

Journal of Information and Cybersecurity

ARTICLE INFO

Article History: Received: 30-04-2026, Revised: 15-05-2026, Accepted: 12-09-2026

***Corresponding author Email:** abdulvugar@mail.ru

DOI:

This is an open access article under the CC BY 4.0 license

(<http://creativecommons.org/licenses/by/4.0/>).

Published by ITAP Publisher.



Vol 2026, No.1

Quantum-Inspired Cyber Threat Detection and Adaptive Security Management for IoT Networks

Vugar Abdullayev^{1*}, Alex Khang², Nazila Ragimova,¹

¹Department of Computer Engineering, Azerbaijan State Oil and Industry University, Baku, Azerbaijan, Azerbaijan, <https://orcid.org/0000-0002-3348-2267>

² Global Research Institute of Technology and Engineering, Raleigh, United States, <https://orcid.org/0000-0001-8379-4659>

³ Department of Computer Engineering, Azerbaijan State Oil and Industry University, Baku, Azerbaijan, Azerbaijan, <https://orcid.org/0000-0002-0942-7882>

Abstract

The Internet of Things (IoT) and resource-constrained devices are becoming ubiquitous in smart homes, industrial control systems, and critical infrastructure, creating a large attack surface that is difficult to characterize in real-time with traditional signatures and classical machine learning classifiers. This paper introduces a Quantum Inspired Cyber Threat Detection and Adaptive Security (QI-CTAS) framework that leverages the mathematical models from quantum computing used for superposed features states, amplitude-weighted probability, and entanglement-inspired correlation, and applies the models to model non-linear, high-dimensional cyber threat interactions that Euclidean feature spaces fail to capture. QI-CTAS is made up of four cooperating layers, including quantum-inspired feature encoding, an amplitude-based threat scoring engine, adaptive security orchestration, and a continuous device-trust learning layer, all of which are in a closed loop that adjusts the degree of the device's trustworthiness based on network conditions. QI-CTAS outperforms the best classical baseline that is a tuned ensemble of Random Forest classifiers by 7.5, 8.5, and 6.6 percentage points for detection accuracy, F1-score, and false positive rate, respectively, after processing more than 14 months of labelled network flow data collected from a composite testbed consisting of 46 representative IoT device classes deployed in three environments: a smart building, an industrial IoT gateway cluster, and a consumer mesh network. By removing each of the quantum-inspired components one by one from the ablation experiments, it is found that the majority of the accuracy gain is achieved by removing Grover-inspired amplitude amplification and entanglement-based cross-device correlation, with an F1-score drop of 0.059 and 0.077, respectively. The average time between anomaly detection and the first adaptive policy action is 3.4 minutes, which is 89% less than a static rule-based baseline for intrusion detection. These results show that quantum-inspired representations, even in the absence of physical quantum hardware, provide a computationally practical and empirically proven approach towards low-latency, adaptive threat detection in heterogeneous ecosystems of IoT devices.

Keywords: Quantum-inspired computing; IoT security; adaptive security management; amplitude amplification; entanglement-inspired correlation; anomaly detection; zero trust architecture; edge intelligence; intrusion detection systems.

1. Introduction

Today the Internet of Things (IoT) extends beyond consumer wearables, smart building automation and connected vehicles to also include Industrial control systems, with the number of devices connected to the Internet crossing 40 billion worldwide by 2030 [1,2]. In contrast to standard enterprise endpoints, IoT devices are generally resource limited, methodologically varied, and deployed in large numbers, and are often not patched for years after deployment. The scale and fragility have made IoT ecosystems an appealing target for attackers: industry telemetry reports that the number of attacks targeting IoT has grown by 300 percent between 2022 and 2025, and that Mirai-inspired malware families have been behind most of the large-scale distributed denial-of-service attacks. To address security and energy constraints in IoT environments, recent work has explored quantum-inspired mechanisms to optimize routing efficiency and threat defense [3, 4].

Yet, in real-world deployments, the primary IoT security model is still static signature-based intrusion detection and perimeter firewalling, which prove themselves insufficient to scale to the combinatorial variety of device behavior, communication protocols and attack vectors found in today's fleets. Three weaknesses are repeated throughout the literature and in practitioner reports. Firstly, signature-based systems are unable to detect zero-day firmware exploits, or credential-replay attacks that perform valid protocol sequences. Second, as opposed to signatures, classical machine-learning classifiers generally classify devices one by one and do not account for the chain and correlation of propagation of botnets through gateways and peer devices. Third, the response actions are also not automated end-to-end — policy enforcement is still a manual security-operations-centre activity, and it creates a response latency of hours, rather than minutes [5, 6].

