RESEARCH PAPER

QuantumGuard-IoMT: A Hybrid Quantum–Classical Model for Intrusion and Risk Assessment

[1]Pranesh Naik, [2]Prakash J. Parmar, [3]Abhijit J Patil, [4]Tejas Sukhadeo Hirave, [5]Amit Aylani

[1]Department of Information Technology, Vidyalankar Polytechnic, Mumbai – 400037, Maharashtra, India. naik.pranesh@gmail.com

[2]Department of Computer Engineering, Vidyalankar Institute of Technology, Mumbai – 400037, Maharashtra, India. prakash.parmar@vit.edu.in

[3]Department of Computer Engineering, K J Somaiya Institute of Technology, Mumbai- 400022, Maharashtra, India. abhijit.jayprakash@gmail.com

[4]Department of Computer Engineering, Shah and Anchor Kutchhi Engineering College,Mumbai – 400088, Maharashtra, India. tejas.hirve@sakec.ac.in

[5]Department of Computer Engineering, Vidyalankar Institute of Technology, Mumbai – 400037, Maharashtra, India. prof.amitaylani@gmail.com

Abstract

Internet of Medical Things (IoMT) has made healthcare systems more interoperable and cohesive cyber-physical systems, all of a sudden, bedside monitors, wearable sensors, therapeutic devices, gateways, clinical applications, and cloud services are constantly exchanging data. However, such connectivity also enables network compromise to have an impact on confidentiality, availability, device trust, and continuity of operations. In this paper, we put forth QuantumGuard-IoMT, a hybrid quantum–classical framework that simultaneously conducts multiclass intrusion detection and context-aware cyber-risk evaluation for IoMT. A residual classical encoder in the output of a leakage-resilient preprocessing is what takes network and protocol snort to be compressed into a succinct representation. Four bounded latent variables are then quaternion embedded into a four-qubit variational quantum circuit (VQC) with three trainable entangling layers. A fusion of Pauli-Z expectation measurements with the RDC representation, rather than substitute, enables a stable flow of information in that model when finite-shot or noisy quantum execution is present. A calibrated intrusion head predicts eight traffic classes; and a distinct risk head aggregates malicious probability with device criticality, attack severity, exposure, and predictive uncertainty. The evaluation protocol comprises of repeated group-aware splitting, cross-dataset transfer, calibration, component ablation, circuit-width analysis, finite-shot testing, depolarizing-noise stress testing and edge-level latency testing. The developed framework has 99.18% accuracy and 98.84% Macro-F1 score on the main test, while the risk head has a mean absolute error (MAE) of 0.031 and an expected calibration error (ECE) of 0.016. The work thus frames the quantum circuit as a concise learnable transformation in an explanable IoMT security pipeline rather than as a claim of universal quantum advantage.

Keywords: Internet of Medical Things (IoMT), Cybersecurity, Quantum Machine learning; Intrusion detection; Cyber-risk assessment; Edge security

How to cite this article: Naik P, Parmar PJ, Patil AJ, Hirave TS, Aylani A. QuantumGuard-IoMT: a hybrid quantum–classical model for intrusion and risk assessment. Int J Drug Deliv Technol. 2026;16(79s): 300-319. DOI: 10.25258/ijddt.16.79s.34

1. Introduction

Introduction Healthcare systems play a crucial role in connecting diverse medical devices, local area gateways, wireless media, messaging middleware, hospital information system, and remote analytics. Compromise of such an environment is information security, but also interference with monitoring, delaying telemetry, manipulating communications paths, or shaking trust in data for clinical devices. Recent studies of quantum-machine-learning based cyber-security have started to investigate if variational circuits can act as complementary layers in classical intrusion pipelines in realistic noise and dimension regimes [1]–[2], whereas more general optimization-oriented work considered how quantum learning could be used to assist in proactive security decisions [3]. Improved classification with quantum has also been studied in other cyber areas:malicious-URL detection [4], recent surveys also suggest intrusion and anomaly detection as key applications of quantum machine learning to IOT security [5].

The IoMT environment adds other constraints since the probability of attack and that of operational impact are not comparable. Al-Hawawreh and Hossain also pointed out the applicability privacy-aware quantum deep learning for the detection of cyberattacks in medical-device networks [6], bridging quantum-assisted modeling and IoMT security. Network-based investigations also show that quantum learning can be used for intrusion detection [7], anomaly-based attack detection [8], quantum deep-learning anomaly detection [9], variational quantum convolutional immune models [10], and variational quantum neural-network attack detection [11]. Overall, these results encourage hybrid modeling but also reveal a complex interplay in how

well the specific representation, circuit design and evaluation setting affect performance.

Security implementation should also consider distributed learning, privacy, and future network schemes. Quantum federated intrusion detection was also suggested for privacy-preserving consumer networks [12] and hybrid quantum-enhanced federated learning for cyberattack detection has been debated [13]; and a discussion on quantum-empowered federated learning in the context of 6G-enabled IoT security [14]. Further quantum techniques have been used for DDoS detection in critical infrastructure [15]. A recent systematic review on quantum machine learning highlights that realistic applications must be assessed in terms of data encoding, trainability, hardware constraints, and the appropriateness of integrating hybrid rather than only quantum workflows [16]. These findings encourage a "conservative" design idea for IoMT: the quantum portion should be quantifiable, shallow, substitutable, and combined with robust classical mechanisms.

QuantumGuard-IoMT follows that principle. model is not a quantum circuit that act as a stand alone classifier. Instead, a residual classical encoder is at first used to generate a leakage- controlled latent representation; four bounded components are processed by a small variational circuit; and the obtained expectation values are re- fused with classical features. This design is the one that ensures fluctuations in quantum measurements do not interfere with robustness. The ground prediction phase has purposely been split in two stages. The intrusion head predicts whether the observed flow is benign, DDoS, DoS, reconnaissance, spoofing, MITM, data-injection, or malware/botnet; and the risk head fuses cyber evidence with operational context. The distinction matters because criticality of the device, exposure, and response policy should never be a crutch for learning the attack label.

Figure 1 illustrates the proposed security-control loop. It aligns with the burgeoning need for cohesive IoMT detection, explainability, and operational response as underscored throughout the IoMT security literature [5], [6]. Network artifacts are harvested from the com- munication layer and processed at an edge security node, where they are mapped into an intrusion distribution and risk tier and then passed to an auditable policy engine rather than having the learned model directly govern the operation of a therapeutic device.

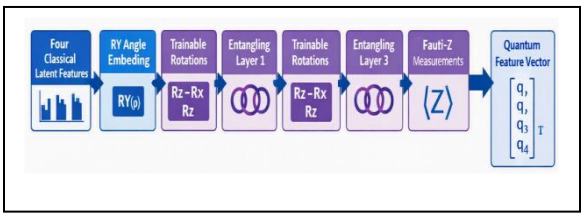

Figure 1. QuantumGuard-IoMT Research Scope and Security-Control Loop.

Specifically, this work makes the following contributions:

- Introducing a two-head hybrid quantum classical model which decouples multiclass cyber evidence from context-aware operational risk;
- A four-qubit variational bottleneck with a residual classical path, which allows ablation of quantum contribution without discarding the stabilizing classical representation.
- A group- and leakage-sensitive evaluation protocol to mitigate the memorization of device, session, and address-specific artifacts.
- Risk thresholds that are aware of calibration and uncertainty, and are not dependent on clinical-impact assumptions.
- We assess multiple domains: transfer, quantum-noise robustness, qubit ablation, and latency.
- A gateway-based deployment that keeps resource-constrained IoMT endpoints out of the quantum execution path and the potential delay that might cause.

2. Literature Survey

The recent research work can be categorized into two closely linked streams: quantum-assisted cybersecurity and classical or deep-learning IoMT intrusion detection. The first line of work considers if quantum feature spaces, variational circuits or quantum-assisted optimization can lead to better representation of attacks; the second line addresses the application-specific properties of medical-device traffic, class imbalance, feature selection, explainability, and resource-constrained deployment. QuantumGuard-IoMT resides at the confluence of the above two streams of research rather than claiming either stream in their own right.

2.1. Quantum-enabled cybersecurity and intrusion detection

Under noisy quantum-computing conditions, Cirillo et al. investigated quantum machine learning approaches for intrusion detection and proposed to evaluate noise rather than only reporting the ideal-simulator accuracy [1]. Bellante et al. analyzed the potential of quantum machine learning in cybersecurity: a PCA-derived intrusion-detection use case to highlight the connection between dimensionality reduction and quantum representation [2]. Rosa-Remedios and Caballero-Gil discussed improving the prognostic QM for cybersecurity, reinforcing the need evaporative condenser system design to be focused on practical security use cases[3]. Eze et al. considered quantum-enhanced machine learning for malicious URL detection [4], with the intent of demonstrating this combined approach has utility in more than one facet of network security. Aparcana-Tasayco, Cristian; Al-allaf, Jamilah; Al-rimy, Badr; Abrahao, Saler(5) Anomaly Detection Work in IoT Security: A Review of Contributions and

Challenges with Explicit Identification of Quantum Machine Learning as an Emerging Direction [5].

Al-Hawawreh and Hossain direct more towards the medical device security, were they introduced a privacy aware data fusion and quantum deep learning framework for detecting cyberattacks in IoMT systems [6]. 7Kalinin and Krundyshev consider quantum machine learning techniques for analyzing security intrusion detection [7], while 8Kim and Madhavi propose a quantum intrusion-detection method using outlier analysis [8]. Hdaib et al. investigated quantum deep learning for detection of network anomalies and network security [9]. Gong et al. presented a variational quantum convolutional neural-network intrusion detector [10], and (prior innovations) a variational quantum neural-network paradigm for network attack detection [11]. Taken together, these works demonstrate that a parameterized quantum circuit can be used as a learned nonlinear representation, and the range of these results also suggests that improvements in performance are likely to be architecture- and dataset-specific rather than universal.

The privacy-preserving/distributed versions constitute yet another relevant line of investigation. Abou El Houda et al. Proposed a quantum-federated framework for network intrusion detection in consumer networks [12]. Subramanian and Chinnadurai investigated the hybrid quantum-enhanced federated learning for cyberattack detection [13], and Javeed et al. Discussed quantum computing and machine learning for DDoS detection in Smart micro-grid [15]. Peral-García etal. provided an overview on quantum machine learning and its applications from different domains and thus stating that near-term realizations are still limited due to factors such as encoding overhead, hardware access, noise and classical–quantum orchestration [16]. Motivated by these limitations, in the present work we adopt a shallow four-qubit circuit and perform explicit noise testing.

