VOL. 21, NO. 7, APRIL 2026 ISSN 1819-6608
ARPN Journal of Engineering and Applied Sciences

www.arpnjournals.com

QUANTUM-INSPIRED CYBER THREAT DETECTION VIA ENTANGLEMENT-BASED FEATURE FUSION

Sareddy Hemalatha and K. Chandrasekhar
Department of CSE, JNTUA College of Engineering, Pulivendula, Andhra Pradesh, India
E-Mail: hemareddysareddy@gmail.com

ABSTRACT

The increasing sophistication of cyberattacks requires intrusion detection systems that can model nonlinear and complex network-traffic patterns beyond the capabilities of conventional machine-learning designs. This paper proposes a Hybrid Quantum-Inspired Cyber Threat Detection (Hybrid QCTD) framework that combines amplitude-based quantum feature encoding, a variational quantum circuit (VQC), and a lightweight classical classifier. Evaluated on the CICIDS2017 dataset, the proposed model achieved an accuracy of 0.9550, an F1-score of 0.8779, a precision of 0.8108, a recall of 0.9571, and an ROC-AUC of 0.9931. Random Forest and XGBoost achieved higher accuracy and F1-scores, whereas the Hybrid QCTD model maintained strong detection performance and stable training convergence. A prediction-level example further illustrates that the hybrid model may produce different decisions from the classical baselines for complex network flows; however, such examples should be interpreted together with the aggregate test metrics. These findings demonstrate the feasibility of entanglement-based feature fusion for intrusion detection and motivate further evaluation on noisy quantum hardware and resource-constrained, real-time cybersecurity systems.

Keywords: quantum-inspired machine learning, intrusion detection system (IDS), variational quantum circuit (VQC), cybersecurity, CICIDS2017, hybrid quantum-classical model, entanglement-based feature fusion, network traffic analysis, cyber threat detection.

Manuscript Received 9 March 2026; Revised 2 April 2026; Published 10 June 2026

1. INTRODUCTION

The rapid expansion of digital infrastructure has been accompanied by a substantial increase in cyber threats affecting critical sectors such as finance, healthcare, defense, and smart cities. Modern cyberattacks are often multi-vector, dynamic, and stealthy, allowing them to exploit limitations in conventional security mechanisms [1], [2]. Traditional intrusion detection systems (IDSs) based on statistical heuristics or signature-matching techniques may be ineffective against previously unseen or polymorphic attacks because of their limited adaptive learning capability [3]. Machine learning (ML) and artificial intelligence (AI) have therefore been integrated into cybersecurity systems to learn behavioral patterns, adapt to changing traffic, and support near-real-time threat detection [4]. Deep-learning models, including convolutional neural networks (CNNs) and recurrent neural networks (RNNs), have demonstrated strong intrusion-detection performance, but they often require substantial computational resources, large labeled datasets, and long training times [5], [6]. In addition, purely classical architectures may struggle to represent complex high-dimensional correlations in network-traffic data, particularly for rare or previously unseen attacks [7]. These limitations have motivated research into quantum machine learning (QML), which combines quantum-mechanical principles with classical data-driven optimization [8].

Quantum computing uses superposition, entanglement, and interference to process information in high-dimensional state spaces, creating new possibilities for pattern-recognition tasks [9], [17]. Variational quantum circuits (VQCs) and quantum-kernel methods can construct expressive nonlinear feature spaces and decision boundaries [10]-[13]. These properties are relevant to cybersecurity applications in which subtle relationships among network-flow features must be identified while controlling model complexity and overfitting [12], [16].

Building on these developments, this paper proposes a Quantum-Inspired Cyber Threat Detection (QCTD) model based on entanglement-driven feature fusion and hybrid quantum-classical optimization. The framework applies amplitude-based feature encoding, transforms the encoded features through a VQC, and performs final classification using a lightweight classical head. The model is evaluated on the CICIDS2017 dataset and compared with Random Forest, XGBoost, and SVM-RBF baselines using the same data split and evaluation metrics.

2. LITERATURE REVIEW

A. Classical ML Intrusion Detection

Early machine-learning-based network intrusion detection systems (NIDSs) commonly used shallow classifiers such as support vector machines (SVMs), Random Forest, and k-nearest neighbors (KNN), trained on benchmark datasets including KDD99, UNSW-NB15, and CICIDS2017. These methods generally provide strong and interpretable baselines, but their performance can deteriorate when network traffic is high-dimensional, nonstationary, or affected by multi-vector attacks that change feature distributions over time. UNSW-NB15 was

introduced to address several limitations of KDD99 by providing more contemporary attack categories and network-flow attributes [4]. CICIDS2017 contains labeled benign and attack flows across scenarios such as DoS/DDoS, PortScan, Botnet, and Web attacks [5]. Although classical models perform well on these datasets, generalization to rare, evolving, or cross-domain attacks remains challenging.

B. Deep Learning Advances

Deep-learning architectures such as CNNs, RNNs, LSTMs, and GRUs can learn nonlinear and spatio-temporal relationships directly from network flows or payload-derived features [2], [3], [6]. Their improved representational capacity can increase detection performance, but these models are often parameter-intensive, data-hungry, and computationally expensive. Such requirements may limit their suitability for latency-sensitive or edge-based deployments. These trade-offs motivate compact feature-learning methods that can capture higher-order relationships without relying on very large classical networks.

C. Variational and Kernel-Based Quantum Machine Learning (QML)

QML offers alternative representations through quantum feature maps, quantum kernels, and variational quantum circuits operating in high-dimensional Hilbert spaces [7]-[15]. Havlíček et al. demonstrated supervised learning with quantum-enhanced feature spaces using a quantum variational classifier and a quantum-kernel estimator [7]. Variational quantum algorithms are particularly relevant to the noisy intermediate-scale quantum (NISQ) era because they combine parameterized quantum circuits with classical optimization [8], [13], [14]. However, practical issues such as barren plateaus, limited trainability, noise, circuit depth, and ansatz selection remain important. Problem-inspired circuit structures, careful initialization, and layer-wise training are commonly investigated to mitigate these limitations.