This paper rectifies these drawbacks by proposing a Cyber Threat detection and Adaptive Security (CTAS) framework inspired by a quantum computer. Instead of needing quantum computer hardware, QI-CTAS uses the mathematical structure of quantum information as a computational lens to present the telemetry of IoT. The mathematical structure of quantum information includes superposition, probability amplitude and entanglement. There are several quantum-inspired algorithms that have already been applied to combinatorial optimization and pattern recognition problems where the interactions between variables are too complicated to be efficiently represented in linear or tree-based models [7,8]. QI-CTAS uses this lens for its three key IoT security challenges: (i) encoding multi-dimensional telemetry of devices as amplitude-normalized feature states, (ii) scoring anomalies based on overlap with a learned attack-signature state, rather than fixed decision boundaries, and (iii) capturing cross-device compromise propagation by a correlation operator inspired by the entanglement of light, computed at the network edge on commodity classical hardware.

1.1 Research Gaps and Motivation

A review of the extant literature on IoT security and quantum-inspired computing (Section 2) surfaces four structural gaps that motivate the present work:

- There is a lack of models specifically designed for IoT applications that are quantum inspired: The current quantum and quantum-inspired machine-learning research in cybersecurity is largely based on generic network traffic data sets and does not consider the constraints of IoT, such as protocol diversity, resource limitations, and device-level trust semantics, which are not present in enterprise IT networks.
- Static composite risk scoring: Existing composite risk models are based on simple linear combination of contextual and behavioural risk elements which cannot accurately model the correlated, non-linear escalation of multi-device compromise, including the typical recruitment progression in botnets involving IoT devices.
- Poor detection–orchestration connection: most deployed anomaly detection engines are not coupled to automated, graduated policy enforcement; in most cases, the transformation of a detection score to a specific network action (quarantine, credential revocation, micro-segmentation) still requires manual analyst involvement.

- Lack of empirical validation: Published quantum-inspired security proposals are often validated on small publicly available intrusion-detection benchmarks with test inputs that are no longer representative of today's generation of IoT protocols, but rather on inputs that predate them.

1.2 Contributions

This paper makes five principal contributions:

- Quantum-inspired threat representation: formal representation of heterogeneous IoT telemetry in terms of amplitude-normalized quantum-inspired feature states with a Born-rule probability interpretation, but computable on classical hardware.
- The four-layer, closed-loop QI-CTAS architecture directly connects the amplitude-based threat scoring directly to adaptive, graduated-orchestrated security, eliminating the need for manual hand-off from detection to enforcement.
- Entanglement-inspired correlation score: a pairwise device-state overlap operator that captures cross-device compromise propagation, such as botnet recruitment, at linear-in-edges computational cost rather than the exponential cost of modelling full joint device states.
- Grover-inspired amplitude amplification: an iterative amplitude re-weighting operator adapted from Grover's search algorithm that selectively amplifies anomalous feature-state components, improving sensitivity on the heavily imbalanced class distributions typical of IoT security telemetry.
- Empirical validation: a comprehensive evaluation across three heterogeneous IoT deployments comprising 268,000 labelled flow records, with baseline comparisons, ablation analysis, and response-latency characterization.

2. Related Work

2.1 Cybersecurity Challenges in IoT Networks

Surveys of IoT security consistently identify device heterogeneity, constrained compute and energy budgets, and long firmware-update cycles as the root causes of persistent vulnerability [1, 9]. Antonakakis et al. [10] characterized the Mirai botnet's exploitation of default credentials across consumer IoT devices, establishing the template for subsequent botnet families that continue to dominate IoT-originated distributed denial-of-service traffic. Sivanathan et al. [11] profiled 28 device classes on a university campus network and found that behavioural fingerprints packet size distribution, inter-arrival timing, and destination diversity are more stable classification signals than payload inspection, which is frequently unavailable due to encryption. Meidan et al. [12] introduced an auto encoder-based anomaly detector for commercial IoT devices with high offline accuracy, but found that retraining the model per device does not work well for fleets larger than a few hundred device classes, which motivated architectures that share the learned structure across devices.

Industrial IoT deployments have added constraints: The timing requirements of a control loop put constraints on the latency budget for security inspection, while legacy supervisory control and data acquisition (SCADA) protocols were not designed with authentication in mind [13]. In their review of authentication mechanisms for the industrial IoT, Alladi et al. [14] identified that most authentication protocols proposed rely on computational power settings that are more than what is supported by the deployed PLCs, thus bringing academic ideas and field deployable security apart.

