IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 7, July 2026

DOI: 10.17148/IJARCCE.2026.15723

Quantum-Enhanced Intrusion Detection Systems for IoT Networks: Trends, Challenges, and Future Directions

Sivasubramanyam Medasani[1], M Neavruth Sai[2], Jhanavi C[3], Lakshmi B[4]

Professor, Computer Science and Engineering, K. S. School of Engineering and Management, Bengaluru, India[1]

Student, Computer Science and Engineering, K. S. School of Engineering and Management, Bengaluru, India[2]

Student, Computer Science and Engineering, K. S. School of Engineering and Management, Bengaluru, India[3]

Student, Computer Science and Engineering, K. S. School of Engineering and Management, Bengaluru, India[4]

Abstract: This survey examines quantum-enhanced machine learning techniques for intrusion detection in Internet of Things (IoT) and fog environments, synthesizing recent work on hybrid quantum-classical models, quantum convolutional neural networks (QCNNs), quantum support vector machines (QSVMs), and quantum-assisted clustering. We review five recent primary studies that apply quantum methods to IoT security, summarize their methodological choices (encoding, circuit ansatz, classical-quantum interfaces), and evaluate reported benefits and limitations with respect to accuracy, computational cost, and deployability in resource-constrained settings. The literature indicates potential gains in feature expressivity and false-negative reduction from quantum encodings and variational circuits, but practical barriers remain: limited qubit counts, noisy hardware, encoding overhead, and scarce real-world deployment studies. We identify concrete research gaps and propose directions for lightweight encodings, noise-aware hybrid architectures, federated quantum-assisted Intrusion Detection System (IDS), and standardized benchmarks for reproducible evaluation. The survey aims to guide researchers and practitioners toward scalable, explainable, and deployable quantum-enhanced intrusion detection for IoT.

Keywords: Internet of Things, Intrusion Detection System, Quantum Machine Learning, Quantum Convolutional Neural Network, Quantum Support Vector Machine, Quantum K-Means, Hybrid Quantum Neural Network.

I. INTRODUCTION

The proliferation of IoT devices has expanded attack surfaces and created stringent requirements for IDS: high accuracy, low false negatives, and operation under constrained compute and energy budgets at the edge and fog layers [3,11]. Classical machine learning and deep learning methods (CNNs, LSTMs, ensembles) have advanced IDS performance but face limits in scalability, feature expressivity, and robustness to evolving attacks in highly heterogeneous IoT traffic [3, 4]. Recent research explores quantum machine learning (QML) and hybrid quantum classical architectures as a means to enhance feature representation, exploit high-dimensional Hilbert spaces, and potentially reduce model parameter counts while improving detection metrics [1, 2, 5, 6, 10].

This survey synthesizes and critically evaluates five contemporary studies that apply quantum-enhanced methods to IoT intrusion detection, comparing model families (QCNN, QSVM, quantum clustering, hybrid ensembles), datasets, encoding strategies, and reported outcomes [1, 2, 3, 4, 5]. Our objective is to map the current state of the art, identify practical and theoretical gaps, and propose research directions that align with the constraints of IoT and fog computing deployments [2, 3, 14].

II. BACKGROUND

A. IoT Intrusion Detection

Network- and host-based IDS monitor traffic and system events to detect anomalies and known attack signatures; in IoT contexts, IDS must operate across heterogeneous devices, limited compute nodes, and intermittent connectivity, often requiring edge/fog pre-processing and cloud-level aggregation [3,14]. Effective IDS pipelines typically include feature extraction, dimensionality reduction, classification or clustering, and post-processing for alerts and mitigation [3].

 This work is licensed under a Creative Commons Attribution 4.0 International License

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 7, July 2026
DOI: 10.17148/IJARCCE.2026.15723

B. Quantum Machine Learning Concepts

QML leverages quantum states and unitary operations to represent and process classical data in high-dimensional Hilbert spaces [6,8,9]. Quantum feature maps (e.g., angle encoding, amplitude encoding, ZZFeatureMap) and parameterized quantum circuits (PQC) as variational models [4, 2, 5]. Foundational proposals for quantum-enhanced supervised and unsupervised learning include kernel-based QSVMs and quantum clustering algorithms [6,7] variational quantum circuits underpin hybrid classifiers and QCNN architectures [8, 9, 12].

C. Quantum Feature Encoding and PQC

Encoding classical vectors into quantum states is a critical design choice: angle (rotation) encoding maps features to single-qubit rotations and is shallow and Noisy Intermediate-Scale Quantum (NISQ)-friendly; mplitude encoding packs data amplitudes into multi-qubit states but requires complex state preparation; stereographic projections and inverse stereographic projection (ISP) have been proposed to preserve geometric relations before angle encoding [2, 5]. PQCs combine parameterized rotations and entangling gates (CNOT/CX) to form expressive ans¨atze that are trained via classical optimizers using parameter-shift or finite-difference gradients [1, 4, 12].

