| Volume 02 | Issue 01 | June 2026 | ISSN: 3139-1737 |

International Journal of Advances and Applications in Convergent Studies & Transdisciplinary (IJ-AACST)

Bridging Discipline Bridging Knowledge

Website: www.aacstij.org | Email: journaleditor@aacstij.org

Quantum Machine Learning in Cybersecurity: Recent Advances, Challenges, and Perspectives

Neha Samundre

Dept. of Computer Science,
Dr. S. C. Gulhane Prerna College of Commerce, Science and Arts
nehasamundre@yahoo.in

ORCID: https://orcid.org/0009-0002-7726-8649

ABSTRACT

Quantum machine learning (QML) has emerged as a promising intersection of quantum computing and artificial intelligence, with particular relevance to cybersecurity because modern defense systems must process large, noisy, high-dimensional, and adversarial data streams in near real time. Recent literature shows growing interest in applying QML to intrusion detection, anomaly detection, malware and botnet classification, encrypted-traffic analytics, and privacy-preserving distributed defense. At the same time, the field remains early-stage: many reported gains come from hybrid or quantum-inspired methods, small-scale experiments, or simulator-based studies rather than fault-tolerant quantum hardware. This paper synthesizes recent advances from 2022–2026 and argues that the most credible near-term role of QML in cybersecurity is not wholesale replacement of classical machine learning, but targeted augmentation of security pipelines where quantum feature maps, kernels, autoencoders, or transfer-learning modules may improve representation quality or robustness. The paper also identifies the central bottlenecks slowing practical adoption, including hardware noise, qubit scarcity, data-encoding overhead, benchmarking weaknesses, deployment latency, and the absence of consistent evidence for end-to-end quantum advantage. Finally, it outlines a forward-looking agenda centered on hybrid architectures, rigorous evaluation, privacy-aware distributed learning, and co-design with post-quantum cryptography. The overall conclusion is cautiously optimistic: QML is strategically important for cybersecurity research, but operational impact will depend on advances in hardware, reproducible benchmarking, and careful alignment with real security workflows. [1] [2] [3]

KEYWORDS: anomaly detection, botnet detection quantum machine learning, cybersecurity, federated learning, intrusion detection, post-quantum cryptography, quantum transfer learning.

1. INTRODUCTION

Cybersecurity systems now operate in environments defined by massive telemetry volumes, encrypted traffic, adversarial behavior, concept drift, and increasingly automated attacks. Traditional rule-based and classical machine-learning defenses remain important, but they often struggle when threat patterns are subtle, rare, high-dimensional, or intentionally evasive. QML has therefore attracted interest because quantum representations and hybrid quantum-classical pipelines may enable richer feature mappings, faster similarity estimation, and improved handling of complex structure in security data. Recent surveys position QML as relevant to intrusion detection, malware analysis, botnet detection, encrypted-traffic analytics, cloud security, and security optimization. [1]

A second reason this topic matters is the broader “quantum transition” now underway in cybersecurity. Quantum computing threatens much of today’s public-key cryptography, and NIST’s Post-Quantum Cryptography project has already released core PQC standards and advised organizations to begin migration. However, post-quantum cryptography and QML address different layers of defense: PQC protects confidentiality and authentication against future quantum attackers, while QML is being explored as a possible enhancement to live detection, analytics, and response. In other words, PQC secures cryptographic foundations, whereas QML may strengthen operational cyber defense. [8] [1]

This paper reviews recent advances, challenges, and future directions in QML for cybersecurity. Rather than assuming that any reported accuracy gain proves "quantum advantage," it adopts a critical perspective: Which cybersecurity tasks are actually benefiting? Under what assumptions? And what technical or organizational barriers must be overcome before QML becomes operationally meaningful? Recent work on theoretical evaluation frameworks argues that these questions are essential, especially because many near-term QML results rely on simulators or heuristic variational models whose long-term scaling remains uncertain. [2]