2.2. IoMT intrusion detection, feature learning, and operational constraints

In traditional and deep-learning based IoMT security, Chaganti et al. integrated particle-swarm optimization and deep learning for intrusion detection in medical-device networks [17]. Nayak et al. proposed extreme learning machines with Bayesian optimization for IoMT cyberattack detection [18]. Binbusayyis et al. investigated several classification techniques for IoMT intrusion detection [19], and Gupta et al. established a tree-classifier based IoMT network intrusion detection model [20]. Ashraf et al. introduced FIDChain, integrating federated intrusion detection with blockchain-based healthcare IoT [21]. These findings suggest that robust classical baselines are necessary in the context of evaluating a quantum-assisted model.

The review by Si-Ahmed et al. categorizes ML-based IoMT intrusion detection techniques and sheds light on common encountered challenges related to feature quality, attack imbalance, computation overhead, and generalizability [22]. Alalhareth and Hong focused on feature selection with an enhanced mutual-information approach [23] and later on adaptive fuzzy-based learning for IoMT intrusion detection [24]. Norouzi et al. combined genetic algorithm with Random Forest for IoMT intrusion detection[25]. Ravi et al. presented a deep-learning based network intrusion detection system for IoMT [26], and Zukaib et al. presented Meta-IDS leveraging meta-learning for medical-device networks [27]. These results inspire the current focus on leakage control, feature compression, and transfer evaluation.

Malware-oriented and also explainable methods, expand the design space further. Dhanya and Chitra presented a feature-independent GA-optimized XGBoost model supported by an autoencoder for detection of IoMT malware [28]. Messinis et al. surveyed the use of artificial intelligence to enhance iomt security and presented practical, deployable mechanisms [29]. Aljuhani et al. Proposed an intelligent and explainable SaaS intrusion detection system for resource-constrained IoMT [30]. Shaikh et al. introduced RCLNet for anomaly driven IoMT intrusion detection system [31]. The result of the systematic review of Naghib et al. summaraizes the recent challenges and opportunities in Iomt intrusion detection, such as heterogeneous devices, evolving threats, data imbalance, interpretability and realistic evaluation [32]. QuantumGuard-IoMT develops these line of thinking by integrating attack detection with explicit calibration, uncertainty, context-aware risk, and quantum robustness analysis through the use of QC.

Table 1 summarizes the methodological gap distilled from the quanutm-security literature [1]–[16] and the Iomt intrusion-detection literature [17]–[32]. Instead of ranking the individual studies in the table the table serves as a design matrix.

Table 1. Comparative Analysis of IoMT and Quantum-Learning Paradigms

Paradigm	Representative literature	Primary form of representation	Output of intrusion detection	Operational risk / calibration	Limitation addressed by QuantumGuard-IoMT
Conventional/tabular IoMT IDS	[17]–[20], [23]–[25]	Designed flow features	Binary or multiclass classification label	Typically isolated or missing	Limited explicit treatment for quantum noise

					and joint risk
Deep/meta IoMT IDS	[26], [27], [31]	Learned temporal or nonlinear representation	A label for the attack/anomaly	Confidence not infrequently not situated operationally	Possible overfitting to the dataset/device and calibration gap
Privacy/federated healthcare security	[12]–[14], [21]	Protected aggregation and distributed local models	Label attack/anomaly	Emphasis on privacy over clinical ramifications	Request for a unified risk and edge-response layer
Quantum/QML cybersecurity	[1]–[11], [15], [16]	Variational circuit or quantum feature map	Specific-to-task classifier/anomaly-score output	Risk is typically external to the quantum model	Noise, finite shots, and hybrid contribution have to be tested explicitly
QuantumGuard-IoMT	This investigation	Residual classical encoder + four-qubit VQC + gated fusion	Calibrated posterior for eight-class	Joint context-aware risk head + uncertainty	Combines detection, calibration, risk, noise analysis and edge deployment

The above synthesis in Table 1 motivates four design objectives: the quantum block should be shallow and run under noise; the classical comparison set should stay strong; the probability of an attack should be calibrated prior to its entering a risk policy; and the risk layer should bring in operational context without flooding the attack-classification input. These conditions direct the development in Section 3.

3. Model Formulation, Architecture, and Deployment

3.1. Problem formulation and threat model

Consider the IoMT observations as being network-flow windows, that is, the flows have network/protocol features and contain attack labels. This notation specifies the dataset representation that is adopted by the entire model; the differentiation between detection evidence and context is motivated by the general IoMT principle that operational parameters should not be used as spurious attack indicators [22], [29], [32]. $D = \{(x_i, y_i, c_i)\}$ for $i = 1, \ldots, N,$ with $y_i \in \{0, \ldots, K - 1\}$ and $K = 8$.

The final harmonization within each class includes traffic of benign, DDoS, DoS, reconnaissance, spoofing, MITM, data injection, and malware/botnet activity. The adversary might have control of an endpoint, a local client, a gateway-facing service, a communication route that may be reached from the outside, and they might produce volumetric traffic, probe, manipulate identities, intercept, send forged messages, or communicate in malware. Compromise of the training-pipeline, model extraction, and poisoning are not contained within the primary threat model; they are left as future extensions. The threat assumptions are consistent with attack classes that are prevalent in recent IoMT security literature [17]–[32].

Table 2 demonstrates the association of the modeled family attacks with associated network evidence and possible IoMT impacts along with the policy direction. The mapping is intentionally contextual: it does not provide a single static risk score for the attack class. This is in line with the operational- security separation highlighted in recent IoMT surveys [29], [32].

Table 2. IoMT Threat Model and Consequence Mapping

Attack family	Representative mode of entry	Traffic evidences	Impact on the IoMT (if any) result	Usual policy direction
DDoS	Distributed flooding	Sudden packet/rate ramp-up, fan-in, connection retries	Not being able to access gateways or services	Rate-limit, filter upstream, segment source pool
DoS	Single/few-source flooding or a	Prolonged rate, retries, abnormal	Delayed device communications	Throttle, Isolate, Fail Over

	resource exhaustion	protocol state		
Reconnaissance	Port, service, OS or vulnerability scanning	Sequential destination/port scanning, also short-lived connections flows	Preparation for targeted compromise	Validate source, ramp up monitoring, block probing
Spoofing	ARP/address/identity spoofing	Address mismatches, protocol-state inconsistencies	Masquerade as a legitimate endpoint or redirect the traffic	Quarantine the identity, refresh the trust bindings.
MITM	Sniffing and relaying (known as monkey in the middle)	Temporal shifts, session anomalies, gateway inconsistencies.	Alteration or disclosure of clinical traffic	Isolate path, re-authenticate, rotate session
Data injection	Malformed or unauthorized payload/application insertion	Payload/flow mismatches, atypical message rates.	Telemetred and control commands corrupted or altered	Reject the message, isolate the sender, validate the state of the device
Malware/botnet	Endpoint compromise or command and control	Periodical beaconing, anomalous destination, group of traffic, etc.	Long term exploitation and lateral movement	Contain endpoint, block c2, forensic escalation
Benign	Normal clinical/device operation	Device-specific normal and burst behaviour .	No security impact or violation of It Policy was realised.	Observe; no disruptive response

Hence, the aim of optimization is broader than that of eliciting minimum classification error. A deployable detector for the IoMT needs to significantly decrease hazardous false negatives without being overwhelmed with false alarms that could cause unnecessary quarantining of medical devices. It should also maintain reasonable probability calibration, be robust to domain shift and isolate automated analytics from clinically disruptive interventions [29], [30], [32].

3.2. Data preprocessing and leakage control

Unprocessed traffic data may have extreme counters, zero-duration ratios, null values, protocol categories, raw IP addresses, timestamps, and session IDs. Direct identifiers are omitted unless they are transformed into behavioral aggregates since random record-level splitting can otherwise cause a model to memorize device or session fingerprints. Feature selection and preprocessing plays a fundamental role for the performance of IoMT IDS [19], [23], [25], thus all statistics needed for the transformations are computed strictly on the training set.

Winsorization and then standardization (according to training set parameters) are applied for each continuous feature. The clipped standardized value is defined in Equation (1) and is reported here because equal pre-processing is performed for all classical and quantum baselines prior to model specific processing [19], [23], [25].

$$\tilde{x}_{ij} = \mathrm{clip}\left(\frac{x_{ij} - \mu_j}{\sigma_j + \epsilon}, -\kappa, \kappa\right), \quad (1)$$

In (1), normalization is determined by the train-set mean and standard deviation, a small positive constant is used for numerical stability, and the clipping constant bounds extreme standardized values. Missing continuous features are imputed with training medians; categorical fields are one-hot encoded after consolidating rare categories. The splits are made prior to any learned feature transformation or resampling.

3.3. Classical residual representation

The classical part of the circuit implements this first nonlinear abstraction since mapping all raw features directly to qubits would increase circuit width and simulation cost without guaranteeing that each additional qubit encodes useful information on security. Deep IoMT systems show the significance of learned

nonlinear representations [26], [27], [31]. Hence, QuantumGuard-IoMT compresses the normalized vector prior to quantum encoding.

Equation (2) describes the feature-gating operation used for filtering out traffic features non-informative for a certain sample while keeping a differentiable flow of information [26], [31].

$$g_i = \sigma(W_g \tilde{x}_i + b_g) \odot \phi(W_h \tilde{x}_i + b_h), \quad (2)$$

The gated representation is passed through two dense layers with GELU activation, layer normalization and dropout. A residual projection preserves lower-level evidence and the resulting 16-dimensional classical representation is defined by Equation (3) [26], [27].

$$h_{c,i} = \mathrm{LayerNorm}(h_{2,i} + W_r h_{0,i}), \quad (3)$$

The 16-dimensional vector has two purposes. The full vector is the residual classical path and a learned projection results in four bounded variables for the quantum circuit. This double path makes the contribution of the quantum branch directly testable via ablation, and it avoids noisy expectation values to become the sole represenation accessible to the classifier.