D. QML Opportunities and Gaps in Cybersecurity

End-to-end quantum intrusion detection remains an emerging research area, although early studies indicate that quantum kernels and VQCs may provide expressive nonlinear feature transformations with compact parameterizations [16]. Important open issues include fair comparison with tuned classical baselines, consistent preprocessing and class-imbalance handling, and transparent reporting of qubit count, circuit depth, gate count, shot budget, runtime, and noise effects. The proposed QCTD architecture addresses part of this gap by: (i) encoding network-flow features into quantum states, (ii) applying entangling VQC layers for feature fusion, (iii) using a lightweight classical classification head, and (iv) benchmarking the hybrid model against classical baselines under a common evaluation protocol.

Table-1. Comparative synthesis.

Family	Representative works/data sets	Strengths	Limitations	Typical metrics & notes
Classical ML (SVM/RF)	UNSW-NB15; CICIDS2017	Interpretable, fast training; good baselines	Limited in high-dimensional, nonstationary settings; feature-engineering intensive	Accuracy, precision, recall, and F1-score; sensitive to preprocessing
Deep Learning (CNN/RNN)	CICIDS2017 analyses	Captures nonlinear & temporal patterns; strong accuracy	Parameter-intensive; long convergence; data-hungry	Higher F1/ROC-AUC but high compute/latency in production
QML-Quantum Kernels	Havlíček *et al.*, Nature 2019	Rich feature maps via Hilbert-space lift; fewer hyper-params	Hardware access and noise issues; kernel-estimation cost	Promising margins; need apples-to-apples NIDS benchmarks
QML-VQCs (Hybrid)	VQA surveys (Cerezo 2021)	Trainable entanglement; compact parameterization; hybrid optimization	Barren plateaus; ansatz design critical; depth limited on NISQ	Track circuit depth, gate count, and fidelity alongside F1/ROC-AUC

2.1 Research Gaps

2.1.1 Lack of apples-to-apples benchmarking

Objective comparisons between QML models (quantum kernels or VQCs) and well-tuned classical SVM or neural-network baselines are limited. Differences in preprocessing, class-imbalance handling, data splits, and evaluation metrics make it difficult to determine whether observed gains arise from the quantum component or from the surrounding experimental pipeline.

2.1.2 Insufficient analysis of encoding strategies

Network-flow features can be represented using amplitude, angle, or IQP-style encodings, but systematic comparisons of these choices are limited. Their effects on

accuracy, convergence, robustness, qubit requirements, and shot budgets require further investigation.

2.1.3 Limited guidance on ansatz design for cybersecurity

There is limited guidance on how problem-inspired VQC topologies, entanglement patterns, and parameter-sharing strategies should be aligned with groups of network-flow features, such as header, temporal, and volumetric attributes. Poorly selected circuits may be either under-expressive or unnecessarily deep.

2.1.4 Limited evaluation of hardware and noise effects:

Many QML studies are evaluated using ideal simulators. Fewer studies quantify readout error, decoherence, two-qubit gate error, transpilation overhead, and reproducibility on NISQ hardware [14], [18], [19].

2.1.5 Opaque resource-accuracy trade-offs

Circuit depth, two-qubit gate count, number of shots, runtime, memory consumption, and predictive metrics such as F1-score and ROC-AUC are not consistently reported together. Consequently, the operational cost of a QML model relative to its detection performance is often unclear.

2.1.6 Class imbalance and rare attacks

The behavior of QML models on long-tail attack classes, such as infiltration and Heartbleed, remains insufficiently studied. Standardized protocols for resampling, cost-sensitive learning, SMOTE, and focal loss in hybrid QML pipelines are still lacking [22], [23].

2.1.7 Distribution shift and generalization

Cross-dataset and cross-day generalization are not routinely evaluated. Examples include training on UNSW-NB15 and testing on selected CICIDS2017 traffic, or training on one day of CICIDS2017 and testing on attacks collected on another day.

2.1.8 Explainability for QML-based IDS

Unlike established classical explanation tools such as SHAP and LIME, interpretation methods for VQCs and quantum kernels are still developing. Translating quantum observables into network features that are meaningful to security analysts remains an open problem [25].

2.1.9 Adversarial robustness of hybrid QML

The vulnerability of hybrid QML systems to evasion and poisoning attacks has not been studied extensively in NIDS applications. The effectiveness of countermeasures such as adversarial training and randomized smoothing also requires further evaluation [24].

2.1.10 Online and continual learning

IDSs must adapt to evolving threats. Continual and federated QML approaches must address catastrophic forgetting, privacy, communication overhead, and the cost of exchanging parameters, gradients, or expectation values.

2.1.11 Multimodal and multi-stage feature fusion

There is limited empirical evidence on entanglement-based fusion of heterogeneous cybersecurity cues, including flow statistics, TLS fingerprints, and system logs, or on the contribution of quantum correlations to the detection of coordinated and multi-stage intrusions.

2.1.12 Latency and energy profiling

End-to-end measurements of latency, throughput, and energy consumption for hybrid quantum-classical IDS pipelines are rarely reported under realistic operating conditions. Such profiling is essential before deployment in high-throughput security operations centers or edge environments.

2.2 Problem Statement

The increasing scale and diversity of cyber threats, including DDoS, data exfiltration, and zero-day attacks, place a substantial burden on traditional IDSs based on handcrafted rules or conventional machine-learning representations. Such systems may fail to capture highly nonlinear relationships among network-flow features or to adapt to multi-vector attacks. Although deep learning can improve detection performance, its model complexity, training cost, and data requirements may limit deployment in resource-constrained environments such as IoT networks and edge devices [1], [2].

QML provides an alternative approach by using superposition, entanglement, and interference to construct expressive feature representations in high-dimensional Hilbert spaces [7]-[15]. However, cybersecurity-oriented QML research still lacks standardized benchmarking, clearly justified encoding strategies, and sufficiently detailed evaluation on realistic datasets such as CICIDS2017 and UNSW-NB15 [4], [5], [16]. In addition, many hybrid studies do not report circuit-level resource measures, making reproducible comparison with classical methods difficult.