2. Foundations of Quantum Machine Learning for Cybersecurity

QML combines quantum computation with machine learning, usually in hybrid architectures. In practical systems, classical processors often handle preprocessing, optimization, batching, and postprocessing, while the quantum component implements a feature map, kernel estimation, variational circuit, quantum autoencoder, or other trainable quantum subroutine. Cybersecurity is a natural application area because many defensive tasks involve distinguishing benign from malicious behavior in noisy, high-dimensional data such as flows, logs, endpoint events, and IoT telemetry. Recent surveys emphasize that QML methods most commonly used in cybersecurity include quantum support vector machines (QSVMs), quantum neural networks (QNNs), variational quantum circuits (VQCs), quantum generative models, and quantum-enhanced anomaly-detection pipelines. [1]

From a theoretical standpoint, the strongest long-term case for QML in cybersecurity comes not from today's noisy hardware, but from fault-tolerant quantum algorithms with provable performance properties. Bellante et al. argue that meaningful evaluation should compare quantum and classical models not only on accuracy, but also on runtime, approximation error, and data-access assumptions. Their framework is important because it highlights a recurring weakness in applied QML papers: many claim promise on cybersecurity tasks, but relatively few establish the precise conditions under which a genuine computational advantage would emerge. [2]

At present, then, QML in cybersecurity should be understood in three layers. First are quantum-inspired approaches that borrow ideas from quantum optimization without requiring quantum hardware. Second are hybrid NISQ-era methods, where small variational circuits or quantum kernels are tested on simulators or noisy devices. Third are fault-tolerant future algorithms, which may eventually offer provable asymptotic gains on selected learning tasks. The recent literature spans all three layers, but the most practical cybersecurity deployments today remain hybrid or quantum-inspired rather than fully quantum-native. [2] [5]

3. Recent Advances

3.1 Intrusion detection and anomaly detection

One of the most active directions is intrusion and anomaly detection. A 2024 *Scientific Reports* paper proposed a Quantum Intrusion Detection System using Outlier Analysis (QIDS-OA) for DDoS detection. The model used angle embedding, parameterized quantum circuits, entangling layers, and entropy-based outlier analysis. On a generated dataset of 100,000 samples with 87 features, it reported 99.897% detection accuracy, outperforming benchmark AMM-CNN and ANN baselines in that experimental setup. The paper is notable because it illustrates how QML is being used not merely for classification, but for security-oriented state discrimination and anomaly scoring. [3]

A second important advance is the use of quantum autoencoders for anomaly detection. Hdaib, Rajasegarar, and Pan proposed three hybrid frameworks that combine a quantum autoencoder with one-class SVM, random forest, and k-nearest neighbors. Evaluated on KDD99, IoT-23, and CIC IoT 2023 datasets, the best-performing configuration—QAE plus quantum kNN—achieved 97% accuracy and 98% F1-score on CIC IoT 2023. This study is significant because it moves beyond simple QSVM-style comparisons and explores quantum-assisted representation learning for dimensionality reduction before downstream anomaly detection. [4]

Cross-domain generalization has also improved. A 2026 *Scientific Reports* study introduced a hybrid quantum transfer learning framework for cybersecurity threat detection and categorization. By combining classical pre-trained models with variational quantum circuits and standardized preprocessing across heterogeneous datasets, the framework achieved 83% accuracy on UNSW-NB15, 91% on CICIDS2017/CSE-CIC-IDS2018, and 86% on

TON_IoT. This matters because practical security systems rarely operate on a single static dataset; they must transfer knowledge across environments, infrastructures, and attack families. The paper suggests QML may become useful not just for isolated benchmarks, but for more realistic cross-domain detection settings. [6]

3.2 Distributed and privacy-aware cyber defense

Another major trend is the fusion of QML ideas with federated learning. A 2024 *Scientific Reports* study proposed a hybrid quantum-enhanced federated learning model for cyber-attack detection that integrates a spatio-temporal attention network, hierarchical model aggregation, and a quantum-inspired federated averaging technique. On the UNSW-NB15 dataset, the method reported 98.31% test accuracy, 98.15% precision, 98.50% recall, and 98.32% F1-score. The paper is especially relevant because it frames quantum-enhanced learning as part of decentralized, privacy-preserving cyber defense rather than as a standalone classifier. [5]

This direction aligns well with real-world security constraints. Many organizations cannot centrally pool sensitive network traffic, endpoint traces, or sector-specific attack data. Federated and privacy-aware cyber analytics are therefore becoming strategic priorities. Quantum-inspired aggregation and optimization may, in principle, help improve convergence or coordination in such decentralized settings, especially in IoT and large-scale distributed systems where privacy, bandwidth, and heterogeneity are central concerns. Even if these methods remain quantum-inspired rather than hardware-quantum, they show how "quantum" thinking is already influencing practical cyber-defense architectures. [5]