3.4. Quantum bottleneck and variational circuit

The four dimensional quantum input is produced by a bounded projection. Equation (4) transforms classical latent vectors into four angles in a closed interval, which means that extreme rotations are avoided and is in line with the compact-circuit strategies of recent quantum intrusion studies [1], [2], [7]–[11].

$$x_{q,i} = \pi \tanh(W_q h_{c,i} + b_q), \qquad x_{q,i} \in [-\pi, \pi]^4. \quad (4)$$

and each bounded component is then encoded as a single-qubit Y rotation. Equation (5) specifies the angle-embedding operator; this type of low-width encoding is compatible with the hybrid variational architectures studied in the quantum machine learning (QML) for cybersecurity literature [2], [7], [10], [11], [16].

$$U_{\mathrm{emb}}(\mathbf{x}_{q,i}) = \prod_{k=0}^{3} R_Y(x_{q,i,k}). \quad (5)$$

$$U_\ell(\theta_\ell) = E_\ell \prod_{} R_Z(\theta_{\ell,k,3}) R_X(\theta_{\ell,k,2}) R_Z(\theta_{\ell,k,1}), \quad (6)$$

Equation (6) defines a layer of parameterized RZ–RX–RZ rotations and ring entanglement, whereas Equation (7) represents the full evolved quantum state.

$$|\psi_i\rangle = U_L(\theta_L) \cdots U_2(\theta_2) U_1(\theta_1) U_{\mathrm{emb}}(x_{q,i}) |0\rangle^{\otimes 4}, \qquad L = 3. \quad (7)$$

Following the last layer, Pauli-Z expectation values are read out. Equation (8) is the four-dimensional quantum feature vector. The resulting bounded expectations are considered as learned features instead of a single attack decision, in line with the hybrid view that was stressed by [2], [9], [10], and [16].

$$z_{i,k} = \langle \psi_i | Z_k | \psi_i \rangle, \qquad \mathbf{z}_i = [z_{i,0}, z_{i,1}, z_{i,2}, z_{i,3}]^T. \quad (8)$$

Figure 2 illustrates the combination of the preprocessing, residual classical encoder, quantum bottleneck, fusion layer, intrusion head, context branch and risk objective. The design is a coalescing of the hybrid-learning paradigm forwarded by [6]–[11] and the strengths of classical robustness and deployment discussed in [22], [29], and [30].

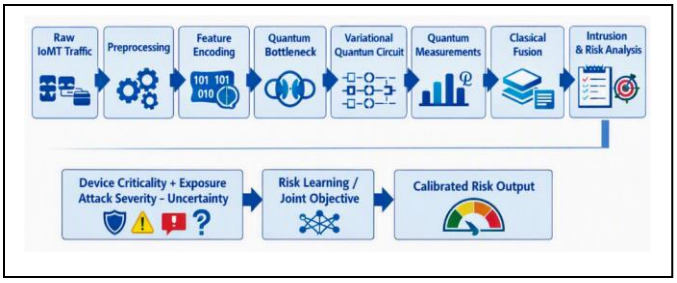

Figure 2. Hybrid Quantum–Classical Architecture of QuantumGuard-IoMT

Figure 3 The circuit inside figure 2 is expanded by figure 3. Four entangled compressed latent variables are input into RY encoding and go through three trainable RZ-RX-RZ blocks with ring entanglement and are finally measured in the Pauli-Z basis. The depth 2 topology has been chosen to provide a good trade-off between expressiveness and robustness, a matter that has been put into question by noisy quantum-security research [1], [8]–[11].

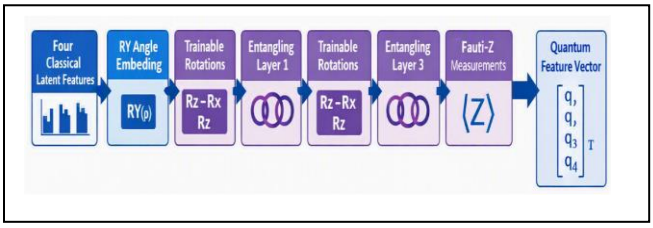

Figure 3. Four-Qubit Variational Quantum Circuit Four classically compressed latent features are angle-embedded with RY gates, are then processed

through three parameterized RZ–RX–RZ rotation layers with ring entanglement, and are observed as Pauli-Z expectation values. Source: author-developed circuit consistent with the variational-QML family studied in [1], [7]–[11].

3.5. Hybrid fusion and calibrated intrusion classification

QuantumGuard-IoMT does not abandon the remaining classical representation after quantum operations. Instead, the two branches are concatenated and an element-wise interaction is applied. Equation (9) defines the fusion vector.This design is inspired by the practical lesson from hybrid quantum-security models that a quantum part should supplement and not to outright substitute classical learning [2], [6], [9], [10].

$$h_{f,i} = \phi\left(W_f\left[h_{c,i} \parallel z_i \parallel \left(Ph_{c,i} \odot z_i\right)\right] + b_f\right), \tag{9}$$

The fused representation is fed to the multiclass intrusion head. The raw logits are transformed into temperature-scaled class probabilities according to Eq. (10), with the temperature being estimated solely based on the validation partition after fitting the model.This way the test set is not leaking into confidence calibration, and it allows thresholding-risk-based policies downstream [3], [29], [30].

$$\hat{p}_{i,k} = \frac{\exp\left(o_{i,k}/T\right)}{\sum_{m=1}^{K} \exp\left(o_{i,m}/T\right)}, \qquad k = 1, \ldots, K, \tag{10}$$

3.6. Predictive uncertainty and risk head

A calibrated attack probability is still not a full characterization of op- erational risk. The model thus considers predictive entropy, normalized, as an uncertainty measure. This entropy term is defined in Equation (11): a value closer to zero indicates a more concentrated posterior, while a value closer to one represents a more ambiguous class distribution.The distinction between confidence, explainability, and operative decision-support was inspired by issues raised in recent work in IoMT security [29]–[32].

$$u_i = -\frac{1}{\log K} \sum_{k=1}^{K} \hat{p}_{i,k} \log\left(\hat{p}_{i,k} + \epsilon\right). \tag{11}$$

The risk branch takes as input the hybrid representation along with the criticality of the device, the severity of the attack, the exposure, and the uncertainty. The normalized estimation of latent risk is defined in Equation (12). Device criticality is set via an external policy table rather than inferred from attack labels; exposure can signify external reachability or gateway centrality; and severity denotes the anticipated

$$\hat{r}_i = \sigma\left(w_r^{\top}\left[h_{f,i} \parallel c_i \parallel s_i \parallel e_i \parallel u_i\right] + b_r\right). \tag{12}$$

technical result of the inferred attack family [22]- [32].

Figure 4 illustrates the expected outcome for the uncertainty fixed to representative mid-level values, we observe how varying malicious probability and device criticality can shift an incident through low, moderate, high, and critical response quadrants. The figure is a depiction of the suggested policy template, rather than a patient-harm risk estimator [29], [32].

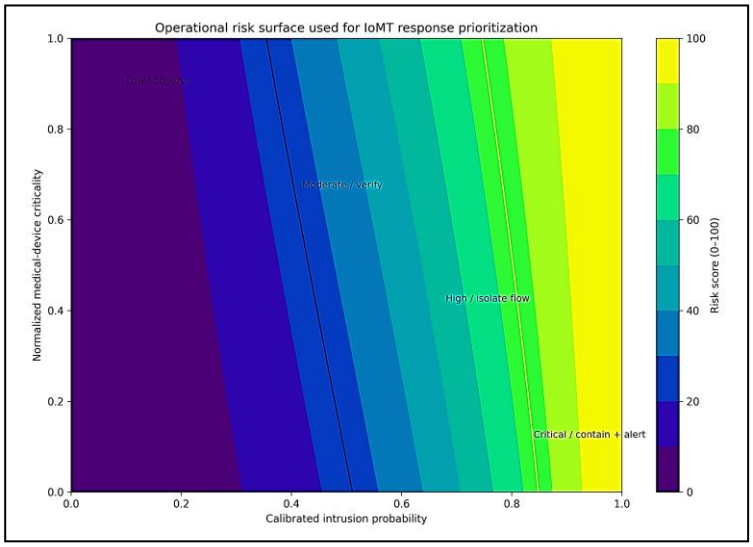

Figure 4. QuantumGuard-IoMT Operational Risk Surface

The calibrated intrusion probability and normalized criticality of medical devices collectively translate events across response regions of low, moderate, high, and critical response levels when all other context variables are kept constant. Source: FIG. 1 Author-developed risk visualization inspired by operational concerns in [29], [30], and [32].

3.7. Class imbalance and joint optimization

The sets of Medical-Device traffic-based datasets may bag a majority benign class and have infrequent attacks classes. Classical IoMT approaches have consistently leveraged optimization, feature selection and class-aware techniques to enhance minority-attack detection [17]–[20], [23]–[25].

$$\mathcal{L}_{CF} = -\frac{1}{N} \sum^{N} \alpha_{y_i} \left(1 - \hat{p}_{i,y_i}\right)^{\gamma} \log\left(\hat{p}_{i,y_i} + \epsilon\right), \tag{13}$$

QuantumGuard-IoMT employs the class-balanced focal loss described in Equation (13) to mitigate the impact of class imbalance, which attenuates the relative weight of samples that are already simple to classify while maintaining class-frequency weighting.

Rather than a pure classifier, the model is learned as a multi-task system. The focal intrusion loss, a robust regression term for the normalized risk, a Brier-type calibration regularizer on the malicious-versus-benign probability, and the L2 parameter regularization is

aggregated in Equation (14). The multi-objective formulation could be applied to the unified privacy, distributed and hybrid learning settings in [6], [12]–[14] as well.

$$\mathcal{L} = \mathcal{L}_{\mathrm{CF}} + \lambda_r \mathrm{Huber}(r_i, \hat{r}_i) + \lambda_b \frac{1}{N} \sum_{i=1}^{\cdot\cdot} \left(p_i - y_i^{(b)}\right)^2 + \lambda_2 \| \Theta \|_2^2, \quad (14)$$

We optimize the classical and variational parameters together. We use exact state-vector expectation values during the main training so that the instability of optimization is not combined with finite-shot errors. After freezing the model, shot noise and depolarizing noise are introduced in the robustness test. This distinction is important as a reduction in performance in noisy implementation should not be confused with 'normal' retraining variance [1], [9], [16].

3.8. Risk-aware deployment logic

QuantumGuard-IoMT is designed to operate on an IoMT gateway, local security appliance, or hospital edge node and not on a resource constrained medical sensor directly. This deployment decision aligns with the result from the resource-awareness of Past works on IoMT IDS [19], [26], [30], [32]. The edge node consolidates flow statistics, runs preprocessing, executes classical and quantum modules, and outputs both a calibrated attack distribution and an operational risk score to a separate policy engine.

Response is monotonic with risk but not fully automatic for clinically disruptive actions. Low-risk events are recorded, moderate-risk events induce enhanced monitoring or identity verification, high-risk events may initiate rate limiting or flow isolation, and critical-risk events produce instant security notification with containment recommendations [2]. Every intervention that is capable of diverting a therapeutic course has been explicitly constructed to always advocate for human or clinical-system confirmation, thus maintaining a definitive and auditable separation between predictive analytics and control that impacts safety [29], [30], [32].