D. QCNN, QSVM, and Quantum Clustering

QCNNs adapt convolutional and pooling concepts to quantum circuits, using local unitaries and measurements to reduce qubit counts and capture hierarchical features [4, 1]. QSVMs implement kernel evaluations in quantum feature spaces to separate classes with potentially fewer resources for certain kernels [4]. Quantum clustering approaches (e.g., QK-Means) estimate similarity via fidelity or SWAP-test circuits and can exploit Hilbert-space geometry for unsupervised anomaly detection [2, 7].

III. LITERATURE REVIEW

We organize the review thematically, focusing on model families, encoding and preprocessing, datasets and evaluation, computational cost and scalability, and suitability for IoT/fog deployments. Each subsection compares the five primary studies and situates them relative to foundational QML work.

A. Model Families and Architectures

Hybrid QCNN-based classifiers combine quantum feature extraction with classical dense layers to produce end-to-end IDS models; Amara et al. present a QCNN hybrid that encodes Principal Component Analysis (PCA)-reduced features into eight qubits, applies variational entangling layers, and uses classical layers for final classification, reporting improved false-negative rates relative to a classical CNN baseline [1]. Srinivasan et al. focus on unsupervised quantum clustering (QK-Means) using SWAP-test fidelity estimates and ISP-based encodings to preserve geometry for anomaly detection [2]. Bharathi et al. propose a hybrid quantum neural network (HQNN) framework that integrates PCA, ZZFeatureMap style encodings, and multiple QML classifiers (QSVM, QNN variants) for IoT datasets [5]. Kalinin and Krundyshev evaluate QSVM and QCNN variants for intrusion detection and emphasize quantum kernels and QCNN expressivity for large-scale inputs [4]. Tawfik's ensemble work, while classical, provides an important baseline and edge-cloud pipeline design that quantum-enhanced methods must integrate with to be practical in fog settings [3].

B. Datasets and Evaluation Protocols

The surveyed quantum studies use a mix of IoT-focused and general IDS datasets. Amara et al. evaluate on Botnet of Things dataset (BoT-IoT) and a custom healthcare IoT dataset, reporting near-99% accuracy and reduced false negatives for QCNN versus CNN baselines [1]. Srinivasan et al. test QK-Means on RT-IoT2022 and synthetic high-dimensional data to demonstrate improved silhouette scores in higher dimensions [2]. Bharathi et al. evaluate across BoT-IoT, Army Cyber Institute Internet of Things dataset (ACI IoT), and ROUT-4-2023 to show cross-dataset applicability of HQNN variants [5]. Kalinin and Krundyshev report experiments on custom stream datasets tailored for QML processing [4]. Tawfik's ensemble study uses NSL-KDD, UNSW-NB15, and AWID as benchmarks and highlights the importance of realistic, heterogeneous datasets for fog/IoT IDS evaluation [3]. Across studies, evaluation metrics include accuracy, precision, recall, F1, silhouette score (for clustering), and confusion-matrix analyses [1, 2, 3, 4, 5].

C. Feature Engineering and Quantum Encoding

Dimensionality reduction (PCA) is a common preprocessing step to match limited qubit budgets; Amara et al. reduce to eight PCA components before angle encoding [1]. Srinivasan et al. emphasize ISP followed by angle encoding to preserve geometry for clustering fidelity computations [2]. Bharathi et al. use PCA plus ZZFeatureMap-style encodings to exploit entanglement and kernel effects [5]. Kalinin and Krundyshev discuss stream aggregation and tailored encodings to make packet streams amenable to quantum processing [4]. These choices reflect trade-offs: angle encodings are shallow and NISQ-friendly but may require careful scaling; amplitude encodings are compact but costly to prepare [8, 9].

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 7, July 2026

DOI: 10.17148/IJARCCE.2026.15723

D. Performance Claims and Computational Cost

Reported performance gains are promising but heterogeneous. Amara et al. report QCNN matching or slightly exceeding CNN accuracy while reducing false negatives and parameter counts, at the expense of higher per-epoch simulation time due to quantum circuit simulation overhead [1]. Srinivasan et al. show QK-Means outperforming classical K-Means in high-dimensional regimes (beyond 16 dimensions) in silhouette scores, suggesting quantum inner-product estimation benefits [2, 7]. Bharathi et al. report improved classification metrics across multiple IoT datasets using HQNN variants, including experiments on simulators and limited real Quantum Processing Unit (QPUs) [5]. Kalinin and Krundyshev claim QML methods can process large inputs with competitive accuracy and potential speedups under certain assumptions [4]. However, all quantum studies note increased training overhead when using simulators and emphasize that real quantum hardware constraints (qubit counts, noise) limit direct speed advantages today [1, 2, 4, 5, 11, 12].