3.3 Malware, botnets, and specialized threat analytics

Outside intrusion detection, QML is also being tested on malware and botnet analysis. Suryotrisongko and Musashi evaluated a hybrid quantum-classical deep learning model for DGA-based botnet detection using PennyLane embeddings and layered quantum circuits. The best configuration reached 94.7% accuracy on a small-sample experiment and 93.9% on a larger setting, with one hybrid setup outperforming its classical deep-learning counterpart in a limited case. However, the broader result was mixed: overall performance was often still inferior to classical baselines. This is a valuable finding because it shows the field is not uniformly progressing toward clear superiority; some QML gains are conditional and architecture-specific. [7]

Recent surveys further indicate that cybersecurity QML applications now extend to malware and botnet detection, adversarial threat analytics, cloud security, encrypted-traffic analysis, and generative simulation. The significance of this expansion is conceptual as much as empirical: the field is moving from "Can QML classify attacks?" to "Which parts of the cybersecurity pipeline—feature extraction, clustering, anomaly scoring, transfer, privacy-preserving learning, or simulation—benefit most from quantum methods?" That reframing is a sign of maturation. [1]

3.4 Toward a research taxonomy and evaluation discipline

A further recent advance is the emergence of a structured **taxonomy** for QML in cybersecurity. The 2025 survey by Sai et al. organizes the field across supervised, unsupervised, and generative paradigms, linking algorithms such as QSVMs, QNNs, VQCs, and QGANs to intrusion detection, anomaly detection, malware analysis, encrypted traffic analytics, and cloud-security use cases. This type of taxonomy is important because cybersecurity research often suffers from fragmented terminology and incomparable benchmarks; a common map of methods and tasks can improve reproducibility and research focus. [1]

Relatedly, Bellante et al. make a methodological contribution by proposing a framework for assessing when QML might truly outperform classical ML in cybersecurity. Their case study on PCA-based intrusion detection does not merely chase higher benchmark accuracy; it asks what errors are introduced by quantum subroutines, how runtime compares with classical methods, and what hardware/software advancements are required for meaningful benefit. This is a crucial step toward more rigorous science in the area. [2]

4. Comparative Snapshot of Recent Studies

Study	Task	Method	Dataset(s)	Headline Result	Source
Kim & Madhavi (2024)	DDoS / intrusion detection	QIDS-OA, QNN + angle embedding + outlier analysis	Generated network dataset	99.897% accuracy	[3]
Hdaib et al. (2024)	Anomaly detection	Quantum autoencoder + QkNN / QSVM hybrids	KDD99, IoT-23, CIC IoT 2023	97% accuracy, 98% F1 on CIC IoT 2023	[4]
Subramanian & Chinnadurai (2024)	Federated cyber-attack detection	STAN + hierarchical aggregation + quantum-inspired averaging	UNSW-NB15	98.31% test accuracy	[5]
Suryotrisongko & Musashi (2022)	Botnet DGA detection	Hybrid quantum-classical deep learning	Botnet DGA dataset	Up to 94.7% accuracy in one setting	[7]
Alsubai et al. (2026)	Cross-domain threat detection	Hybrid quantum transfer learning	UNSW-NB15, CICIDS2017, CSE-CIC-IDS2018, TON_IoT	83–91% accuracy across domains	[6]
Sai et al. (2025)	Survey/taxonomy	QSVM, QNN, VQC, QGAN taxonomy	Multi-domain review	Structured map of cybersecurity QML use cases	[1]
Bellante et al. (2025)	Evaluation framework	Fault-tolerant QML assessment	PCA-based intrusion detection case study	Identifies conditions needed for quantum advantage	[2]

5. Challenges and Limitations

5.1 Hardware immaturity and simulation bias

The strongest practical obstacle is that many reported QML cybersecurity results are obtained on simulators or under idealized conditions. The 2026 transfer-learning paper explicitly states that experiments were performed on a PennyLane simulator and that real deployment would face noise, decoherence, gate errors, and qubit-connectivity limits. Likewise, the anomaly-detection and hybrid botnet papers note that current NISQ hardware still hinders realization of the theoretical speedups often associated with quantum models. This means that much of the present literature demonstrates **algorithmic feasibility**, not operational readiness. [6] [4] [7]