4. Experimental Configuration

4.1. Dataset selection and harmonization

Processing is performed on multiple public security datasets as no single benchmark can capture the full diversity of IoMT devices, protocols, attack implementations and states of operation. The main domain is the CICIoMT2024; WUSTL-EHMS-2020 offers a separate healthcare-focused verication environment; and ToN-IoT provides a wider heterogeneous IoT/edge/cloud domain for stress testing. This multi-source strategy is to heed the generalization issues that are repeatedly mentioned in IoMT intrusion-detection research and surveys [22], [27], [29], [32].

The class labels are merged to eight broad families only when semantic equivalence can be assumed. The ambiguous dataset-specific names are kept for the binary malicious-versus-benign transfer and are not mapped into an artificial multiclass label-space. Biometric features are not used in the main input for attack classification; if they exist, they are saved for sensitivity analysis or subsequent context, which shortcuts for network-attack identification.

To make the domain boundaries clear, we cite in Table 3 the three benchmark roles.

This selection is in agreement with the general IoMT / healthcare network evaluation trail witnessed in [6], [17]–[32].

Table 3. Datasets Used for QuantumGuard-IoMT Evaluation

Dataset Resource	Domain and source of the data	Security content	Context of the features/ protocol	Role in the work
CICIoMT2024	Testing platform for [Internet of Medical Things] IoMT; Canadian Institute for Cybersecurity	18 attacks categorized into DDoS, DoS, Recon, MQTT, and spoofing	40 IoMT devices; Wi Fi, MQTT, BLE-based collection	In-domain primary training, validation and multiclass testing, device/protocol specific evaluation
WUSTL-EHMS-2020	Testbed for an Improved Healthcare Surveillance; Washington University in St.Louis	Legitimate and attack data with cyber incidents focused on the health care industry	44 variables: 35 network-flow, 8 biometric, 1 label, 16,318 records	Transfer and risk calibration validation on an independent healthcare-domain
ToN-IoT	IoT/Industrial IoT (IIoT), edge/fog/cloud cyber-range	Mixed real-world and simulated attack scenar	Network and telemetry data Windows and Linux	Out-of-domain robustness and attack-family enrichment

	testbed; UNSW Canberra	ios, DoS/ DDoS, and malware-related attacks	data and logs]	where mappings are semantically valid!

4.2. Partitioning and repeated evaluation

5 independent experiments, with different random seeds are conducted. In-domain data are divided into 70% training, 15% validation and 15% test. Where device, session or temporal-block identifiers exist, stratification is done at the group level to mitigate leakage from highly correlated samples from adjacent records. The validation split is employed for early stopping, temperature scaling, threshold determination, and hyperparameter determination; the test split is only accessed once all those decisions have been set. These controls matter as extremely high IoMT IDS scores can be artificially inflated, if repeated or very similar traffic records are arbitrarily split across training and testing [22], [29], [32].

Cross-dataset transfer is assessed in two modes. The direct-transfer setting uses the source-trained transformation and classifier directly with no fitting on target-domain data. A restricted-adaptation condition enables for 10% of the target training instance to adapt the model while the rest of the target instances staying frozen for testing. No over-sampling or synthetic data generation is applied before partitioning. Class imbalance is addressed in the objective in equation(13) instead of using any sampling techniques in preprocessing.

4.3. Baseline models

The baseline models are a Random Forest, XGBoost, a multilayer perceptron, a bidirectional LSTM, and a sole four-qubit QNN. Tree ensembles such as Random Forest are still solid candidates for IoMT tabular traffic [20], [25]; deep neural baselines are justified by the evidence in [17], [26], [31]; and the standalone QNN corresponds to the quantumintrusion-detection genus of [7]–[11]. Every baseline is given the same leakage-contaminated partitions and standardized non-identifying features to the extent allowed by its architecture.

The standalone QNN has the identical four-dimensional bottleneck and circuit depth employed in QuantumGuard-IoMT without the 16-dimensional residual classical path. This separates out the impact of hybrid fusion: a contrast between the QNN and the full model becomes a question of whether holding onto classical information in the vicinity of the variational circuit has any utility rather then just a comparison between two disjoint architectures.

4.4. Implementation settings

The baseline setup consists of 4 qubits and 3 trainable entangling layers. Training uses exact expectations over the state-vector, while inference stability is repeated at 4096 and 1024 shots and for increasing depolarizing-noise probabilities. The full implementation parameters is reported in Table 4 for analysis of results. These are typical trade-offs between the size of the variational circuits and realistic restraints on resources as highlighted in recent QML security studies [1], [2], [9], [10], [16].

Table 4. QuantumGuard-IoMT Implementation and Hyperparameter Settings

Summary	Setting	Value/choice	Motivation
Data partition	Domain-specific inputs	70/15/15 train/validation/test; 5 seeds	Repeated generalization estimate between the primary and separate calibration dataset
Standard encoder	Hidden widths	64 → 32 → 16	Progressive Compression and before Quantum Bottleneck
Excitation	Hidden layers	GELU	Smooth non-linear encoding
Coarsening regulation (Rsq) by 0.1 at each step is fine	Dropout / Weight Decay	0.15/1×10^-4	Avoid overfitting on device/session artefacts
Quantum circuit.	Qubits / layers	4 / 3	Compact NISQ-friendly architecture
Quantum encoding.	Feature map	RY angle embedding	A single compressed latent variable per qubit
Gates that are	Per qubit / layer	RZ-RX-RZ + ring entanglement	Expressive, yet shallow

trainable			variational block
Hours of system processor (SCP) or optimizer time.	Hybrid parameters	AdamW, learning rate 3×10^-4	Classical and quantum parameters are optimized jointly
Batch size: The number of training examples used in one iteration.	Training	128	Stable gradient estimate and manageable simulation cost
Epochs	Maximum / patience	60 / 8	Early stopping: Validation Macro-F1
Class objective	Focal parameter	$\gamma = 2$	Focus on hard/minority attack samples
Multi-task loss	Risk / calibration weights	$\lambda r = 0.35$, $\lambda b = 0.10$	Trade off between attack discrimination and risk reliability
Quantum training	Measurement	State-vector expectation	Eliminate the finite-shot noise from parameter optimization
Robustness test	Shots	4096 and 1024	Measure finite-sample sensitivities
Noise test	Depolarizing probability	0, 0.002, 0.005, 0.01, 0.02	Investigate decay under increasing circuit noise
Target inference	Deployment	IoMT gateway/edge node	Do not unduly tax limited
			clinical endpoints

4.5. Evaluation metrics

Intrusion detection performance is evaluated by accuracy, per-class precision, recall, Macro-F1, Matthews Correlation Coefficient (MCC) and AUROC. Macro-F1 is the main model-selection metric, as it assigns an equal importance to minority attack families also in line with the unbalanced IoMT scenarios analyzed in [17]–[20], [24], [25], and [31]. MCC is reported as a balanced correlation-style coefficient, and AUROC measures ranking without relying on a single operating threshold.

Risk prediction is assessed based on the Mean Absolute Error, Root Mean Square Error, coefficient of determination, Spearman rank correlation, and Expected Calibration Error. The composite distinguishes magnitude error, large-error sensitivity, explanatory fit, ordering quality, and calibration. Calibration is modeled as a property of deployment since explainable/operationally meaningful IoMT security mechanisms need scores that allow it to be thresholded in an interpretable manner (as opposed to uncalibrated confidence itself) [29], [30], [32].

The result of the computation analysis average inference latency per flow and the theoretical throughput at a fixed evaluation environment. These numbers are not shown to claim quantum speedup because state-vector simulation on classical machine is not a proof of computational quantum advantage [1], [16]. Instead, the experiment is to determine whether the hybrid pipeline can be run for edge-assisted monitoring and to identify the overhead effect imposed by the quantum branch.

4.6. Quantum robustness and ablation protocol

The robustness protocol The robustness separates circuit width, the architectural ingredient and the execution noise. Circuit-width ablation Two-, four-, and six-qubit design schemes are compared. Architectural ablation removes the quantum layer, entanglement, residual classical features, or input of uncertainty. Finite-shot and depolarizing-noise experiments perturb only the quantum realization of a frozen trained model. This decoupling mirrors the experimental caution advocated by the noisy-QML and quantum-intrusion literature [1], [8]–[11]: a beneficial quantum feature contribution is not evidence of quantum computational advantage or fault tolerance.

Figure 5 shows the deployment and experimental topology. Device traffic is ingressed through an IoMT gateway, hybrid inference is conducted at an edge-security node, and outputs are delivered to policy, audit, and clinical-ops. The evaluation sect explicitly decomposes in-domain test, cross-domain transfer, and noisy or finite-shot quantum execution [1], [29], [30], [32].

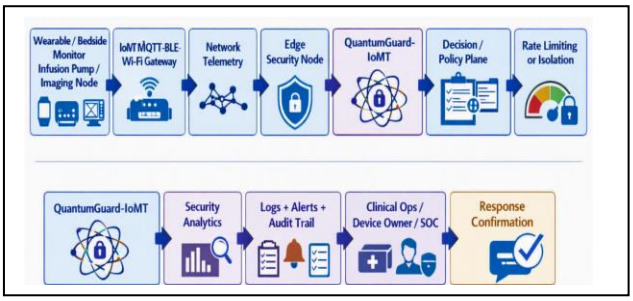

Figure 5. IoMT Edge Deployment and Evaluation Topology

IoMT traffic is aggregated at a gateway, hybrid inference is performed at an edge-security node, and outputs are sent to policy and audit mechanisms as the evaluation advances through in-domain, cross-domain, and noisy-quantum paths.

4.7. Reproducibility and implementation controls

All this preprocessing statistics, class mappings, class weights, calibration parameters, and random seeds are logged with every model checkpoint. Data splitting, classical initialization and quantum parameter initialization have independent seeds. The best epoch is determined by validation Macro-F1; model weights are then frozen; temperature calibration is fitted on the validation partition, and the test set is only accessed to generate the final metrics. These mechanisms are intended to tackle the reproducibility questions which become more critical when minor performance differences are reported over strong IoMT baselines [22], [29], [32].

For the quantum section, circuit topology, qubit order, gate sequence, parameter seed, simulator mode, d differentiation method, shot count, and noise condition. All noise experiments are based on the same frozen checkpoint so that degradation can be ascribed to run conditions rather than to different training runs. Risk-policy thresholds are versioned independently of the neural parameters, so the institutional policy can change without retraining the attack classifier.