Therefore, this study develops a Quantum-Inspired Cyber Threat Detection (QCTD) framework that integrates amplitude-based feature encoding, entanglement-based feature fusion through a VQC, and a lightweight classical classifier. The framework is designed to provide a reproducible evaluation using standard security metrics, including accuracy, precision, recall, F1-score, and ROC-AUC. Quantum-resource characteristics, such as qubit count, circuit depth, gate count, and fidelity, should also be reported in future hardware-oriented evaluations.

3. PROPOSED METHODOLOGY

The proposed Quantum-Inspired Cyber Threat Detection (QCTD) framework combines quantum-classical learning components to detect multi-vector cyberattacks. It applies amplitude-based feature encoding, a variational quantum circuit (VQC) with entangling layers, and classical post-processing for final classification. The methodology consists of five principal stages, as illustrated in Figure-1.

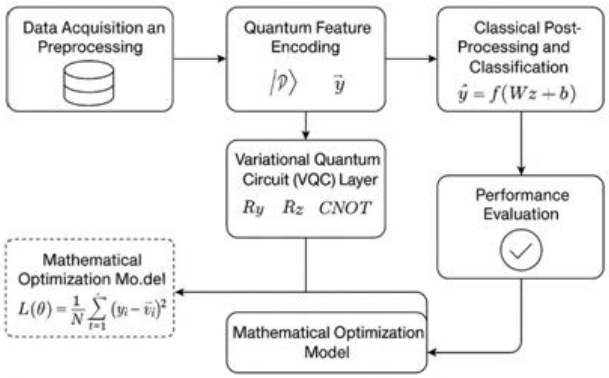

Figure-1. Quantum-Inspired Cyber Threat Detection (QCTD) framework.

A. Data Preprocessing and Acquisition

Network-intrusion datasets, including CICIDS2017 and UNSW-NB15, contain labeled benign traffic and attacks such as DDoS, PortScan, Botnet, and Infiltration. In the experiments reported in this paper, CICIDS2017 is used for model development and evaluation.

The feature vector is denoted by $X = \{x_1, x_2, ..., x_n\}$.

Continuous attributes are normalized, categorical values are encoded, and principal component analysis (PCA) is used to reduce the feature dimension to a size compatible with the available qubits, typically four to eight.

The class labels are denoted by $Y = \{y_1, y_2, ..., y_n\}$,

where each label represents benign or attack traffic. Stratified sampling is applied to preserve the class distribution across the training, validation, and test subsets.

B. Quantum Feature Encoding

Amplitude encoding maps a normalized n-dimensional classical feature vector into a quantum state using log2(n) qubits:

$|\psi\rangle = \sum (i = 1 \text{ to } N)\ x_i|i\rangle$,

where xi denotes the normalized amplitude associated with computational-basis state |i>. The encoded state satisfies the normalization condition so that the squared amplitudes sum to one.

This representation allows the feature vector to be processed through quantum-state evolution before measurement.

C. Variational Quantum Circuit (VQC) Layer

The quantum-encoded data are processed by a parameterized circuit composed of rotation gates (Ry and Rz) and entangling CNOT gates:

$U(\theta) = \prod_i R_\gamma(\theta_i)\ R\_z(\varphi_i)\ CNOT(i, i + 1)$.

The entangling layers support feature fusion by modeling correlations among encoded variables. Circuit parameters are optimized using a hybrid gradient-based method such as Adam or SPSA [20], [21].

The measured expectation values <Zi> form a compact feature embedding for the classical classification head.

D. Classical Post-Processing and Classification

The vector of measured quantum expectation values z in R^m is supplied to a small fully connected neural-network head for final classification:

$\hat{y} = f(Wz + b)$,

For binary classification, f(.) is a sigmoid activation; for multiclass classification, it may be replaced by a softmax activation.

This hybrid design combines expressive quantum feature transformation with a conventional and readily trainable classification layer.

E. Performance Evaluation

The Hybrid QCTD model is compared with Random Forest, XGBoost, and SVM-RBF using the same preprocessing pipeline and data partitions.

The predictive metrics are accuracy, precision, recall, F1-score, and ROC-AUC. Relevant implementation measures include the number of qubits, circuit depth, gate count, training time, and convergence behavior.

The quantum layer is implemented using PennyLane and Qiskit simulators [18], [19]. Evaluation on physical quantum hardware is reserved for future work.

F. Mathematical Optimization Model

The optimization objective can be expressed using the mean-squared classification loss: $L(\theta) = (1/N) \sum (i = 1 \text{ to } N) [y_i - \hat{y}_i(\theta)]^2$,

where y_i is the true label and $\hat{y}_i(\theta)$ is the model output obtained from the measured observable O using trainable parameters θ.

3.1 Experimental Setup and Parameters

The Hybrid Quantum-Inspired Cyber Threat Detection (Hybrid QCTD) model was evaluated on the CICIDS2017 dataset, which contains contemporary benign and malicious network-traffic patterns. Experiments were conducted on a workstation equipped with an Intel Core i7 processor, 32 GB RAM, and an NVIDIA RTX-3060 GPU with 12 GB memory. Python 3.10, TensorFlow, scikit-learn, XGBoost, PennyLane, and Qiskit were used to implement the classical and quantum components.

Continuous features were normalized, categorical features were encoded, and the data were divided into 70% training, 15% validation, and 15% test subsets. The Hybrid

QCTD architecture comprised a quantum feature encoder, a VQC with three entangling layers, and a fully connected classical classification head. The model was trained for 20 epochs using the Adam optimizer, a learning rate of 0.001, and a batch size of 256. Random Forest, XGBoost, and SVM-RBF were used as classical baselines, with hyperparameters selected by grid search. Each experiment was repeated three times, and the average results were reported. The evaluation metrics were accuracy, precision, recall, F1-score, and ROC-AUC.

Table-2. Experimental parameters for the Hybrid QCTD and baseline models.