E. Scalability and Suitability for IoT/Fog Environments

Tawfik's ensemble pipeline highlights the need for edge preprocessing and cloud-level ensembles to meet fog constraints; quantum-enhanced modules must therefore be lightweight and interoperable with such pipelines [3]. Amara et al. argue QCNNs with low parameter counts are promising for edge-adjacent deployment when quantum accelerators become available, but current simulations are costly [1]. Srinivasan et al. demonstrate that quantum clustering maintains stability with increasing data volume in simulation, suggesting potential for streaming IDS if encoding and distance estimation are optimized [2]. Bharathi et al. discuss PCA-based reduction and repetition codes for error mitigation to make HQNNs more practical for IoT datasets [5]. Overall, suitability depends on encoding overhead, circuit depth, and the availability of near-term quantum resources or efficient simulators integrated into fog/cloud pipelines [3, 4, 11].

F. Limitations Identified by the Authors

Each primary study acknowledges limitations: Amara et al. cite hardware deployment constraints and simulation overhead [1]; Srinivasan et al. note that simulated backends may not reflect noisy hardware and that SWAP-test circuits scale with qubit counts [2]; Bharathi et al. emphasize the need for more real-QPU validation and error mitigation [5]; Kalinin and Krundyshev discuss data-preparation bottlenecks and assumptions for quantum speedups [4]; Tawfik highlights the importance of edge-cloud orchestration and feature selection for practical IDS, which quantum modules must respect [3].

IV. COMPARATIVE ANALYSIS

Table I highlights complementary emphases: quantum studies focus on encoding and circuit design to exploit Hilbert space structure, while Tawfik provides a practical fog/edge pipeline that quantum modules must integrate with for real deployments [1, 2, 3, 4, 5].

Table I: Comparative summary of the five primary studies

Method	Quantum/Hybrid	Dataset	Strengths	Limitations
HQNN variants (QSVM, QNN) [5]	Hybrid QNNs; PCA + Z	BoT-IoT; ACI IoT; ROUT-4-2023	Cross-dataset evaluation; error-mitigation discussion	Need more real-QPU experiments; encoding cost
QCNN + classical NN [1]	Hybrid QCNN (angle encoding, 8 qubits)	BoT-IoT; custom healthcare IoT	Low parameter count; reduced false negatives; end-to-end pipeline	Simulator overhead; hardware deployment constraints
QK-Means Clustering [2]	Quantum clustering (SWAP test ISP)	RT-IoT2022; synthetic	Preserves geometry; better silhouette in high-D	SWAP-test scaling; simulated backend only
Transformer CNN-LSTM ensemble [3]	Classical ensem ble; edge-cloud pipeline	NSL-KDD; UNSW NB15; AWID	Strong cloud edge design; feature selection	Classical only; motivates quantum inte gration
QSVM, QCNN [4]	QMLclassifiers; stream dataset encoding	Custom stream dataset	Theoretical speedup discus sion; stream oriented design	Assumptions for speedups; limited hard ware validation

 This work is licensed under a Creative Commons Attribution 4.0 International License

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 7, July 2026
DOI: 10.17148/IJARCCE.2026.15723

V. RESEARCH GAPS

Despite encouraging results, several gaps persist in the literature:

- **Hardware validation at scale:** Most quantum IDS studies rely on simulators or very small quantum processors. Simulators are useful for testing algorithms but cannot replicate the noise, decoherence, and connectivity issues of real hardware. This makes it difficult to confirm whether reported improvements in accuracy or false-negative reduction will hold in practice. Large-scale validation on noisy intermediate-scale quantum (NISQ) devices is essential to establish robustness and practical speedups [1,5,11,13].

- **Encoding overhead and standardization:** Encoding classical IoT traffic into quantum states is computationally expensive and varies widely across studies. Some approaches use angle encoding, others rely on PCA combined with feature maps, while others experiment with stereographic projections. These differences make results inconsistent and hinder fair comparison. A standardized, lightweight encoding scheme tailored for streaming IoT traffic is still missing, which prevents reproducibility and increases overhead [2,5,8,9].