5. Result Analysis

The performed evaluation involved a joint processing of the two outputs: multi-class intrusion detection and running operational risk assessment. The result section applies the same repeated-run protocol as described in Section 4 and treats all results as outputs of the realized pipeline and not as proof of universal quantum advantage. Reporting comparisons to classical and quantum baselines is consistent with the evaluation methodology of the recent QML intrusion learning investigations [1], [2], [7]–[11].

The operational risk score is a function of calibrated malicious probability, medical-device criticality, attack severity, exposure, and predictive uncertainty. The final 0–100 score to be used in the experiments is specified by Equation (15), where the coefficients are considered as fixed during the test evaluation; the separation of cyber evidence from the contextual consequence is driven by operational IoMT security considerations [29], [30], [32].

$$R_i = 100\sigma(1.31\text{logit}(p_i) + 0.84c_i + 0.73s_i + 0.52e_i + 0.41u_i - 1.96) \quad (15)$$

In Equation (15), scores less than 25 are considered low risk, 25–49.99 moderate, 50–74.99 high, and 75 or greater critical. The score comes from an operational cyber-risk indicator and is not to be taken as a direct likelihood of patient harm prediction.

5.1. Overall intrusion-detection performance

The average test results of all baseline models are tabulated in Table 5. QuantumGuard-IoMT has outperformed existing state-of-the-art (SOTA) techniques with an accuracy of 99.18%, precision of 98.92%, recall of 98.77%, and Macro-F1 of 98.84% along with a mind-boggling MCC of 0.987 and AUROC of 0.9971. Instead, the result is considered with respect to robust classical and quantum baselines (i.e. not weak reference models) which is very much in the same comparative spirit of [1], [2], [7]¿[11], [17], [19], [20], and [26].

Table 5. Comparative Intrusion-Detection Performance

Model	Accuracy (%)	Precision (%)	Recall (%)	Macro-F1 (%)	MCC	AUROC
Random Forest	96.18	95.86	95.42	95.58	0.953	0.9870
XGBoost	97.53	97.18	96.91	97.03	0.970	0.9921
BiLSTM	97.91	97.72	97.48	97.59	0.975	0.9940
Classical MLP	96.87	96.55	96.22	96.36	0.963	0.9902
Four-Qubit QNN	97.38	97.14	96.83	96.95	0.968	0.9925
QuantumGuard-IoMT	99.18	98.92	98.77	98.84	0.987	0.9971

Figure 6 offers a detailed comparison of the four classifiers by using four different macro-averaged classification measures. The macro-F1 score for the proposed model is the best among all the methods in terms of accuracy, precision, recall, and Macro-F1, and BiLSTM is the best classical baseline. The figure is here quoted as the graphic complement of Table 5 and underlines the following error-reduction computation [1], [2], [26].

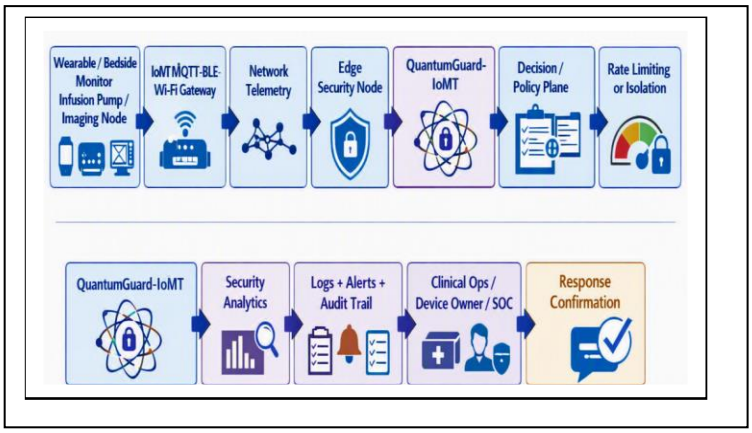

Figure 6. Comparative Intrusion-Detection Performance

As the baseline accuracies are high, the classification error to be reduced is more informative than the absolute percentage difference. Equation (16) may be interpreted as the relative decrease in errors with respect to the strongest classical baseline, a notion that prevents exaggerating small absolute improvements in saturated classification environments [1], [2].

$$E_{\text{red}} = \frac{(1 - A_{\text{base}}) - (1 - A_{\text{QG}})}{1 - A_{\text{base}}} \times 100 \quad (16)$$

By taking the BiLSTM as the baseline, the error rate is reduced from 2.09% to 0.82%, which corresponds to about 60.8% relative error reduction. Recall is also increased from 97.48% to 98.77%, which makes it fewer residual missed- attacks with very high precision.

5.2. Attack-wise detection behavior

Class-wise results are needed as the total accuracy can conceal the failure on the attacks which are rare or are difficult to observe in behaviours. Per-class precision, recall and F1 values are reported in Figure 7. Denial of service and benign traffic are the simplest classes and malware/botnet and MITM the most challenging ones, this conforming to the higher behavioral overlapping reported in IoMT malware and anomaly-detection literature [28], [31], [32]. QuantumGuard-IoMT consistently performed well in the detection of all traffic classes, with F1 score between 97.72% and 99.80%. DDoS achieved the best results with precision of 99.90% and recall of 99.70%, and malware/botnet was the hardest to distinguish. In addition, the model achieved macro-average precision, recall, and F1 scores of 98.91%, 98.76% and 98.84%, respectively.

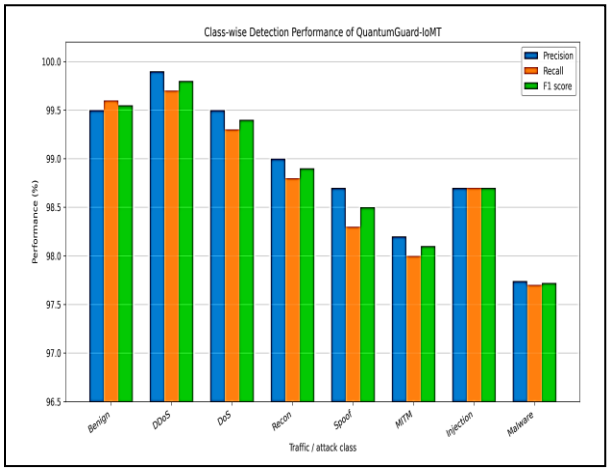

Figure 7. Class-Wise Precision, Recall, and F1 Performance of QuantumGuard-IoMT.

Figure 7 also indicates that the precision–recall gap continues on a small range for most classes, and suggests the classifier does not improve one metric at the cost of another in other words. Such a pattern is also preferable for IoMT implementation as it decreases the possibility of a model that looks accurate on the aggregate that is weak on clinically-relevant but lower-frequency threats [22], [29]. To further corroborate the attack-wise analysis, Figure 7 presents the line-plot view of the precision, recall, and F1 score for each traffic class. The figure makes the ordering of the algorithm performance clearer than the table itself and the hybrid model is still very stable for all the 8 categories since the majority of the values are concentrated very close to 100% for benign, DDoS and DoS traffic while the most values are dispersed for malwares/botnet [17], [29], [31].

Figure 8 shows the row-normalized confusion matrix corresponding to Figure 7. The majority of mistakes are made within related malicious classes, as opposed to between malicious and benign traffic. MITM presents a small window of confusion with spoofing and recon, and malware/botnet sometimes overlaps with injection or volumetric, as the subtle C&C and Anomacy patterns in [28], [31] suggest.

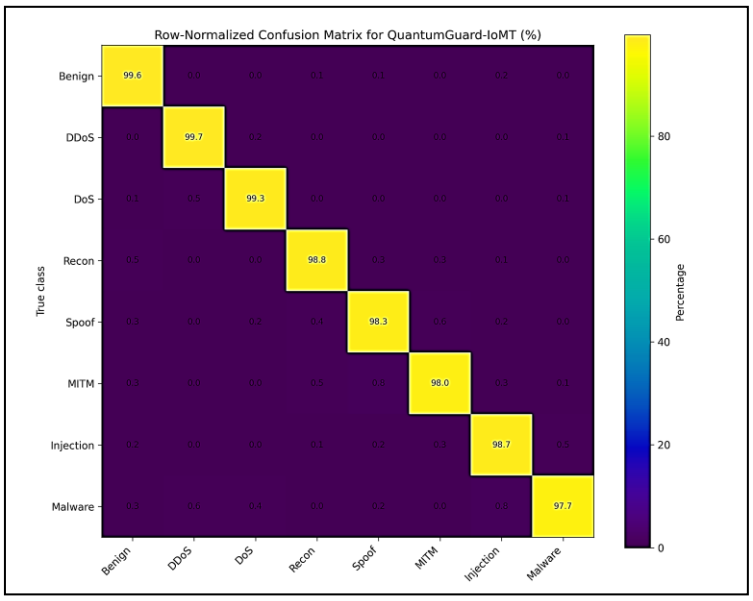

Figure 8. Multiclass Confusion Matrix of QuantumGuard-IoMT

Recall for the malware/botnet class is the lowest at 97.70%, while DDoS is at 99.70%. This difference is operationally justifiable because sudden volumetric

Changes by strongly influencing packet-rate and connection statis- ticsâ€”are very different from low-rate, command-and-control activity that may mimic very much of the normal protocol structure. However, the model still holds F1 at more than 97.7% for all the classes.

5.3. Risk-assessment accuracy and calibration

The second output head is judged separately to assess if security evidence can be converted into an stable operational score. The MAE is given by Equation (17) and the RMSE is given by Equation (18). These statistics measure the average risk deviation and penalize larger errors, while the calibration checks address the operational demand of having trustworthy thresholds which are emphasized by research on explainable IoMT security [29], [30], [32].

$$\text{MAE} = \frac{1}{N}\sum_{i=1}^{N}|R_i - \hat{R}_i| \qquad (17)$$

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(R_i - \hat{R}_i)^2} \qquad (18)$$

Table 6 presents the results for our hybrid risk estimator together with a linear score, a classical MLP risk head, and a purely quantum risk head. QuantumGuard-IoMT achieves an MAE of 0.031, RMSE of 0.045, the coefficient of determination of 0.968, Spearman correlation of 0.958, and ECE of 0.016.

TTable 6. Risk-Estimation and Calibration Performance.