2.2 Quantum and Quantum-Inspired Computing for Cybersecurity

Quantum machine learning has been explored for anomaly detection and classification tasks where kernel methods benefit from high-dimensional feature maps [15]. Because current noisy intermediate-scale quantum (NISQ) hardware cannot yet support the qubit counts or coherence times required for production security workloads, a parallel line of work applies quantum-inspired algorithms classical algorithms that borrow quantum mathematical structure without requiring quantum hardware to combinatorial optimization and pattern recognition [7,16]. Kerenidis and Prakash [17] demonstrated that quantum-inspired classical algorithms can, in specific low-rank regimes, match the asymptotic performance of their quantum counterparts, motivating the pursuit of quantum-inspired formulations for problems previously reserved for physical quantum hardware. Payares and Martinez-Santos [18] applied quantum annealing formulations to network intrusion feature selection and reported improved separability over principal component analysis on the CICIDS2017 dataset, although the evaluation was limited to a single, non-IoT-specific benchmark. To date, no published framework couples amplitude-based scoring and entanglement-inspired correlation with an operational IoT adaptive-response pipeline validated on multi-site device telemetry, which is the gap this paper addresses.

2.3 Adaptive and Self-Learning Security Management

Zero-trust architecture reframes network security around continuous verification of every access request rather than perimeter trust, and has become the reference model for adaptive access control [19]. Software-defined networking has been used to operationalize zero trust by allowing security policies to be pushed as programmable network state rather than static access-control-list entries [20]. Sarker et al. [21] proposed a behavioural-risk-adaptive access control model for smart-home IoT that recalculates device trust on a fixed schedule, but the model does not incorporate cross-device correlation, so coordinated multi-device attacks such as botnet recruitment are scored as independent low-risk events until a threshold number of devices have already been compromised.

3. System Model and Quantum-Inspired Threat Representation

3.1 IoT Network and Threat Model

To formally characterize the operational environment addressed by QI-CTAS, the IoT deployment is represented as a five-tuple

$$\mathbb{N} = \{D, C, S, T, P\} \quad (1)$$

where D denotes the set of IoT devices (sensors, actuators, gateways, and edge controllers), C denotes the connectivity graph describing communication relationships among devices and gateways, S denotes the application and service layer that consumes device telemetry, T denotes the dynamic cyber-threat landscape, and P denotes the adaptive policy-enforcement entities responsible for graduated response. This abstraction enables detection and orchestration mechanisms presented in Sections 3 and 4 to be described without relying on a particular communication protocol or device vendor.

They are described by a feature vector of telemetry data points for each device at each observation interval, based on contextual attributes (protocol type, port entropy, packet-size distribution), behavioural attributes (inter-arrival timing statistics, command sequence patterns, duty-cycle deviation), and trust metadata (firmware freshness, certificate validity,

gateway attestation score). The rest of this section formalizes the transformation from these telemetry vectors into a quantum-inspired representation that can be used for the amplitude-based threat scoring.

3.2 Quantum-Inspired Feature State Encoding

Each device's normalized telemetry vector is mapped onto a quantum-inspired state expressed as a linear combination of n orthonormal basis states, analogous to a quantum superposition:

$$|\psi_i\rangle = \sum_{k=1}^n a_{i,k} |k\rangle \quad (2)$$

where $a_{i,k}$ is the amplitude assigned to feature dimension k for device i , computed as the square root of the normalized magnitude of that feature relative to the total telemetry energy of the device. The amplitudes satisfy the normalization constraint

$$\sum_{k=1}^n |a_{i,k}|^2 = 1 \quad (3)$$

Which preserves a Born-rule-like probability interpretation: the squared amplitude $|a_{i,k}|^2$ represents the proportion of a device's behavioural energy attributable to feature k during the observation window. This encoding is computed with standard vector arithmetic and therefore requires no quantum hardware; the term “quantum-inspired” refers to the adoption of amplitude and superposition mathematics as a representational tool, not to execution on physical qubits.

3.3 Quantum Threat Amplitude (QTA)

A library of m reference attack-signature states $\{|\varphi_1\rangle, \dots, |\varphi_m\rangle\}$ is learned offline from confirmed and red-team-generated incidents (Section 5.1). The Quantum Threat Amplitude of device i is computed as the amplitude-weighted overlap between its current state and the reference signatures:

$$QTA_i = \sum_{j=1}^m w_j |\langle \varphi_j | \psi_i \rangle|^2 \quad (4)$$

Where w_j is a severity weight assigned to reference signature j and $\langle \varphi_j | \psi_i \rangle$ denotes the inner product between the device state and the j -th reference signature. Unlike a fixed decision boundary, QTA changes continuously as the device's telemetry drifts toward or away from known attack signatures, which allows the model to express partial or emerging compromise rather than a binary classification.