Parameter	Value / Setting
Dataset	CICIDS2017
Train/Val/Test Split	70% / 15% / 15%
Computing Hardware	Intel i7 CPU, 32 GB RAM, NVIDIA RTX-3060 GPU
Quantum Backend	PennyLane + Qiskit Simulator
Classical Frameworks	TensorFlow, scikit-learn, XGBoost
Optimizer	Adam
Learning Rate	0.001
Batch Size	256
Epochs	20
Quantum Circuit Type	Variational Quantum Circuit (VQC)
Quantum Layers	3 entangling layers
Classical Layers	Fully connected dense layers
Baseline Models	Random Forest, XGBoost, SVM-RBF
Hyperparameter Tuning	Grid Search
Evaluation Metrics	Accuracy, Precision, Recall, F1-score, ROC-AUC

4. RESULTS AND DISCUSSION

The Hybrid QCTD model was evaluated on the CICIDS2017 test set using accuracy, precision, recall, F1-score, and ROC-AUC. The analysis includes training-loss convergence, validation performance, comparison with classical baselines, ROC analysis, multi-metric visualization, and a prediction-level example. The results are presented in the following subsections.

4.1 Hybrid QCTD Model Performance

Table-3. Hybrid QCTD performance metrics.

Accuracy	0.955
F1-score	0.877883311
Precision	0.810776942
Recall	0.957100592
ROC-AUC	0.993058597

Table-3 summarizes the test performance of the Hybrid QCTD model on CICIDS2017. The model achieved an accuracy of 0.9550, indicating that 95.5% of the evaluated flows were classified correctly. Its F1-score of 0.8779 reflects the balance between precision (0.8108) and recall (0.9571). The high recall indicates that the model detected a large proportion of attack samples, whereas the lower precision shows that false-positive predictions remain a concern. The ROC-AUC of 0.9931 demonstrates strong overall discrimination between benign and attack traffic across classification thresholds.

4.2 Training Loss Convergence Analysis

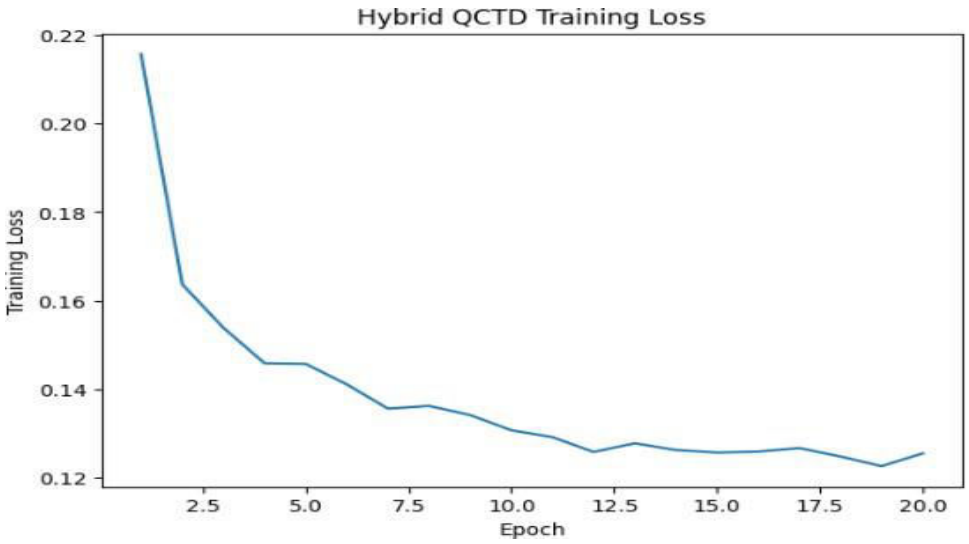

Figure-2. Hybrid QCTD Training Loss Curve.

Figure-2 shows the training-loss curve of the Hybrid QCTD model over 20 epochs. The loss decreases rapidly during the initial epochs, from approximately 0.215 to 0.165 by the third epoch, and then declines more gradually before stabilizing near 0.125-0.135. Minor fluctuations are consistent with stochastic optimization. The overall trend indicates stable convergence without visible divergence during the reported training period.

4.3 Validation Accuracy and F1-Score Performance

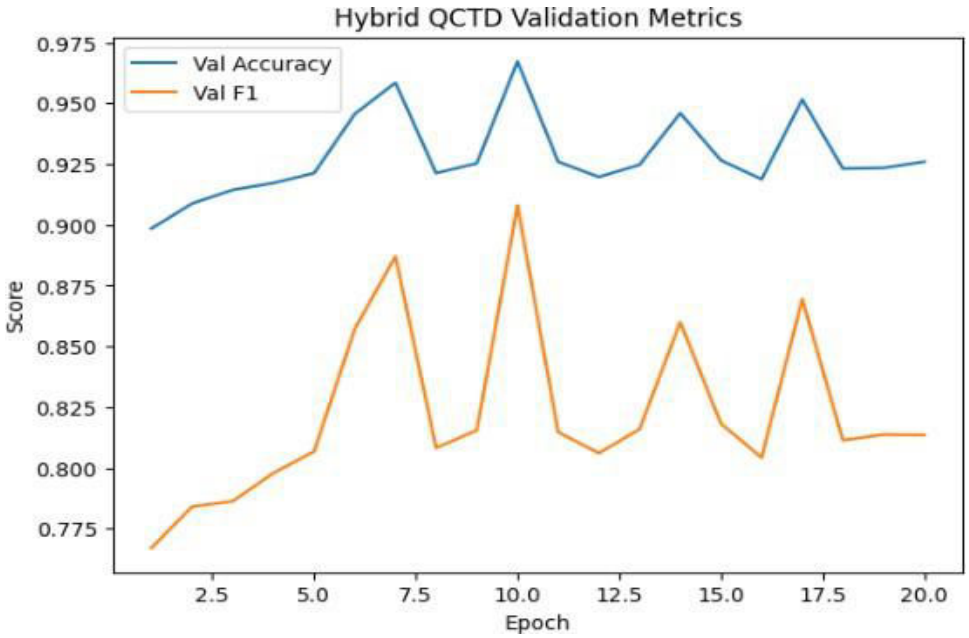

Figure-3. Hybrid QCTD Validation Accuracy and F1 Metrics.

Figure-3 presents the validation accuracy and F1-score over the training epochs. Validation accuracy increases from approximately 0.90 and generally remains between 0.92 and 0.97 after the initial epochs. The validation F1-score fluctuates more strongly, with peaks near 0.88-0.90 and later values around 0.82-0.84. These variations may reflect class imbalance and interactions between the variational quantum parameters and the classical weights. Overall, the model maintains strong validation accuracy and an F1-score above 0.80 throughout most of the training process. Together with Figure-2, these results indicate stable learning and reasonable generalization on the validation subset.