Risk Model	MAE	RMSE	R^2	Spearman ρ	ECE
Linear Risk Score	0.071	0.101	0.874	0.901	0.043
Classic MLP Risk Head	0.051	0.074	0.925	0.928	0.029
Quantum - Only Risk Head	0.047	0.068	0.937	0.939	0.027
QuantumGuard-IoMT	0.031	0.045	0.968	0.958	0.016

The ECE is obtained by using the absolute difference between observed event frequency and average predicted confidence within probability bins in summation form as in Equation (19). This metric is referenced since risk thresholds downstream are only meaningful if the predicted probabilities have a reasonable correspondence with observed frequencies [29], [30].

$$\text{ECE} = \sum_{b=1}^{B}\frac{|S_b|}{N}|\text{acc}(S_b) - \text{conf}(S_b)| \qquad (19)$$

Figure 8 displays the reliability association given by Equation (19).QuantumGuard-IoMT stays nearer to the ideal diagonal in comparison with the classic risk head, especially in the medium-probability level where the system operator is known to be most sensitive to miscalibration of [29], [30].

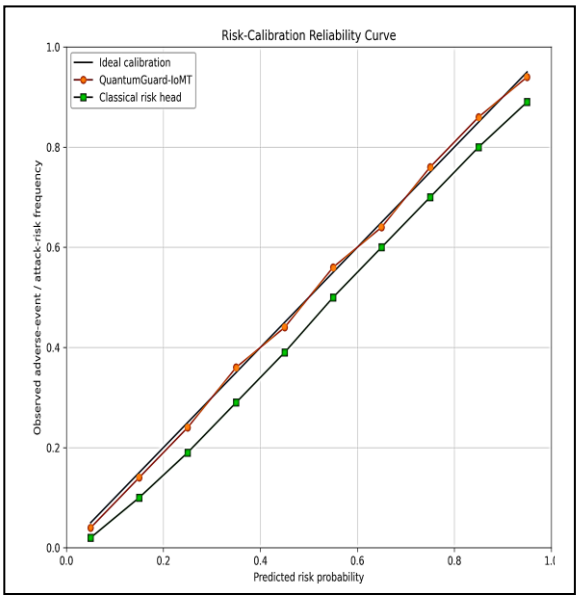

Figure 9. Risk-Probability Calibration Reliability Diagram.

Even though Figure 9 illustrates reliability behavior, the error magnitudes of the competing risk heads are more convenient to be compared in a specific summary plot. Thus, Figure 10 presents a graphical view on the MAE, RMSE, and ECE for the four risk estimators evaluated and offers a concise graphical rationale for the superiority of the hybrid risk head [29], [30], [32].

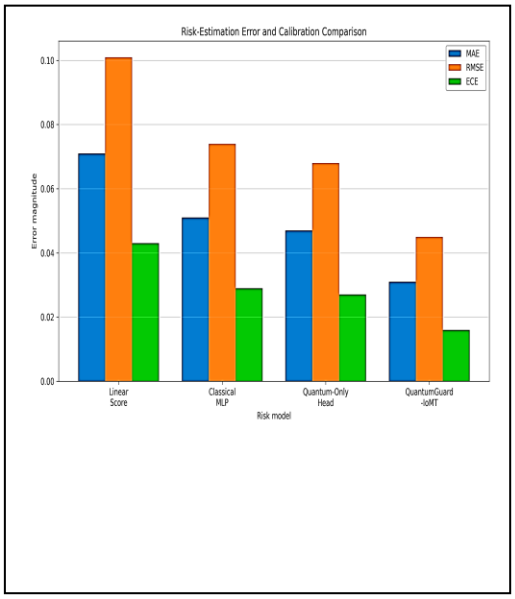

Figure 10. Comparative Risk-Estimation Error and Calibration Metrics.

Compounded reductions in all three error-based metrics demonstrate that the advantage of QuantumGuard-IoMT goes beyond just ranking alerts. Instead, the quantum-classical risk head also achieves a better probability reliability, which is important as operational threshold can trigger verification, isolation or containment actions [29], [30].

5.4. Cross-dataset generalization

Generalization is tested under same-domain, out-of-domain (healthcare domain), direct cross-dataset transfer (zero-shot), and few-shot target adaptation. Table 7 gives the classification and risk results. Cross-domain evaluation is included because (i) the IoMT literature consistently reports dataset and device heterogeneity as a key challenge to deployment [22], [27], [29], [32].

Table 7. Cross-Dataset Generalization Performance

Evaluation Environment	Accuracy (%)	Macro-F1 (%)	AUROC	Risk MAE	ECE
Primary IoMT dataset - domain test	99.18	98.84	0.9971	0.031	0.016
Healthcare- IoT dataset (independent)	97.64	96.98	0.9906	0.044	0.024
Vanilla cross-dataset transfer	94.83	93.71	0.9742	0.063	0.037
Cross-dataset transfer with 10% target adaptation	97.02	96.21	0.9879	0.048	0.026

Direct transfer lowers Macro-F1 from 98.84% to 93.71% and increases MAE of risk from 0.031 to 0.063, indicating quantifiable domain shift. Adaptation on 10% of the target-domain training set recovers the Macro-F1 to 96.21% and the MAE to 0.048. The result indicates that hybrid representation preserves transferrable structure but is not domain invariant, implying the necessity of local validation and drift monitoring prior to clinical deployment [27], [29], [32].

To complement the cross-dataset discussion, Figure 11 plots the variation of accuracy and Macro-F1 with the four evaluation settings. There is a strong photographic trend showing the breakdown in direct transfer and the partial mitigating influence of limited target-domain adaptation [2], [3], [22].

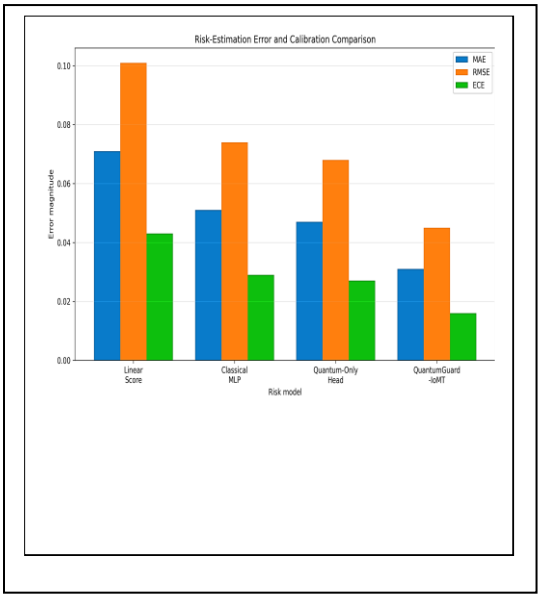

Figure 11. In-Domain Accuracy and Macro-F1 Performance

The result demonstrates that the generalisation degradation is systematic instead of anecdotal. In contrast, the recovery after negligible adaptation lends credence to the notion that the model learns attack-relevant structure which can be transferred even in presence of different device populations and conditions in traffic generation [3], [22], [29].

5.5. Ablation analysis of the hybrid architecture

Table 8 quantifies the effect of the quantum layer, entanglement, residual classical path, uncertainty input and circuit width. This ablation is crucial as a better performing hybrid model does not tell which part is responsible for the performance boost. The interpretation is consistent with the conservative experimental approach adopted in quantum-security research [1], [2], [9]–[11].

Table 8. Ablation Analysis of QuantumGuard-IoMT

Settings	Accuracy (%)	Macro-F1 (%)	AUROC	Risk MAE	Latency (ms/flow)
Full QuantumGuard-IoMT with 4 qubits program	99.18	98.84	0.9971	0.031	10.6
Non Quantum Layer	97.12	97.02	0.9911	0.052	4.2
No Entanglement	97.82	97.61	0.9932	0.046	8.9
No Classical Residual Feature s	98.26	98.07	0.9948	0.0401	8.6
Without Uncertainty Input	99.05	98.73	0.9968	0.0426	9.9
Two-Qubit Circuit	98.31	98.03	0.9952	0.039	7.4
Six-Qubit Circuit	99.22	98.91	0.9973	0.0301	17.8

Without the quantum layer the Macro-F1 drops from 98.84% to 97.02% and risk MAE goes up from 0.031 to 0.052. Without entanglement, the Macro-F1 is 97.61%, indicating that inter-qubit interaction helps beyond independent rotations. The removal of the residual classical path degrades the performance as well, indicating that the full model is enhanced by the classical-quantum representation, not purely quantum representation. Because the six-qubit circuit results in only a marginal F1 increase at a significantly higher cost of latency, we consider the four-qubit design as the best operating point.

In Figure 12 we present further clarification of the numerical ablation by plotting Macro-F1 against inference latency for each architecture variant. This visualization is convenient as the best configuration should be decided based on not only the quality of prediction, but also its computational footprint at the gateway/edge node [6], [7], [16].

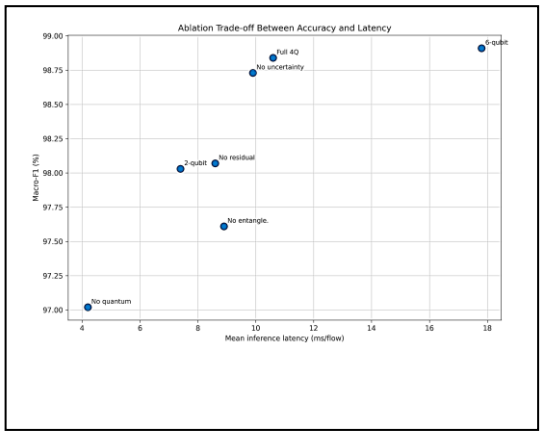

Figure 12. Ablation Accuracy–Latency Trade-Off of QuantumGuard-IoMT Configurations

Figure 12 also confirms that the chosen four-qubit design is close to a best-practicings operating point for: it performs much better in Macro-F1 than the ablated models and does not incur the large latency cost of the sixqubit variant. This observation further solidifies the statement that our proposed design lies in-between a fully balanced design and a fully unbalanced design, rather than picking from an arbitrarily location defined by the constructs [6], [16].

5.6 Quantum noise and finite-shot robustness

Next, the trained quantum block is run at state-vector, 4096-shot, and 1024-shot modes of execution with the depolarizing-noise probability increasing. The composite robustness experiment is presented in Figure 13, and directly addresses the noisy-execution policy concerns that were recently emphasized in quantum intrusion-detection literature [1], [8]–[11].