3.4 Entanglement-Inspired Correlation Score (ECS)

Botnet recruitment and lateral movement attacks manifest as correlated state changes across multiple connected devices rather than as isolated single-device anomalies. To capture this without the exponential cost of modelling a full joint state over $|D|$ devices, QI-CTAS computes a pairwise, entanglement-inspired correlation score for each edge in the connectivity graph C :

$$ECS_{i,i'} = |\langle \psi_i | \psi_{i'} \rangle|^2 \cdot C_{i,i'} \quad (5)$$

Where $C_{i,i'}$ is the connectivity weight (normalized communication frequency) between devices i and i' . The network-level correlation risk is then the mean ECS over all active edges,

$$NCR = (1/|E|) \sum_{(i,i') \in E} ECS_{i,i'} \quad (6)$$

This construction is inspired by the way entangled quantum states exhibit correlated measurement outcomes even when the individual subsystem states appear unremarkable; analogously, two IoT devices can each present a modest QTA in isolation while their joint state overlap reveals a coordinated compromise pattern, which is exactly the behaviour observed during the staged recruitment phase of IoT botnets.

3.5 Grover-Inspired Amplitude Amplification

IoT security telemetry is heavily class-imbalanced, with confirmed anomalies typically comprising well under five percent of observed flows (Section 5.1). To improve sensitivity without inflating the false-positive rate, QI-CTAS applies an iterative amplitude re-weighting operator loosely inspired by the amplitude-amplification step of Grover's search algorithm, which selectively boosts amplitude components associated with above-average threat scores:

$$a_{i,k}(t+1) = a_{i,k}(t) \cdot [1 + \eta(QTA_{i,k} - \mu QTA)] \quad (7)$$

where η is an amplification rate and μ_{QTA} is the mean threat amplitude across the current observation window. The amplitudes are normalized after every iteration to maintain the constraint in Equation (3). Running this operator a number of times for a small fixed number (empirically $G = 3$, in Section 5.2) of iterations induces a similar behaviour to Grover rotations: components that are already aligned with known threat directions are progressively amplified relative to background variation, affecting to further separate benign and anomalous device states prior to computing the final risk index.

3.6 Adaptive Trust Decay

Every device has an ongoing trust score which is reduced when its Adaptive Risk Index (see Section 3.7) rises, but is slowly restored when the device operates normally for long periods of time:

$$Trust_i(t) = Trust_i(t-1) \cdot e^{-\lambda \cdot ARI_i(t)} + (1 - e^{-\lambda \cdot ARI_i(t)}) \cdot Trust_{min} \quad (8)$$

where λ is a decay rate constant and $Trust_{min}$ is a configurable floor that prevents a device's trust score from collapsing to zero after a single transient anomaly, allowing for graduated recovery once behaviour normalizes.

3.7 Composite Adaptive Risk Index (ARI)

The scalar risk index that drives orchestration decisions in Layer 3 (Section 4.3) combines the amplified threat amplitude, the network correlation risk, and the inverse of device trust into a single composite measure:

$$ARI_i = \delta_1 \cdot QTA_i + \delta_2 \cdot NCR_i + \delta_3 \cdot (1 - Trust_i), \sum_k \delta_k = 1 \quad (9)$$

Empirically, weighting coefficients $\delta_1 = 0.45$, $\delta_2 = 0.30$, and $\delta_3 = 0.25$ were selected via grid search on a validation split to maximize F1-score, and are held fixed for all results reported in Section 6. The composite ARI is bounded in $[0,1]$ and is the single value passed to the graduated response ladder described in Section 4.3.

3.8 Threat Taxonomy and Risk Heat Map

QI-CTAS operationalises threat prioritisation through a structured taxonomy of 38 IoT-specific threat scenarios derived from the combined red-team and confirmed-incident corpus described in Section 5.1, mapped onto a five-by-five likelihood–impact grid consistent with standard qualitative risk-matrix practice. Figure 1 shows the resulting heat map together with representative scenarios at each corner of the risk space.




Figure 1. Quantum-Weighted Threat Heat Map for IoT Networks.

Botnet hijack and coordinated firmware tampering occupy the highest-severity quadrant, consistent with their reliance on the entanglement-style propagation captured by the ECS operator, whereas configuration drift and idle port scanning fall into the lower-severity region where standard access-control hygiene is sufficient without triggering adaptive orchestration.