4.4 Comparative Evaluation with Classical Models

Table-4. Comparison of Metrics.

Model	Accuracy	F1	Precision	Recall	ROC-AUC
Hybrid QCTD	0.95500	0.877883	0.810777	0.957101	0.993059
Random Forest	0.99750	0.992582	0.995536	0.989645	0.999672
XGBoost	0.99800	0.994065	0.997024	0.991124	0.999952
SVM-RBF	0.92175	0.808093	0.690052	0.974852	0.990642

Table-4 compares the Hybrid QCTD model with Random Forest, XGBoost, and SVM-RBF. Random Forest and XGBoost achieve the strongest overall performance, with accuracies of 0.9975 and 0.9980 and F1-scores of 0.9926 and 0.9941, respectively. Their ROC-AUC values are also near perfect. SVM-RBF obtains a lower accuracy of 0.9218 and F1-score of 0.8081, primarily because its precision is substantially lower than that of the other models.

The Hybrid QCTD model achieves an accuracy of 0.9550, an F1-score of 0.8779, and an ROC-AUC of 0.9931. It outperforms SVM-RBF in accuracy, precision, F1-score, and ROC-AUC, but it does not exceed Random Forest or XGBoost on the reported predictive metrics. Its recall of 0.9571 is high in absolute terms, although it is lower than the recall of all three classical baselines in Table-4. These results show that the hybrid model is feasible and competitive with SVM-RBF, while further optimization is required to match the ensemble methods.

4.5 ROC–AUC Comparison across Models

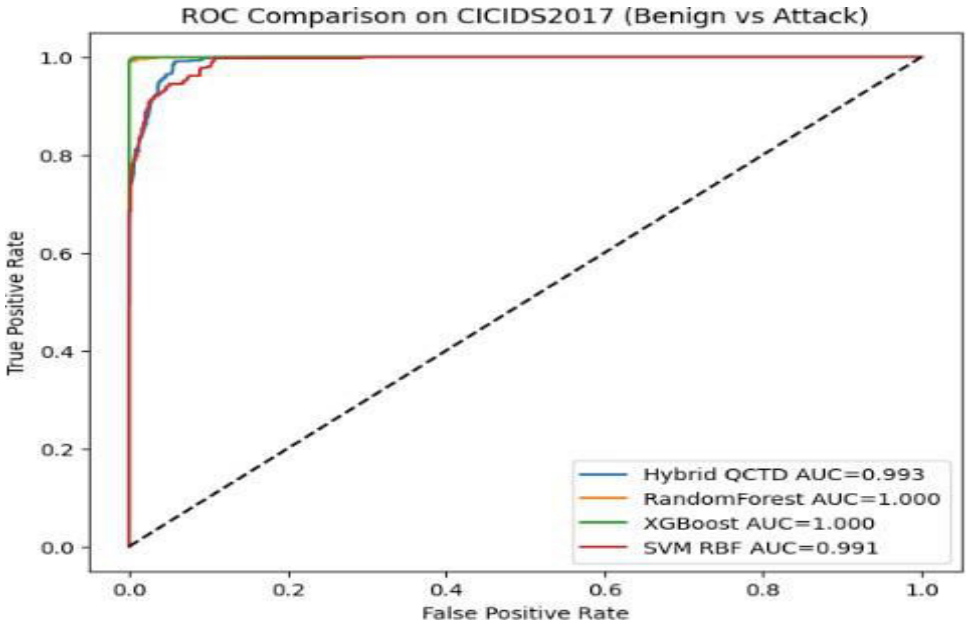

Figure-4. ROC comparison on CICIDS2017 (benign vs. attack).

Figure-4 compares the receiver operating characteristic (ROC) curves of the Hybrid QCTD model and the classical baselines. The Hybrid QCTD model achieves an AUC of 0.993, indicating strong separation between benign and attack traffic. Random Forest and XGBoost achieve higher AUC values of approximately 1.000, while SVM-RBF achieves an AUC of 0.991. The ROC analysis therefore confirms that all four models have strong ranking ability, with the ensemble methods providing the best discrimination on this test set.

4.6 Multi-Metric Comparison of Hybrid QCTD and Classical Models

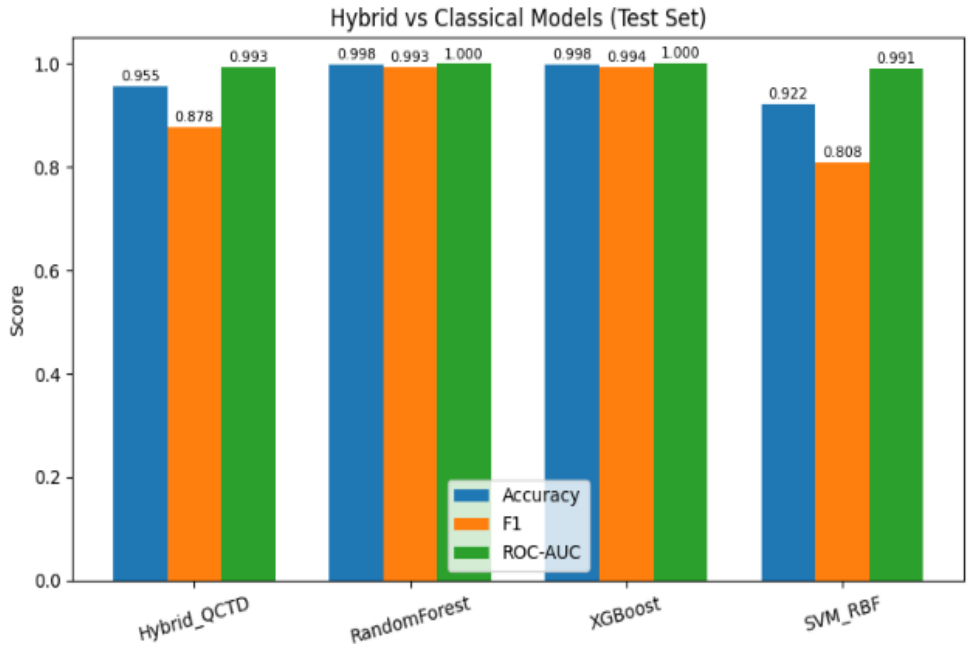

Figure-5. Multi-metric comparison of the Hybrid QCTD and classical models.