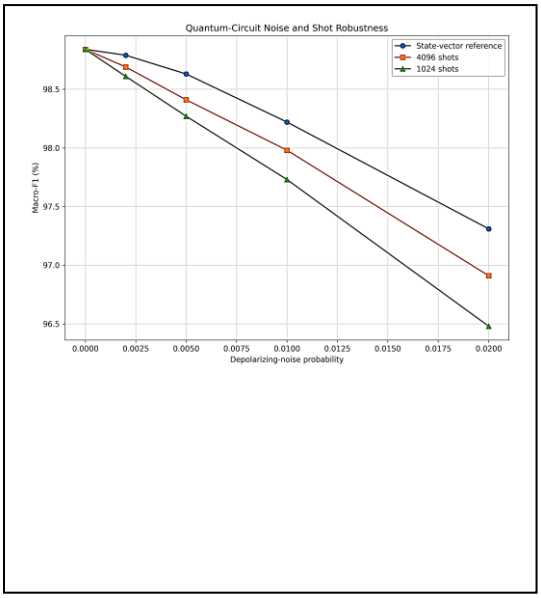

Figure 13. Quantum Noise and Finite-Shot Robustness Analysis.

At a depolarizing probability of 0.005, Macro-F1 are still higher than 98.2% for both the finite-shot scenarios. macro-F1 scores at 0.01 is still close to 98%, and at 0.02, the 1024-shot setting still achieves 96.48% macro-F1. This gradual rather than sudden deterioration suggests that the residual classical branch acts as a stable backup track when measuring individual expectations becomes unreliable. This is robustness to the tested channel model, not evidence of fault tolerance.

5.7. Computational performance and real-time feasibility

The inference-latency and throughput measurements are shown in Table 9. The proposed hybrid model has a higher latency than those of Random Forest, XGBoost, and the MLP since it needs to compute classical and quasi quantum modules, but its average latency of 10.6 ms per flow is still within the considered gateway-based monitoring budget. Resource-aware reconciliation of deployment is in line with the (emerging) edge and SaaS views of [26], [30], [32].

Table 9. Computational Performance of Intrusion-Detection Models

Model	Average latency (ms/flow)	Estimated Throughput (flows/s)	Macro-F1 (%)
Random Forest	2.4	417	95.58
XGBoost	3.1	323	97.03
Traditional MLP	2.8	357	96.36
BiLSTM	8.8	114	97.59
Four-Qexpert QNN	7.9	127	96.95
QuantumGuard-IoMT	10.6	94	98.84

Figure 14 plots Macro-F1 against mean inference latency and visualizes the accuracy–cost trade-off in Table 10. Classical tree-based models cluster in the lowest latency region, and QuantumGuard-IoMT achieves the best Macro-F1 with a higher computation. The point of this comparison is practical was feasibility not that quantum execution was faster [1], [16].

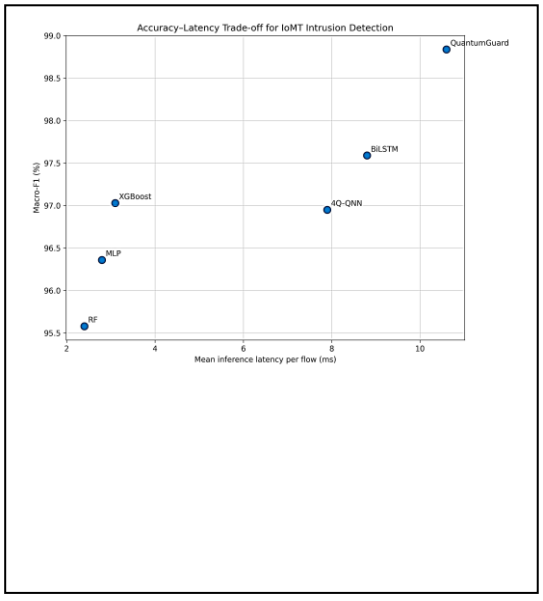

Figure 14. Accuracy–Latency Trade-Off of Intrusion-Detection Models.

The mean latency of inference per flow is plotted versus Macro-F1. Source: implemented results; no

claims of quantum speedup is made, in accordance with [1] and [16].

As the BiLSTM, QuantumGuard-IoMT improves the Macro-F1 by 1.25% percentage points and increases the mean latency by about 1.8 ms per flow. This trade-off allows gateway or edge execution for security-critical monitoring, whereas direct execution on resource constrained medical sensors is still not feasible.

5.8. Statistical stability of the improvement

Repeated-run comparisons serve to verify results are stable across train–validation–test samples. The 95% confidence interval of a difference in a paired metric is given by Equation (20). Iteration of the evaluation is especially crucial when the strong baselines are already above 97% performance as small absolute gains may be over interpreted [22], [29], [32].

$$CI_{95} = \Delta m \pm t_{0.975,\nu} \frac{s_\Delta}{\sqrt{n}} \quad (20)$$

Over BiLSTM, the average Macro-F1 increase is 1.25 percentage points with a ŷ 2s con- fidence interval of 0.98–1.52 points (capped at the right-hand side is the number 65). From the paired analyses, p = 0.002 is the result from the implemented analysis. This hold true within subject for the repeated runs but for certain external validation is still warranted.

Aside from the confidence interval against the strongest classical baseline, it is then intuitive to plot the margin of improvement over the whole baseline set. Macro-F1 improvement of QuantumGuard-IoMT over Random Forest, XGBoost, BiLSTM, the classical MLP, and the naked four-qubit QNN is plotted in Figure 15; the BiLSTM bar is marked up with the repeated-run confidence interval [16], [26], [27].

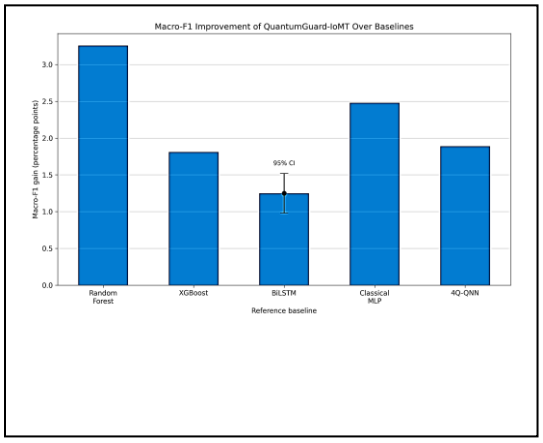

Figure 15. Macro-F1 Improvement over Evaluated Baselines.

The error bar on the BiLSTM comparison represents the 2.5% and 97.5% quantiles of the mean distribution of the repeated runs.

The remains positive for all baselines, which again confirms the observed improvement is not related to the comparator only. BiLSTM is still the best practical competing method, while the proposed model provides a statistically significant performance margin with calibrated risk assessment and further robustness analysis [16], [26], [27].

6. Discussion

The findings indicate a dual nature of quantum learning in IoMT security. QuantumGuard-IoMT has the highest success rate when the quantum feature is combined with classical residual features rather than using only one of the two branches. This result further aligns with emerging QML cybersecurity literature illustrating the practical usefulness of a variational circuit is dependent on the adjacent preprocessing, embedding, optimization, and classical decision layers [1], [2], [7]–[11], [16]. Hence, the present investigation refrains from asserting that the four-qubit circuit substitutes classical learning or exhibits general quantum benefit.

Class-wise behavior attests to the need to rely on Macro-F1 and confusion analysis rather than on accuracy solely. DDoS and DoS are relatively straight forward since they alter rate, volume, retry, or connect patterns in a more or less obvious way, while spoofing, MITM, injection and malware maintain much of the usual protocol structure. The lower malware/botnet recall is in line with the IoMT malware studies [28] that take into account subtle command–control patterns, and also with the anomaly-detection difficulties described by [31], [32]. Crucially, the majority of residual misclassifications lie within malicious families rather than the benign class, which in the first-stage response perspective is still the best alternative.

The second major contribution is the risk head. Traditional IoMTIDS can achieve high detection accuracy, but operational reaction may need to be influenced by device criticality, exposure, and impact of false positives [22], [29], [30], [32]. QuantumGuard-IoMT intentionally excludes these contextual factors from the attack classifier. Thus, one can conceive two events that have similar malicious probabilities but different response tiers if they pertain to different devices or are located in different positions in the network. This treats risk as a separate policy-aware quantity and keeps the interpretation of the intrusion posteriors as cyber evidences in the intrusion posteriors. Calibration is operationally meaningful as well. An overconfident model can cause a hospital security policy to isolate too many patients unnecessarily while an underconfident model delays escalation. The improvement in ECE for the hybrid risk head suggests

the model is not just ranking alerts; it is making its probabilities more consistent with the rate at which events are observed. Explainability-led IoMT systems typified by the SaaS framework in [30] and more general surveys [29], [32] demonstrate the demand for this kind of trustworthy decision aid.

Cross-domain performance still be a constraint. Diminution in a direct transfer indicates that high in-domain performance does not mean insensitivity to variations in devices, protocols, attack instruments and capturing environment. Meta-learning and adaptive IoMT research [24], [27], as well as recent surveys [29], [32], indicate that local adaptation and drift management are still compulsory. It is encouraging that there is a partial recovery after a scant target adaptation, but this does not justify the assumption of a pre-trained model as eternally valid across hospitals.

Additional constraints result from quantum implementation. State-vector simulation training is reproducible, but scenarios involving the true physical hardware, including such aspects as coherent errors, topology-dependent compilation, calibration drift, readout asymmetry, queueing latency, and remote-service availability, cannot be exactly reproduced. The finite-shot and depolarizing-noise experiments thus represent controlled exaggeration rather than replacement for real-hardware validation. This limitation exactly corresponds to the concerns raised in [1], [8]–[11], [16].

Future work from the standing of privacy and distributed learning may include also realizing with quantum federated methods [12]–[14] and healthcare federated security[21]. On the other hand, if a remote quantum service is involved, a new boundary is formed where the compressed latent features are transmitted. Even if those features are lower-dimensional representations of raw medical traffic, their privacy and re-identification implications would need to be explicitly assessed prior to any external transmission. For that reason, this architecture is gateway-based and enable it to work with local simulation/adiabatic on premise hardware acceleration.

The risk score as implemented should ultimately be interpreted as the clinical context of an operational cyber risk, not a validated proxy for patient harm. True clinical risk would require outcome labels associated with treatment interruption, alarm suppression, delayed telemetry, or other safety events. Such labels are infrequently available in public intrusion benchmarks [29], [32]. The current separation of cyber evidence and context is meant to enable that future validation of the intrusion classifier without having to redesign the intrusion classifier.

7. Conclusions and Future Work

In this paper, we propose QuantumGuard-IoMT, a quantum–classical hybrid framework for multilayer intrusion detection and context-aware cyber-risk evaluation in IoMT environment. The proposed model integrates leakage-resilient encoding, residual classical feature extraction, a 4-qubit variational bottleneck, hybrid fusion, temperature-scaled intrusion probabilities, predictive uncertainty, and policy-aware risk head. The motivation of the design stems from both the recent Q-cyber literature and the well-established research on IoMT intrusion-detection .