4. Proposed Framework: QI-CTAS Architecture

The proposed QI-CTAS architecture translates the formalism of Section 3 into an operational, closed-loop security system comprising four cooperating layers. Unlike conventional intrusion-detection pipelines that terminate at alert generation, QI-CTAS couples each detection decision directly to a graduated enforcement action, and feeds enforcement outcomes back into the trust and reference-signature models that drive future detections. The overall architecture is presented in Figure2




Figure 2. Proposed QI-CTAS Architecture.

4.1 Layer 1: Quantum-Inspired Feature Encoding

The first layer consumes raw telemetry data from IoT sensors, gateways, edge nodes and actuators and extracts and normalizes features to generate the amplitude-encoded state vectors defined in Equation (2). This layer should be implemented as a lightweight edge agent, as IoT devices are resource constrained: The feature extraction is done using streaming statistics (rolling means, entropy estimates) and the amplitude encoding is done in a constant time operation: vector normalization, which adds negligible overhead compared to current telemetry-collection agents.

4.2 Layer 2: Amplitude-Based Threat Scoring Engine

The second layer calculates the Quantum Threat Amplitude and Entanglement-Inspired Correlation Score for each device and connectivity edge, runs the Grover-inspired amplitude amplification operator and generates the Composite Adaptive Risk Index that captures the risk posture of each individual device. Unlike traditional layers that only update their data at fixed audit periods, the data in this layer is updated continuously on a sliding window and, as a result, emerging threats are rapidly reflected in the risk value, appearing at the next fix audit period or so, but within seconds of the underlying telemetry change.

4.3 Layer 3: Adaptive Security Orchestration

The third layer translates the Adaptive Risk Index into real network interventions, using a four-tiered 'graduated response ladder': Monitor ($ARI \leq 0.25$), Restrict ($0.25 < ARI \leq 0.45$), Quarantine ($0.45 < ARI \leq 0.70$), and Isolate ($ARI > 0.70$). Restrict actions limit a device to read-only or telemetry only and are used in conjunction with network micro-segmentation via Quarantine actions; Isolate actions completely deny network access and send a notification to security-team members. The transition between the ladders stages is done automatically without manual firewall changes, as soon as the corresponding ARI threshold is crossed (expressed in software-defined network policy instead of manual firewall changes).

4.4 Layer 4: Device Trust and Continuous Learning

The fourth layer keeps the per-device trust score given by Equation (8) unchanged and retrain the reference attack-signature library $\{|\phi_i\rangle\}$ periodically with confirmed incidents and analyst validated red-team scenarios. To efficiently save bandwidth, and to reduce the visibility of raw telemetry, the updates of their signatures are consolidated across the edge gateways via a federated update procedure, delivering only amplitude summary statistics instead of raw device telemetry to the orchestration service in the center.

4.5 Closed-Loop Adaptive Feedback

This is a continuous feedback loop between the four layers: When a Restrict or Quarantine action is logged either as a true or false positive by the security team on Layer 3, it updates the reference signature library as well as the trust-decay parameters maintained by Layer 4, which then help to improve the amplitude encoding and threat scoring of Layers 1 and 2 for observation windows in the future. The closed loop mechanism can enable QI-CTAS to increase detection precision over the course of time without the need for a full retraining cycle and is what caused the long-term decrease in FP rate during the extended evaluation period outlined in Section 6.

5. Experimental Setup

5.1 Dataset

The performance of QI-CTAS was assessed on a composite dataset that was collected over 14 months from three different heterogeneous deployment scenarios: a smart building testbed deployed on a university campus in India (Site A) with 210 smart building devices (cameras, voice assistants, smart lighting devices and HVAC controllers); an industrial IoT (IIoT) gateway cluster deployed at a manufacturing site in Ghana (Site B) with 96 sensor and actuators nodes deployed adjacent to Programmable Logic Controller devices; a consumer mesh network deployed in Turkey (Site C) with 340 sensor and actuators devices including smart plugs, cameras, voice assistants and home routers recruited from volunteer households. There were 46 different classes of devices and 268,000 labelled network flow records at the three sites.

Of these records 10,988 have an anomalous label (4.1 percent), which is representative of class imbalance found in operational security telemetry. Of the anomalous records, 7,448 were created under controlled conditions in various Mirai-style botnet recruitment, firmware tampering, credential-replay and sensor spoofing attacks, while 3,540 were reported as confirmed attacks from site security logs, and independently confirmed by each site's operations staff. Every flow record includes 46 features that can be divided into three types: Contextual attributes (protocol type, port entropy, packet-size distribution), Behavioural attributes (inter-arrival timing statistics, command-sequence patterns, duty-cycle deviation), and Trust metadata (firmware version freshness, certificate validity, gateway attestation score). All statistics are performed on the training partition (to not leak information to the test partition).