Figure-5 compares accuracy, F1-score, and ROC-AUC across the four models. Hybrid QCTD achieves values of 0.955, 0.878, and 0.993, respectively. Random Forest and XGBoost lead on all three metrics, whereas SVM-RBF records the lowest accuracy and F1-score. The figure shows that the hybrid model performs better than SVM-RBF on accuracy and F1-score but remains below the ensemble baselines. Because parameter count, runtime, memory use, and end-to-end latency are not reported in the present study, no definitive computational-efficiency advantage can be concluded from these results alone.

4.7 Classical Baseline Performance Metrics

Table-5. Classical baseline performance metrics on the test set.

Model	Accuracy	F1-score	Precision	Recall	ROC-AUC
Random Forest	0.99750	0.992582	0.995536	0.989645	0.999672
XGBoost	0.99800	0.994065	0.997024	0.991124	0.999952
SVM (RBF)	0.92175	0.808093	0.690052	0.974852	0.990642

Table-5 presents the standalone performance of the classical baselines. XGBoost provides the best overall result, with an accuracy of 0.9980, an F1-score of 0.9941, a precision of 0.9970, a recall of 0.9911, and an ROC-AUC of 0.999952. Random Forest performs similarly, with an accuracy of 0.9975 and an F1-score of 0.9926. These results demonstrate the effectiveness of ensemble learning for structured network-flow data.

SVM-RBF achieves an accuracy of 0.9218 and an F1-score of 0.8081. Its recall is high (0.9749), but its precision is substantially lower (0.6901), indicating a larger number of false-positive predictions. Its ROC-AUC of 0.9906 remains strong, but its thresholded classification performance is less balanced than that of the ensemble methods and the Hybrid QCTD model.

4.8 Random Forest Performance Visualization

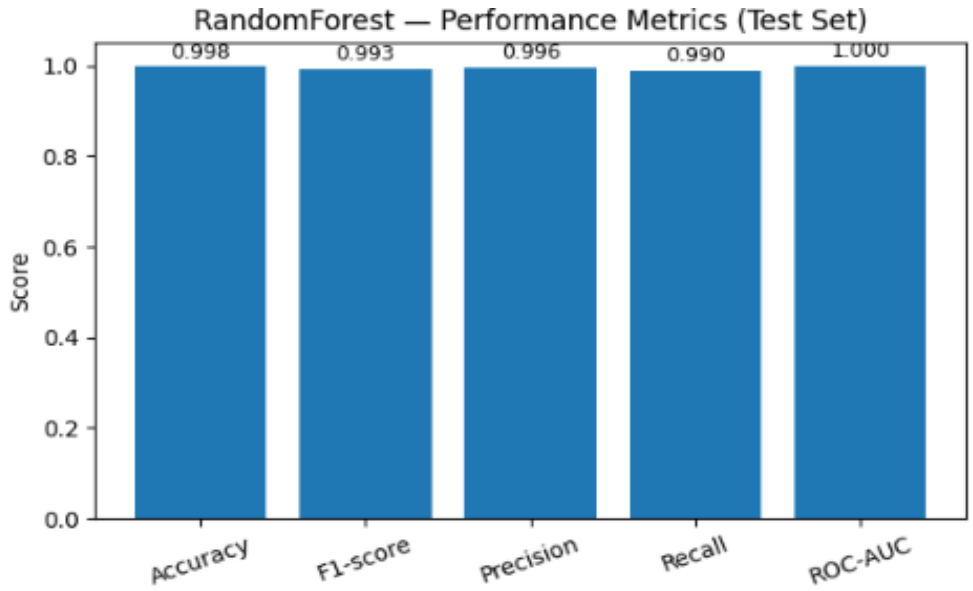

Figure-6. Random forest performance metrics.

Figure-6 summarizes the Random Forest results. The model achieves an accuracy of 0.9975, an F1-score of 0.9926, a precision of 0.9955, a recall of 0.9896, and an ROC-AUC of 0.9997. The consistently high values indicate very low misclassification and strong separation between benign and malicious flows on the test set.

4.9 XGBoost Performance Visualization

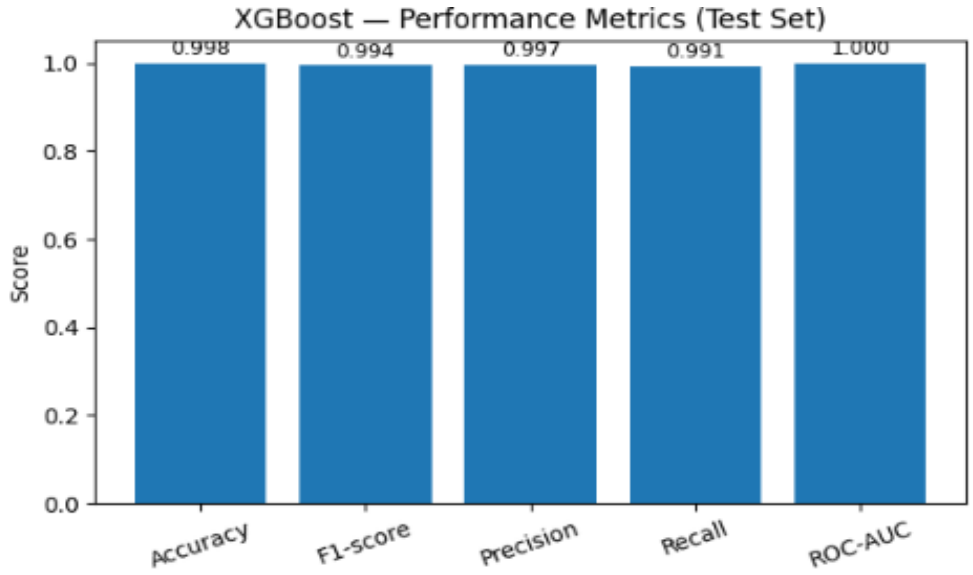

Figure-7. XGBoost performance metrics.