5.2 Data encoding and scalability

Cybersecurity data are often high-dimensional, heterogeneous, streaming, and sparse. Encoding such data into quantum states is expensive and can erase the practical benefit promised by later quantum processing. The 2024 quantum autoencoder paper notes that angle encoding becomes difficult in high-dimensional settings because circuit complexity can grow rapidly, while amplitude encoding is itself intricate and costly to implement. Bellante et al. similarly emphasize that any claim of quantum advantage must account for data access and preparation assumptions, not just the complexity of the quantum subroutine in isolation. [4] [2]

5.3 Benchmarking weakness and limited realism

A recurring issue is the gap between benchmark success and operational realism. The QIDS-OA study reports very high accuracy, but it relies on simulated traffic rather than diverse real-world deployments. The botnet DGA study finds only selective advantage relative to classical deep learning. The Bellante framework warns that many current QML evaluations lack theoretical guarantees, strong baseline comparisons, or realistic scaling analyses. In cybersecurity, such gaps matter because a detector that performs well on a curated benchmark may fail under distribution shift, adversarial manipulation, or heavy production traffic. [3] [7] [2]

5.4 Integration into security operations

Even an accurate quantum-enhanced classifier may still be impractical if it cannot meet real-time latency and integration demands. The 2026 transfer-learning study acknowledges that operational intrusion detection must handle high-speed traffic, strict latency budgets, limited compute resources, and seamless integration with existing security infrastructure. This point is often underemphasized in academic QML work: cybersecurity tools are not judged only by accuracy, but by deployment friction, alert quality, interpretability, resilience, and maintenance cost. [6]

5.5 No clear, general quantum advantage yet

The most important conceptual limitation is that **general quantum advantage for cybersecurity has not yet been established**. Bellante et al. explicitly frame their contribution around identifying the conditions under which such advantage might emerge in the future, implying that it cannot currently be assumed. The 2022 botnet study also found only partial superiority over classical models. Taken together, these works suggest that the field is promising but pre-paradigmatic: claims must remain conditional, task-specific, and benchmark-sensitive. [2] [7]

6. Perspectives and Future Directions

The most realistic near-term perspective is the rise of **hybrid quantum-classical cybersecurity pipelines**. Rather than replacing mature SIEM, IDS, or SOC analytics stacks, QML is more likely to be inserted into selected stages such as feature transformation, anomaly scoring, clustering, transfer learning, or privacy-aware aggregation. This view is consistent with current empirical work, where the best results often come from hybrid autoencoders, hybrid transfer-learning systems, or quantum-inspired federated optimizers rather than purely quantum end-to-end architectures. [4] [6] [5]

A second major direction is **privacy-preserving and distributed defense**. Federated learning is already important in cybersecurity because many stakeholders cannot centralize sensitive data. Quantum-enhanced or quantum-inspired federated methods may help with optimization, convergence, and resilience across

heterogeneous edge, IoT, and enterprise environments. As critical infrastructure becomes more decentralized, this may become one of the most practical arenas for quantum-enhanced cyber analytics. [5]

Third, the field needs **better evaluation protocols**. Future studies should report not only accuracy, precision, recall, and F1, but also latency, resource cost, robustness to adversarial drift, cross-dataset transfer, false-positive burden, and integration overhead. Evaluation should use stronger classical baselines and more realistic datasets. The emerging methodological work by Bellante et al. and the taxonomy work by Sai et al. provide a foundation for that transition from exploratory experimentation to disciplined comparative science. [2] [1]

Fourth, QML should be developed alongside—not instead of—**post-quantum cryptography**. NIST has made clear that organizations should begin migration to PQC now, with ML-KEM, ML-DSA, and SLH-DSA forming the foundation of current standards. In future quantum-era cyber defense, PQC will secure cryptographic trust, while QML may enhance detection, prediction, and response. A mature cybersecurity strategy will likely require both. [8]

