXiv:2410.16308
19. Cirillo F, Esposito C, Seo JT (2026) Benchmarking quantum machine learning methods for intrusion detection on noisy quantum computers. Quantum Mach Intell 8:27. https://doi.org/10.1007/s42484-026-00379-4
20. Bharathi I, Sonai V, Sridevi S (2026) Quantum-driven enhanced machine learning algorithm for intrusion detection in Internet of Things environment. EPJ Quantum Technol 13:20. https://doi.org/10.1140/epjqt/s40507-026-00463-5
21. Kaissar A, Bou Nassif A, Bouridane A (2026) Enhancing network intrusion detection with quantum machine learning: a comprehensive survey of methods, metrics, and applications. Future Internet 18:234. https://doi.org/10.3390/fi18050234
22. Bhoge RR, Keole RR, Karde PP (2026) Integration of quantum supervised learning for intrusion detection system: a review. In: Senjyu T et al (eds) Smart trends in computing and communications. Lecture Notes in Networks and Systems, vol 1459. Springer, Singapore, pp 525–536. https://doi.org/10.1007/978-981-96-7505-0_41
23. Bhoge RR (2026) Comparative analysis of VQC vs QKA on intrusion detection [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.21595715.
24. Abbas A, Sutter D, Zoufal C, Lucchi A, Figalli A, Woerner S (2021) The power of quantum neural networks. Nat Comput Sci 1:403–409. https://doi.org/10.1038/s43588-021-00084-1
25. Pérez-Salinas A, Cervera-Lierta A, Gil-Fuster E, Latorre JI (2020) Data re-uploading for a universal quantum classifier. Quantum 4:226. https://doi.org/10.22331/q-2020-02-06-226
26. Liu Y, Arunachalam S, Temme K (2021) A rigorous and robust quantum speed-up in supervised machine learning. Nat Phys 17:1013–1017. https://doi.org/10.1038/s41567-021-01287-z
27. Cong I, Choi S, Lukin MD (2019) Quantum convolutional neural networks. Nat Phys 15:1273–1278. https://doi.org/10.1038/s41567-019-0648-8

28. Glick JR, Gujarati TP, Córcoles AD, Kim Y, Kandala A, Gambetta JM, Temme K (2024) Covariant quantum kernels for data with group structure. Nat Phys 20:479–483. https://doi.org/10.1038/s41567-023-02340-9
29. Nielsen MA, Chuang IL (2010) Quantum computation and quantum information, 10th anniversary edn. Cambridge University Press, Cambridge
30. Rebentrost P, Mohseni M, Lloyd S (2014) Quantum support vector machine for big data classification. Phys Rev Lett 113:130503. https://doi.org/10.1103/PhysRevLett.113.130503
31. Cotrupi ML, Callahan BR (2025) Towards practical quantum kernels for network intrusion detection. Proc AAAI Symp Ser 7(1):339–342. https://doi.org/10.1609/aaaiss.v7i1.36903
32. Kim TH, Madhavi S (2024) Quantum intrusion detection system using outlier analysis. Sci Rep 14:27114. https://doi.org/10.1038/s41598-024-78389-0
33. Thanasilp S, Wang S, Cerezo M, Holmes Z (2024) Exponential concentration in quantum kernel methods. Nat Commun 15:5200. https://doi.org/10.1038/s41467-024-49287-w
34. Cunningham J, Zhuang J (2025) Investigating and mitigating barren plateaus in variational quantum circuits: a survey. Quantum Inf Process 24:48. https://doi.org/10.1007/s11128-025-04665-1
35. Fontana E, Herman D, Chakrabarti S, Kumar N, Yalovetzky R, Heredge J, Sureshbabu SH, Pistoia M et al (2024) Characterizing barren plateaus in quantum ansätze with the adjoint representation. Nat Commun 15:7171. https://doi.org/10.1038/s41467-024-49910-w
36. Rosa-Remedios C, Caballero-Gil P (2025) Optimizing quantum machine learning for proactive cybersecurity. Optim Eng 26:2321–2353. https://doi.org/10.1007/s11081-024-09934-z
37. Bhoge RR (2026) Quantum-Supervised-Learning-NA_HQKVE_IDS [Computer software]. Zenodo https://doi.org/doi 10.5281/zenodo.22048701.