Figure-7 summarizes the XGBoost results. The model achieves an accuracy of 0.9980, an F1-score of 0.9941, a precision of 0.9970, a recall of 0.9911, and an ROC-AUC of 0.99995. XGBoost is therefore the strongest classical baseline in the reported experiment, with near-perfect discrimination and a well-balanced precision-recall profile.

4.10 SVM-RBF Performance Visualization

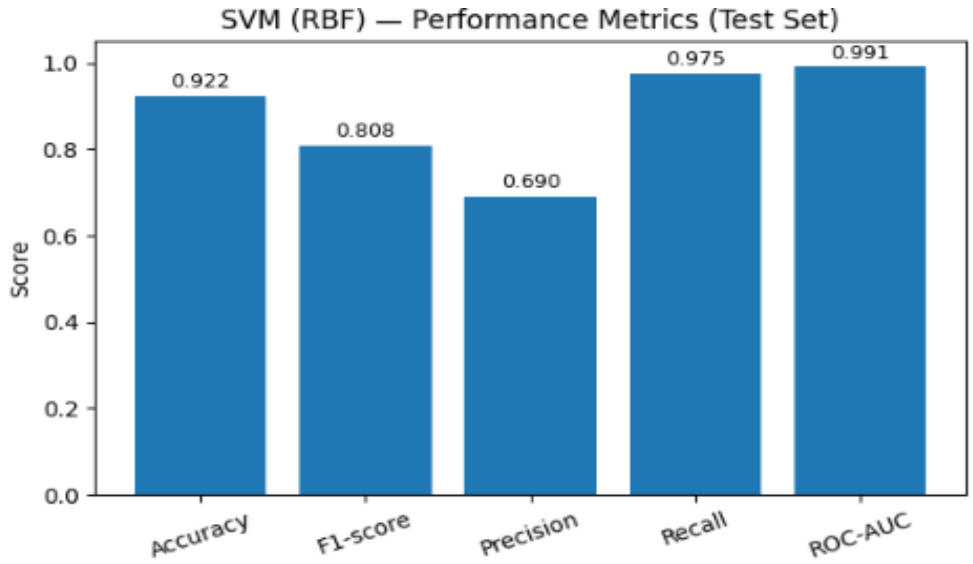

Figure-8. SVM-RBF performance metrics.

Figure-8 summarizes the SVM-RBF results. The model achieves an accuracy of 0.9218, an F1-score of 0.8081, a precision of 0.6901, a recall of 0.9749, and an ROC-AUC of 0.9906. The high recall shows that the classifier detects most attack samples, but the comparatively low precision indicates frequent false positives. This imbalance explains its lower accuracy and F1-score relative to the other evaluated models.

4.11 Prediction-Level Evaluation across Models

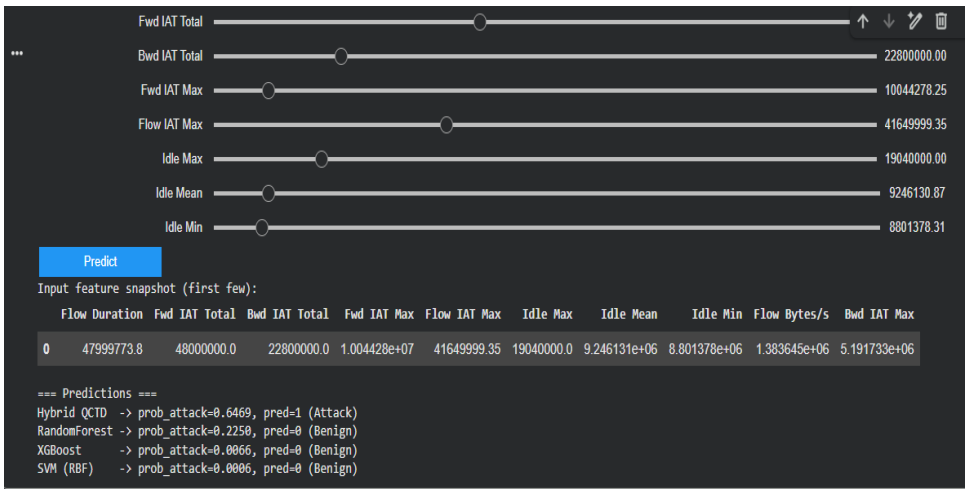

Figure-9. Prediction output snapshot.

Figure-9 shows a prediction snapshot for one network-flow input. The Hybrid QCTD model assigns an attack probability of 0.6469 and predicts Attack. Random Forest assigns 0.2250, XGBoost assigns 0.0006, and SVM-RBF assigns 0.0000; all three classical models therefore predict Benign for this sample.

This example illustrates disagreement among the models and suggests that the Hybrid QCTD decision boundary responds differently to the selected temporal and volumetric features. However, the snapshot does not display the ground-truth class and therefore cannot, by itself, establish which prediction is correct. Prediction-level examples should be interpreted together with the aggregate test metrics and a confusion-matrix or error analysis.

The experimental results show that the Hybrid QCTD framework can learn useful representations for intrusion detection. Its recall of 0.9571 and ROC-AUC of 0.9931 indicate strong sensitivity and ranking performance, while its precision of 0.8108 identifies false-positive reduction as an important area for improvement.

The training-loss and validation curves provide additional evidence of stable learning behavior. Training loss decreases rapidly during the early epochs and then converges without visible divergence. Validation accuracy remains high, although the F1-score fluctuates, indicating

that performance is sensitive to the balance between attack and benign samples.

Comparison with the classical baselines clarifies the current position of the Hybrid QCTD model. Random Forest and XGBoost achieve substantially higher accuracy, precision, recall, and F1-score on this dataset. Hybrid QCTD nevertheless outperforms SVM-RBF on accuracy, precision, F1-score, and ROC-AUC. The present results support the feasibility of the hybrid architecture, but parameter count, runtime, memory, circuit depth, and hardware noise must be reported before claims of computational or resource efficiency can be validated.

SVM-RBF exhibits a different error profile: it achieves high recall but low precision, indicating that it detects most attacks at the cost of more false alarms. Hybrid QCTD provides a more balanced precision-recall trade-off than SVM-RBF, although the ensemble methods remain superior on both measures.