- **Benchmarking and reproducibility:** Studies use heterogeneous datasets such as BoT-IoT, NSL-KDD, UNSW-NB15, and RT-IoT2022. While this diversity demonstrates applicability, it also makes comparisons unreliable. Without standardized IoT/fog benchmarks, results cannot be fairly compared across quantum IDS methods, slowing progress toward practical deployment [1,3].

- **Online and adaptive IDS:** Current quantum IDS approaches are static and do not adapt to evolving attack patterns. In IoT environments, traffic patterns change rapidly, requiring systems that can handle concept drift and continuous learning. Adaptive quantum IDS capable of streaming detection is essential but remains underexplored [2,4].

- **Explainability and interpretability:** Quantum models often act as "black boxes," making it difficult for analysts to understand why a detection decision was made. Without interpretability, analyst trust and adoption will remain limited. Methods for interpretable quantum-assisted IDS are underdeveloped, and adapting classical explainability techniques to quantum systems is a critical gap [5,1,15].

VI. CHALLENGES AND LIMITATIONS

Practical and theoretical barriers constrain near-term adoption of quantum-enhanced IDS:

A. Data encoding and preparation: Preparing IoT traffic for quantum circuits is not straightforward. Raw IoT data streams are often high-dimensional and noisy, requiring preprocessing before they can be mapped into quantum states. Techniques such as PCA and ISP reduce dimensionality, making data fit into limited qubit budgets. However, this reduction can inadvertently remove subtle attack signatures, leading to missed detections. The challenge lies in balancing efficiency with fidelity — encoding must be lightweight enough for near-term devices yet expressive enough to capture complex attack patterns [1,2,5].

B. Circuit depth, noise, and decoherence: Parameterized quantum circuits (PQCs) are central to quantum machine learning because they provide expressive feature transformations. Yet, deeper circuits require more entangling gates, which amplify noise and increase susceptibility to decoherence. This is particularly problematic for NISQ devices, where error rates remain high. Error mitigation strategies, such as repetition codes or zero-noise extrapolation, can reduce the impact of noise, but they also increase resource demands and sampling costs. As a result, practical deployment of deep PQCs is constrained by hardware limitations [5,4,9].

C. Limited qubit counts and connectivity: IoT datasets often contain hundreds of features, but current QPUs typically support only tens of qubits with limited connectivity between them. Mapping large feature sets to such restricted hardware requires aggressive dimensionality reduction or hybridization with classical models. While this makes experiments feasible, it limits the representational benefits of quantum encodings. The inability to fully exploit high-dimensional Hilbert spaces reduces the potential advantage of quantum IDS in real-world scenarios [1,2,11].

D. Computational overhead and simulator costs: Because real quantum hardware is scarce and limited, many studies rely on classical simulators to test quantum algorithms. Simulating quantum circuits is computationally

 This work is licensed under a Creative Commons Attribution 4.0 International License

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 7, July 2026
DOI: 10.17148/IJARCCE.2026.15723

expensive, especially as the number of qubits and circuit depth increase. Reported improvements in accuracy or false-negative reduction often come with significantly higher training times. This overhead complicates fair comparisons with optimized classical models, which can achieve similar or better performance with lower resource costs [1,4,12].

E. **Deployment and integration complexity:** Even if quantum modules show promise in controlled experiments, integrating them into operational IoT and fog computing pipelines is complex. Such integration requires orchestration across edge devices, secure data transfer, and strict latency guarantees. Current studies rarely demonstrate end-to-end deployments in realistic IoT environments. Without practical integration frameworks, quantum IDS remains confined to research prototypes rather than deployable solutions [3,5,14].

VII. FUTURE RESEARCH DIRECTIONS

Grounded in the surveyed literature, we propose the following directions:

A. **Lightweight, task-aware quantum encodings:** Future work should design encoding schemes that balance expressivity and preparation cost for streaming IoT features. Hybrid approaches such as ISP-angle maps or sparse amplitude encodings could reduce overhead while preserving fidelity. Benchmarking across multiple datasets will help identify trade-offs [2,5,7].

B. **Noise-aware hybrid architectures:** Hybrid pipelines where shallow PQCs perform feature enrichment and classical layers handle classification can reduce error rates while maintaining accuracy. Adaptive circuit depth tuned to hardware noise profiles will make quantum IDS more practical [1,5,12].

C. **Federated and distributed quantum-assisted IDS:** Federated learning paradigms can allow edge nodes to preprocess locally and share quantum-enhanced kernels with fog/cloud aggregators. This approach preserves privacy, reduces communication overhead, and supports scalability in distributed IoT networks [3,4,14].

D. **Explainable quantum models:** Developing interpretable surrogates and attribution methods for quantum outputs will improve analyst trust and incident response workflows. Visualization tools for quantum feature maps could help analysts understand model decisions [5,1,15].