Finally, a longer-term perspective is the move from heuristic NISQ experimentation toward **fault-tolerant, provable QML for security analytics**. If scalable quantum hardware and efficient data-loading schemes emerge, then tasks such as high-dimensional anomaly detection, clustering, graph-based attack-path analysis, encrypted-traffic analytics, and large-scale threat hunting may become more favorable to quantum methods. But until then, the most responsible stance is not hype, but carefully validated progress. [2] [1]

7. Conclusion

Quantum machine learning has become one of the most intellectually compelling frontiers in cybersecurity research. Recent work shows encouraging progress in intrusion detection, anomaly detection, federated cyber defense, and cross-domain transfer learning. Hybrid quantum autoencoders, transfer-learning architectures, and quantum-inspired federated optimizers demonstrate that the field is moving beyond toy examples and into more security-relevant scenarios. [4] [5] [6]

Yet the evidence still supports a cautious conclusion. Most gains remain task-specific, benchmark-dependent, and often simulator-based. Data encoding, hardware noise, scalability, latency, and reproducibility continue to limit deployment. Most importantly, there is not yet a universal demonstration that QML outperforms strong classical cybersecurity analytics in operational settings. The field's importance therefore lies not in immediate replacement of classical ML, but in opening a new design space for future cyber defense—especially as quantum hardware matures and cybersecurity itself becomes more distributed, privacy-sensitive, and data-intensive. [2] [7]

In summary, the outlook for QML in cybersecurity is best described as **strategically significant, technically promising, and operationally premature**. The coming years should focus on hybrid deployment, rigorous benchmarking, and co-evolution with post-quantum cryptography. If those foundations are laid well, QML could become an important component of future cyber defense architectures. [1] [8]

8. References

1. Sai, S., Goyal, I., Sharma, S., Manuri, S. H., Chamola, V., & Buyya, R. **Quantum Machine Learning for Cybersecurity: A Taxonomy and Future Directions.** arXiv, 2025. https://arxiv.org/abs/2512.15286

2. Bellante, A., Fioravanti, T., Carminati, M., Zanero, S., et al. **Evaluating the Potential of Quantum Machine Learning in Cybersecurity: A Case-Study on PCA-based Intrusion Detection Systems.** *Computers & Security*, 2025. https://arxiv.org/abs/2502.11173

3. Kim, T. H., & Madhavi, S. **Quantum intrusion detection system using outlier analysis.** *Scientific Reports*, 2024. https://www.nature.com/articles/s41598-024-78389-0

4. Hdaib, M., Rajasegarar, S., & Pan, L. **Quantum deep learning-based anomaly detection for enhanced network security.** *Quantum Machine Intelligence*, 2024. https://link.springer.com/article/10.1007/s42484-024-00163-2

5. Subramanian, G., & Chinnadurai, M. **Hybrid quantum enhanced federated learning for cyber attack detection.** *Scientific Reports*, 2024. https://www.nature.com/articles/s41598-024-83682-z

6. Alsubai, S., Ayari, M., Kryvinska, N., Almadhor, A., Baili, J., et al. **Quantum transfer learning for cross-domain cybersecurity threat detection and categorization.** *Scientific Reports*, 2026. https://www.nature.com/articles/s41598-026-40634-z

7. Suryotrisongko, H., & Musashi, Y. **Evaluating hybrid quantum-classical deep learning for cybersecurity botnet DGA detection.** *Procedia Computer Science*, 2022. https://www.sciencedirect.com/science/article/pii/S1877050921023590

8. NIST. **Post-Quantum Cryptography Project.** National Institute of Standards and Technology. https://csrc.nist.gov/Projects/post-quantum-cryptography

9. Moll M., Leonhard Kunczik **"Comparing quantum hybrid reinforcement learning to classical methods."** Human-Intelligent Systems Integration, 3 (1) (2021), pp. 15-23. https://www.sciencedirect.com/science/article/pii/S1877050921023590

10. Schuld M., Sinayskiy I., Petruccione F**. "An introduction to quantum machine learning."** Contemporary Physics, 56 (2) (2015), pp. 172-185 https://www.tandfonline.com/doi/abs/10.1080/00107514.2014.964942

11. Buffoni L., Filipo Caruso **"New trends in quantum machine learning (a)"** EPL (Europhysics Letters), 132 (6) (2021), p. 60004. https://iopscience.iop.org/article/10.1209/0295-5075/132/60004/meta