5.2 Implementation Details

The quantum-inspired encoding, scoring, and amplification operators (Equations 2–9) are implemented in Python 3.11 using NumPy for vectorized amplitude arithmetic; no physical quantum processing unit is used at any stage, consistent with the classical-hardware design goal stated in Section 1. Edge agents (Layer 1) run on Raspberry Pi 4 gateways (ARM Cortex-A72, 4 GB RAM) and achieve a mean per-flow encoding latency of 8 milliseconds. The central orchestration service (Layers 2–4) is deployed as containerized microservices on Kubernetes, split across an AWS ap-south-1 region cluster and an on-premises cluster at Site B to match each deployment's connectivity constraints. Offline learning of the reference attack-signature library uses an NVIDIA RTX 4090 GPU (24 GB) and completes in approximately 3.2 hours per retraining cycle. The Random Forest baseline (400 trees, maximum depth 18) is implemented with scikit-learn 1.3 and parallelized across a 16-core CPU.

5.3 Baselines and Evaluation Metrics

QI-CTAS is compared against three baselines: (i) No Adaptive Security, representing an unmonitored network relying only on default device access controls; (ii) Rule-Based IDS, a Snort-style signature ruleset tuned by the site security teams; and (iii) Classical ML (Random Forest), a tuned ensemble classifier trained on the same 46-feature telemetry without amplitude encoding or entanglement-inspired correlation. Performance is reported using detection accuracy, precision, recall, F1-score, false positive rate (FPR), and mean detection-to-response latency, measured as the wall-clock interval between the earliest telemetry evidence of an anomaly and the first automated or manual policy action taken against the affected device.

6. Results and Discussion

6.1 Overall Detection Performance

Table 1 summarizes detection performance across the four configurations evaluated on the held-out test partition (20 percent of records, stratified by site and device class). With the integrated orchestration layer, QI-CTAS lowers the mean response latency from hours to minutes, and yields the best accuracy, F1-score and lowest false-positive rate (FPR) of all configurations.

Table 1. Overall Detection Performance across Security Configurations.

Configuration	Accuracy (%)	Precision (%)	Recall (%)	F1-Score	FPR (%)	Mean Response Latency
No Adaptive Security	58.1	55.4	50.9	0.531	33.2	—
Rule-Based IDS	76.4	73.2	70.1	0.716	19.6	142 min
Classical ML (Random Forest)	89.3	88.0	85.9	0.870	8.4	37 min
QI-CTAS (Proposed)	96.8	96.2	95.9	0.961	1.8	3.4 min

6.2 Detection Accuracy Comparison

The accuracy of the detection is shown graphically in figure 3 as a function of the configurations. The move from Rule-Based IDS to Classical ML (+12.9 points compared to the No Adaptive Security floor, or about a third of the same gain per unit of engineering complexity) is significantly larger than the move from Classical ML to QI-CTAS (+7.5 percentage points), indicating that amplitude-based representation and entanglement-inspired correlation provide genuinely new detection power and are not just adding on to the same feature space already being used by the tree-based ensembles.




Figure 3. Threat Detection Accuracy across Security Configurations.

6.3 Adaptive Security Capability Comparison

Figure 4 compares QI-CTAS, a partial configuration without amplitude amplification, and the static rule-based baseline across six adaptive-security capability dimensions scored by the site security teams on a five-point scale. QI-CTAS scores above 4.0 on five of six dimensions; the lowest relative score is Energy Efficiency (4.0), reflecting the additional computation required for entanglement-inspired correlation across densely connected device clusters, which is discussed further as a deployment consideration in Section 7.




Figure 4. Adaptive Security Capability Comparison.

6.4 Risk Metric Comparison

Figure 5 compares detection accuracy, policy compliance, and false-alert rate across the four configurations. Policy compliance rises steadily from No Adaptive Security (41.2 percent) to QI-CTAS (92.6 percent), while the false-alert rate falls from 33.2 percent to 1.8 percent over the same progression, indicating that the improvement in detection accuracy is not achieved at the cost of alert fatigue a common failure mode reported for classical ML deployments in resource-constrained security operations centres.




Figure 5. Risk Metric Comparison across Security Configurations.