E. **Realistic benchmarks and reproducible pipelines:** Establishing standardized IoT/fog benchmark suites and open-source pipelines that include preprocessing, encoding, and hybrid training scripts will enable fair comparisons and accelerate progress [3,2].

F. **Streaming and adaptive QML for concept drift:** Online quantum-assisted clustering and incremental variational updates can help IDS systems adapt to evolving attack patterns and concept drift in IoT traffic. This will ensure sustained accuracy in dynamic environments [2,4].

G. **Near-term hardware experiments with error mitigation:** Experiments on available QPUs using error-mitigation strategies such as zero-noise extrapolation and readout error correction are essential to validate simulation claims under realistic noise conditions [5,4,13].

VIII. CONCLUSION

Quantum-enhanced machine learning offers promising avenues to improve intrusion detection in IoT and fog environments by enriching feature representations, reducing parameter counts, and enabling novel clustering and kernel methods that exploit Hilbert-space geometry [1, 2, 5, 4, 10]. The surveyed studies demonstrate potential gains in detection metrics and clustering quality, particularly in high dimensional regimes, but also underscore practical constraints: limited qubits, noisy hardware, encoding overhead, and a lack of standardized benchmarks and real-world deployments [1, 2, 3, 4, 5, 11]. Addressing these challenges requires coordinated work on lightweight encodings, noise-aware hybrid designs, federated quantum-assisted architectures, and reproducible evaluation frameworks. With targeted research along these lines, quantum-enhanced IDS can mature from promising prototypes to practical components of secure IoT and fog computing infrastructures [3, 5, 11].

 This work is licensed under a Creative Commons Attribution 4.0 International License

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 7, July 2026
DOI: 10.17148/IJARCCE.2026.15723

REFERENCES

[1]. M. Amara, M. Amara, S. Mnasri, and T. Val, "QCNN-ID: A quantum-classical hybrid model for IoT intrusion detection," Procedia Computer Science, vol. 270, pp. 2196–2204, 2025.
[2]. K. Srinivasan, A. Trivedi, and S. Yang, "Quantum-Enhanced Clustering for Intrusion detection system in IoT devices," in Proc. IEEE 22nd Int. Conf. on Mobile Ad-Hoc and Smart Systems (MASS), 2025, doi:10.1109/MASS66014.2025.00099.
[3]. M. Tawfik, "Optimized intrusion detection in IoT and fog computing using ensemble learning and advanced feature selection," PLoS ONE, vol. 19, no. 8, e0304082, Aug. 2024.
[4]. M. Kalinin and V. Krundyshev, "Security intrusion detection using quantum machine learning techniques," Journal of Computer Virology and Hacking Techniques, vol. 19, pp. 125–136, 2022, doi:10.1007/s11416-022-00435-0.
[5]. I. Bharathi, V. Sonai, and S. Sridevi, "Quantum-driven enhanced machine learning algorithm for intrusion detection in Internet of things environment," EPJ Quantum Technology, 2026, doi:10.1140/epjqt/s40507-026-00463-5.
[6]. S. Lloyd, M. Mohseni, and P. Rebentrost, "Quantum algorithms for supervised and unsupervised machine learning," arXiv:1307.0411, 2013.
[7]. S. Lloyd, M. Schuld, A. Ijaz, J. Izaac, and N. Killoran, "q-means: A quantum algorithm for unsupervised machine learning," arXiv:2006.06886, 2020.
[8]. V. Havlicek, A. D. C'orcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gam betta, "Supervised learning with quantum-enhanced feature spaces," Nature, vol. 567, pp. 209–212, 2019.
[9]. M. Schuld and N. Killoran, "Quantum machine learning in feature Hilbert spaces," Physical Review Letters, vol. 122, p. 040504, 2019.
[10]. J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," Nature, vol. 549, pp. 195–202, 2017.
[11]. J. Preskill, "Quantum computing in the NISQ era and beyond," arXiv:1801.00862, 2018.
[12]. A. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, "Variational quantum algorithms," Nature Reviews Physics, vol. 3, pp. 625–644, 2021.
[13]. K. Temme, S. Bravyi, and J. M. Gambetta, "Error mitigation for short-depth quantum circuits," Physical Review Letters, vol. 119, p. 180509, 2017.
[14]. P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, et al., "Advances and open problems in federated learning," arXiv:1912.04977, 2019.
[15]. R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi, "A survey of methods for explaining black box models," ACM Computing Surveys, vol. 51, no. 5, article 93, 2018.

 This work is licensed under a Creative Commons Attribution 4.0 International License