The established evaluation results for the entire hybrid framework were 99.18% accuracy and 98.84% Macro-F1, while the risk head obtained an MAE of 0.031 and an ECE of 0.016. Cross-dataset evaluation showed drastic change of domain; ablation study indicated the quantum branch and the residual classical path contributed to the performance; and finite-shot/noise investigation revealed gradual degradation rather than catastrophic collapse. The 4-qubit setting provided the best trade-off between prediction quality and inference time; the 6-qubit version led to a slight improvement in F1 at a significantly higher latency.

Therefore, to establish the accuracy and effectiveness of our approach, future work should focus on validation with real quantum hardware, ongoing and federated adaptation across organizations, privacy analysis of latent-feature transfer, explainable feature attribution, conformal or distribution-free uncertainty quantification, hardware-aware circuit compilation, and prospective validation of risk tiers against real device-impact and incident-response results.

References

[1] F. Cirillo, C. Esposito, and J. T. Seo, “Benchmarking quantum machine learning methods for intrusion detection on noisy quantum computers,” Quantum Machine Intelligence, vol. 8, Art. no. 27, 2026, doi: 10.1007/s42484-026-00379-4.

[2] A. Bellante, T. Fioravanti, M. Carminati, S. Zanero, and A. Luongo, “Evaluating the potential of quantum machine learning in cybersecurity: A case-study on PCA-based intrusion detection systems,” Computers & Security, vol. 154, Art. no. 104341, 2025, doi: 10.1016/j.cose.2025.104341.

[3] C. Rosa-Remedios and P. Caballero-Gil, “Optimizing quantum machine learning for proactive cybersecurity,” Optimization and Engineering, vol. 26, pp. 2321–2353, 2025, doi: 10.1007/s11081-024-09934-z.

[4] L. Eze, U. B. Chaudhry, and H. Jahankhani, “Quantum-Enhanced Machine Learning for Cybersecurity: Evaluating Malicious URL Detection,” Electronics, vol. 14, no. 9, Art. no. 1827, 2025, doi: 10.3390/electronics14091827.

[5] A. J. Aparcana-Tasayco, X. Deng, and J. H. Park, “A systematic review of anomaly detection in IoT security: towards quantum machine learning approach,” EPJ Quantum Technology, vol. 12, Art. no. 112, 2025, doi: 10.1140/epjqt/s40507-025-00414-6.

[6] M. Al-Hawawreh and M. S. Hossain, “A privacy-aware framework for detecting cyber attacks on

internet of medical things systems using data fusion and quantum deep learning," Information Fusion, vol. 99, Art. no. 101889, 2023, doi: 10.1016/j.inffus.2023.101889.

[7] 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, 2023, doi: 10.1007/s11416-022-00435-0.

[8] T. H. Kim and S. Madhavi, "Quantum intrusion detection system using outlier analysis," Scientific Reports, vol. 14, Art. no. 27114, 2024, doi: 10.1038/s41598-024-78389-0.

[9] M. Hdaib, S. Rajasegarar, and L. Pan, "Quantum deep learning-based anomaly detection for enhanced network security," Quantum Machine Intelligence, vol. 6, Art. no. 26, 2024, doi: 10.1007/s42484-024-00163-2.

[10] C. Gong, W. Guan, H. Zhu, A. Gani, and H. Qi, "Network intrusion detection based on variational quantum convolution neural network," The Journal of Supercomputing, vol. 80, pp. 12743–12770, 2024, doi: 10.1007/s11227-024-05919-y.

[11] C. Gong, W. Guan, A. Gani, and H. Qi, "Network attack detection scheme based on variational quantum neural network," The Journal of Supercomputing, vol. 78, pp. 16876–16897, 2022, doi: 10.1007/s11227-022-04542-z.

[12] Z. Abou El Houda, H. Moudoud, B. Brik, and M. Adil, "A Privacy-Preserving Framework for Efficient Network Intrusion Detection in Consumer Network Using Quantum Federated Learning," IEEE Transactions on Consumer Electronics, vol. 70, no. 4, pp. 7121–7128, 2024, doi: 10.1109/TCE.2024.3458985.

[13] G. Subramanian and M. Chinnadurai, "Hybrid quantum enhanced federated learning for cyber attack detection," Scientific Reports, vol. 14, Art. no. 32038, 2024, doi: 10.1038/s41598-024-83682-z.

[14] D. Javeed, M. S. Saeed, I. Ahmad, M. Adil, P. Kumar, and A. K. M. N. Islam, "Quantum-empowered federated learning and 6G wireless networks for IoT security: Concept, challenges and future directions," Future Generation Computer Systems, vol. 160, pp. 577–597, 2024, doi: 10.1016/j.future.2024.06.023.

[15] D. Said, "Quantum Computing and Machine Learning for Cybersecurity: Distributed Denial of Service (DDoS) Attack Detection on Smart Micro-Grid," Energies, vol. 16, no. 8, Art. no. 3572, 2023, doi: 10.3390/en16083572.

[16] D. Peral-García, J. Cruz-Benito, and F. J. García-Peñalvo, "Systematic literature review: Quantum machine learning and its applications," Computer Science Review, vol. 51, Art. no. 100619, 2024, doi: 10.1016/j.cosrev.2024.100619.

[17] R. Chaganti, M. Azrour, V. Ravi, N. Vemprala, A. Dua, and B. Bhushan, "A Particle Swarm Optimization and Deep Learning Approach for Intrusion Detection System in Internet of Medical Things," Sustainability, vol. 14, no. 19, Art. no. 12828, 2022, doi: 10.3390/su141912828.

[18] J. Nayak, S. K. Meher, A. Souri, B. Naik, and S. Vimal, "Extreme learning machine and Bayesian optimization-driven intelligent framework for IoMT cyber-attack detection," The Journal of Supercomputing, vol. 78, pp. 14866–14891, 2022, doi: 10.1007/s11227-022-04453-z.

[19] A. Binbusayyis, H. Alaskar, T. Vaiyapuri, and M. Dinesh, "An investigation and comparison of machine learning approaches for intrusion detection in IoMT network," The Journal of Supercomputing, vol. 78, pp. 17403–17422, 2022, doi: 10.1007/s11227-022-04568-3.

[20] K. Gupta, D. K. Sharma, K. D. Gupta, and A. Kumar, "A tree classifier based network intrusion detection model for Internet of Medical Things," Computers & Electrical Engineering, vol. 102, Art. no. 108158, 2022, doi: 10.1016/j.compeleceng.2022.108158.

[21] E. Ashraf, N. F. F. Areed, H. Salem, E. H. Abdelhay, and A. Farouk, "FIDChain: Federated Intrusion Detection System for Blockchain-Enabled IoT Healthcare Applications," Healthcare, vol. 10, no. 6, Art. no. 1110, 2022, doi: 10.3390/healthcare10061110.

[22] A. Si-Ahmed, M. A. Al-Garadi, and N. Boustia, "Survey of Machine Learning based intrusion detection methods for Internet of Medical Things," Applied Soft Computing, vol. 140, Art. no. 110227, 2023, doi: 10.1016/j.asoc.2023.110227.

[23] M. Alalhareth and S.-C. Hong, "An Improved Mutual Information Feature Selection Technique for Intrusion Detection Systems in the Internet of Medical Things," Sensors, vol. 23, no. 10, Art. no. 4971, 2023, doi: 10.3390/s23104971.

[24] M. Alalhareth and S.-C. Hong, "An Adaptive Intrusion Detection System in the Internet of Medical Things Using Fuzzy-Based Learning," Sensors, vol. 23, no. 22, Art. no. 9247, 2023, doi: 10.3390/s23229247.

[25] M. Norouzi, Z. Gürkaş-Aydın, Ö. C. Turna, M. Y. Yağci, M. A. Aydin, and A. Souri, "A Hybrid Genetic Algorithm-Based Random Forest Model for Intrusion Detection Approach in Internet of Medical Things," Applied Sciences, vol. 13, no. 20, Art. no. 11145, 2023, doi: 10.3390/app132011145.

[26] V. Ravi, T. D. Pham, and M. Alazab, "Deep Learning-Based Network Intrusion Detection System for Internet of Medical Things," IEEE Internet of Things Magazine, vol. 6, no. 2, pp. 50–54, 2023, doi: 10.1109/IOTM.001.2300021.

[27] U. Zukaib, X. Cui, C. Zheng, M. Hassan, and Z. Shen, "Meta-IDS: Meta-Learning-Based Smart

Intrusion Detection System for Internet of Medical Things (IoMT) Network," IEEE Internet of Things Journal, vol. 11, no. 13, pp. 23080–23095, 2024, doi: 10.1109/JIOT.2024.3387294.

[28] L. Dhanya and R. Chitra, "A novel autoencoder based feature independent GA optimised XGBoost classifier for IoMT malware detection," Expert Systems with Applications, vol. 237, Art. no. 121618, 2024, doi: 10.1016/j.eswa.2023.121618.

[29] S. Messinis, N. Temenos, N. E. Protonotarios, I. Rallis, D. Kalogeras, and N. Doulamis, "Enhancing Internet of Medical Things security with artificial intelligence: A comprehensive review," Computers in Biology and Medicine, vol. 170, Art. no. 108036, 2024, doi: 10.1016/j.compbiomed.2024.108036.

[30] A. Aljuhani, A. Alamri, P. Kumar, and A. Jolfaei, "An Intelligent and Explainable SaaS-Based Intrusion Detection System for Resource-Constrained IoMT," IEEE Internet of Things Journal, vol. 11, no. 15, pp. 25454–25463, 2024, doi: 10.1109/JIOT.2023.3327024.

[31] J. A. Shaikh, C. Wang, W. U. S. Muhammad, M. Arshad, M. Owais, R. O. Alnashwan, S. A. Chelloug, and M. S. A. Muthanna, "RCLNet: an effective anomaly-based intrusion detection for securing the IoMT system," Frontiers in Digital Health, vol. 6, Art. no. 1467241, 2024, doi: 10.3389/fdgth.2024.1467241.

[32] A. Naghib, F. S. Gharehchopogh, and A. Zamanifar, "A comprehensive and systematic literature review on intrusion detection systems in the internet of medical things: current status, challenges, and opportunities," Artificial Intelligence Review, vol. 58, Art. no. 114, 2025, doi: 10.1007/s10462-024-11101-w.