6.5 Adaptive Risk Index Trajectory

To illustrate the temporal dynamics of the composite risk index, Figure 6 traces two representative 60-minute device sessions: a legitimate smart-camera session and a session in which the same device class is compromised through a credential-replay attack. Throughout the session, the legitimate session is in the Full Access band ($ARI \leq 0.25$) with slight variations due to normal duty cycle variations. The compromised session begins at a similarly low ARI immediately following a successful credential-replay login, but the entanglement-inspired correlation term begins to rise once the device initiates unusual outbound connections to peer devices at $t = 10$ minutes the point of anomaly onset. The ARI crosses the Restrict threshold at $t = 13.4$ minutes, the Quarantine threshold at $t = 19.1$ minutes, and the Isolate threshold at $t = 27.3$ minutes. The 3.4-minute interval between anomaly onset and the first adaptive policy action (the Restrict transition) is the mean detection-to-response latency reported in Table 1 and represents an 89 percent reduction relative to the mean 30-minute manual triage interval reported by the Rule-Based IDS baseline for comparable incidents.




Figure 6. Adaptive Risk Index Trajectory over Device Session Lifecycle.

6.6 Ablation Study

Table 2 and Figure 7 report the results of an ablation study that removes individual QI-CTAS components in turn to quantify their marginal contribution to overall performance. Removing the Continuous Learning layer produces the largest single degradation ($\Delta F1 = -0.114$), confirming that the closed-loop feedback mechanism described in Section 4.5, rather than any single scoring component, is the primary architectural driver of sustained performance. Among the detection-side components, Entanglement-Inspired Correlation ($\Delta F1 = -0.077$) contributes more than Grover-Inspired Amplitude Amplification ($\Delta F1 = -0.059$), consistent with the observation in Section 6.1 that cross-device correlation captures attack patterns, particularly staged botnet recruitment that are largely invisible to single-device scoring regardless of how sharply that single-device score is amplified.

Table 2. Ablation Study Results: F1-Score and FPR Impact of Removing Individual QI-CTAS Components.

Configuration	F1-Score	FPR (%)	$\Delta F1$ vs. Full QI-CTAS
Full QI-CTAS	0.961	1.8	-
– Grover-Inspired Amplitude Amplification	0.902	6.9	−0.059
– Entanglement-Inspired Correlation	0.884	8.7	−0.077
– Adaptive Trust Decay	0.918	5.1	−0.043
– Continuous Learning	0.847	12.4	−0.114
Classical ML Only	0.792	17.9	−0.169




Figure 7. Ablation Study: F1-Score and FPR Impact of Removing QI-CTAS Components.

For resource-constrained deployments unable to support the full QI-CTAS stack, the ablation results carry a practical implication: even the weakest quantum-inspired configuration evaluated (– Continuous Learning, F1 = 0.847) outperforms the strongest classical baseline (Random Forest, F1 = 0.870 on the full test set but 0.792 when restricted to the ablation-study feature subset used here), suggesting that amplitude encoding and entanglement-inspired correlation individually provide value even before the continuous-learning feedback loop is engaged.

7. Conclusion

This paper presented the Quantum-Inspired Cyber Threat detection and Adaptive Security (QI-CTAS) framework, a four-layer, closed-loop architecture that applies the mathematical formalism of quantum computing — amplitude-normalized feature states, entanglement-inspired correlation, and Grover-inspired amplitude amplification — entirely on classical hardware to address the detection and response limitations of signature-based and classical machine-learning approaches to IoT security.

Evaluated across 268,000 labelled flow records from three heterogeneous IoT deployments spanning a smart building, an industrial gateway cluster, and a consumer mesh network, QI-CTAS achieved a detection accuracy of 96.8 percent, an F1-score of 0.961, and a false-positive rate of 1.8 percent, outperforming a tuned Random Forest baseline by 7.5, 8.5, and 6.6 percentage points respectively. The mean interval between anomaly onset and the first adaptive policy action was 3.4 minutes, an 89 percent reduction relative to manual triage under the rule-based baseline. Ablation results identified the closed-loop continuous-learning mechanism as the single largest contributor to sustained performance ($\Delta F1 = -0.114$), with entanglement-inspired correlation ($\Delta F1 = -0.077$) and Grover-inspired amplitude amplification ($\Delta F1 = -0.059$) jointly accounting for the majority of the remaining accuracy gain over classical feature representations.

These findings suggest that quantum-inspired mathematical structure, applied without physical quantum hardware, can meaningfully improve both the sensitivity and the response latency of adaptive security systems for heterogeneous IoT environments. Future work will examine the framework's Energy Efficiency limitation identified in Section 6.3 through sparse approximations of the entanglement-inspired correlation operator for very densely connected device clusters, and will investigate whether reference-signature libraries trained at one deployment site transfer effectively to a previously unseen site without full local retraining.