The ROC and multi-metric comparisons confirm that Hybrid QCTD has strong discriminative ability but does not surpass the ensemble baselines. Its performance is most appropriately interpreted as a promising proof of concept that requires further optimization and evaluation under realistic deployment constraints.

The prediction-level snapshot demonstrates that the hybrid and classical models can respond differently to the same input. Because the true label is not shown, the example should not be presented as evidence that the hybrid prediction is correct. A larger error analysis, including confusion matrices and representative true-positive, false-positive, true-negative, and false-negative cases, would provide stronger support.

Overall, the Hybrid QCTD model achieves strong intrusion-detection performance and demonstrates the potential of entanglement-based quantum feature fusion. Further work should focus on improving precision, comparing resource consumption, testing robustness under distribution shift, and evaluating the model on noisy quantum hardware.

5. CONCLUSIONS

The proposed Hybrid Quantum-Inspired Cyber Threat Detection (Hybrid QCTD) framework demonstrates the feasibility of combining amplitude-based encoding, entangling variational quantum circuits, and a classical classification head for intrusion detection. On CICIDS2017, the model achieved an accuracy of 0.9550, an F1-score of 0.8779, a precision of 0.8108, a recall of 0.9571, and an ROC-AUC of 0.9931. These results indicate strong attack-detection sensitivity and discrimination, while the lower precision shows that false positives remain an important limitation.

Random Forest and XGBoost achieved higher values on all reported predictive metrics, whereas Hybrid QCTD outperformed SVM-RBF in accuracy, precision, F1-score, and ROC-AUC. The training and validation trends indicate stable convergence, but the present study does not provide sufficient parameter, runtime, energy, or hardware-noise measurements to establish a computational-efficiency advantage. Future work should evaluate hardware-efficient circuit designs, noisy NISQ execution, multimodal cyber-threat fusion, cross-dataset generalization, and continual learning for evolving attacks.

REFERENCES

[1] H. Hindy, D. Brosset, E. Bayne, *et al*. 2022. A taxonomy and survey of intrusion detection system design techniques, network threats and datasets. IEEE Commun. Surveys Tuts. 24(1): 593-630.

[2] J. Kim, J. Kim, H. Kim and S. Kim. 2018. Deep learning for intrusion detection systems. IEEE Access. 6: 21954-21961.

[3] R. Vinayakumar, M. Alazab, K. P. Soman, *et al*. 2019. Deep learning approach for intelligent intrusion detection system. IEEE Access. 7: 41525-41550.

[4] N. Moustafa and J. Slay. 2015. UNSW-NB15: A comprehensive data set for network intrusion detection systems. In Proc. IEEE MILCOM. pp. 1-6.

[5] I. Sharafaldin, A. H. Lashkari and A. A. Ghorbani. 2018. Toward generating a new intrusion detection dataset and intrusion traffic characterization. In Proc. Int. Conf. Inf. Syst. Secur. Priv. (ICISSP). (CICIDS2017 Dataset, Canadian Institute for Cybersecurity, University of New Brunswick.)

[6] Y. Mirsky, T. Doitshman, Y. Elovici, and A. Shabtai. 2018. Kitsune: An ensemble of autoencoders for online network intrusion detection. in Proc. NDSS.

[7] V. Havlíček, A. D. Córcoles, K. Temme, et al. 2019. Supervised learning with quantum-enhanced feature spaces. Nature. 567: 209-212.

[8] M. Cerezo, A. Arrasmith, R. Babbush, *et al*. 2021. Variational quantum algorithms. Nat. Rev. Phys. 3(9): 625-644.

[9] P. Rebentrost, M. Mohseni and S. Lloyd. 2014. Quantum support vector machine for big data classification. Phys. Rev. Lett. 113(13): 130503.

[10] S. Lloyd, M. Mohseni and P. Rebentrost. 2013. Quantum algorithms for supervised and unsupervised learning. arXiv:1307.0411.

[11] J. Biamonte, P. Wittek, N. Pancotti, *et al*. 2017. Quantum machine learning. Nature. 549: 195-202.

[12] M. Schuld and F. Petruccione. 2021. Machine Learning with Quantum Computers. 2nd ed., Springer.

[13] M. Benedetti, E. Lloyd, S. Sack and M. Fiorentini. 2019. Parameterized quantum circuits as machine learning models. Quantum Sci. Technol. 4(4): 043001.

[14] J. Preskill. 2018. Quantum computing in the NISQ era and beyond. Quantum. 2: 79.

[15] M. A. Nielsen and I. L. Chuang. 2010. Quantum Computation and Quantum Information. Cambridge Univ. Press.

[16] V. Abdi, N. Tahayori and R. Safavi-Naini. 2022. Quantum-based intrusion detection for next-generation networks. IEEE Access. 10: 114523-114534.

[17] A. W. Harrow and A. Montanaro. 2017. Quantum computational supremacy. Nature. 549: 203-209.

[18] H. Abraham *et al*. 2019. Qiskit: An open-source framework for quantum computing. Zenodo.

[19] V. Bergholm, J. Izaac, M. Schuld, *et al*. 2018. PennyLane: Automatic differentiation of hybrid quantum-classical computations. arXiv:1811.04968.

[20] J. C. Spall. 1992. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. IEEE Trans. Autom. Control. 37(3): 332-341.

[21] D. P. Kingma and J. Ba. 2015. Adam: A method for stochastic optimization. In Proc. ICLR.

[22] N. V. Chawla, K. W. Bowyer, L. O. Hall and W. P. Kegelmeyer. 2002. SMOTE: Synthetic minority over-sampling technique. J. Artif. Intell. Res. 16: 321-357.

[23] T.-Y. Lin, P. Goyal, R. Girshick, K. He and P. Dollár. 2017. Focal loss for dense object detection. In Proc. ICCV. pp. 2980-2988.

[24] J. Cohen, E. Rosenfeld and Z. Kolter. 2019. Certified adversarial robustness via randomized smoothing. in Proc. ICML. pp. 1310-1320.

[25] S. M. Lundberg and S.-I. Lee. 2017. A unified approach to interpreting model predictions. In Proc. NeurIPS. pp. 4765-4774.