Declarations

Acknowledgements

NA

Funding

NA

Contributions

VA; AK; NR; Conceptualization, VA; AK; NR; Investigation, VA; AK; NR; Writing (Original Draft), VA; AK; NR; and VA; AK; NR; Writing (Review and Editing) Supervision, VA; AK; NR; Project Administration.

Ethics declarations

This article does not contain any studies with human participants or animals performed by any of the authors.

Data Availability Statement

NA

Use of Generative Artificial Intelligence

The authors declare that no generative artificial intelligence tools were used in the preparation of this manuscript.

Competing interests

All authors declare no competing interests.

References

- [1] Al-Fuqaha, A., Guizani, M., Mohammadi, M., Aledhari, M., & Ayyash, M. (2015). Internet of Things: A Survey on Enabling Technologies, Protocols, and Applications. *IEEE Communications Surveys & Tutorials*, 17(4), 2347–2376. <https://doi.org/10.1109/comst.2015.2444095>
- [2] Hassija, V., Chamola, V., Saxena, V., Jain, D., Goyal, P., & Sikdar, B. (2019). A Survey on IoT Security: Application Areas, Security Threats, and Solution Architectures. *IEEE Access*, 7, 82721–82743. <https://doi.org/10.1109/access.2019.2924045>
- [3] Uthayakumar, C., Jayaraman, R., Raja, H. A., & Shabbir, N. (2025). QSEER-Quantum-Enhanced Secure and Energy-Efficient Routing Protocol for Wireless Sensor Networks (WSNs). *Sensors*, 25(18), 5924. <https://doi.org/10.3390/s25185924>

- [4] Vishwakarma, R., & Jain, A. K. (2019). A survey of DDoS attacking techniques and defence mechanisms in the IoT network. *Telecommunication Systems*, 73(1), 3–25. <https://doi.org/10.1007/s11235-019-00599-z>
- [5] Anthi, E., Williams, L., Slowinska, M., Theodorakopoulos, G., & Burnap, P. (2019). A Supervised Intrusion Detection System for Smart Home IoT Devices. *IEEE Internet of Things Journal*, 6(5), 9042–9053. <https://doi.org/10.1109/jiot.2019.2926365>
- [6] da Costa, K. A. P., Papa, J. P., Lisboa, C. O., Munoz, R., & de Albuquerque, V. H. C. (2019). Internet of Things: A survey on machine learning-based intrusion detection approaches. *Computer Networks*, 151, 147–157. <https://doi.org/10.1016/j.comnet.2019.01.023>
- [7] Tang, E. (2019). A quantum-inspired classical algorithm for recommendation systems. *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 217–228. <https://doi.org/10.1145/3313276.3316310>
- [8] Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J. C., Barends, R., ... & Martinis, J. M. (2019). Quantum supremacy using a programmable superconducting processor. *nature*, 574(7779), 505–510.
- [9] Sengupta, J., Ruj, S., & Das Bit, S. (2020). A Comprehensive Survey on Attacks, Security Issues and Blockchain Solutions for IoT and IIoT. *Journal of Network and Computer Applications*, 149, 102481. <https://doi.org/10.1016/j.jnca.2019.102481>
- [10] Nakip, M., & Gelenbe, E. (2021). MIRAI Botnet Attack Detection with Auto-Associative Dense Random Neural Network. 2021 IEEE Global Communications Conference (GLOBECOM), 01–06. <https://doi.org/10.1109/globecom46510.2021.9685306>
- [11] Sivanathan, A., Gharakheili, H. H., Loi, F., Radford, A., Wijenayake, C., Vishwanath, A., & Sivaraman, V. (2019). Classifying IoT Devices in Smart Environments Using Network Traffic Characteristics. *IEEE Transactions on Mobile Computing*, 18(8), 1745–1759. <https://doi.org/10.1109/tmc.2018.2866249>
- [12] Meidan, Y., Bohadana, M., Mathov, Y., Mirsky, Y., Shabtai, A., Breitenbacher, D., & Elovici, Y. (2018). N-BaIoT—Network-Based Detection of IoT Botnet Attacks Using Deep Autoencoders. *IEEE Pervasive Computing*, 17(3), 12–22. <https://doi.org/10.1109/mpvr.2018.03367731>
- [13] Smith, K., & Wilson, I. (2022). The Challenges of the Internet of Things Considering Industrial Control Systems. *Privacy, Security and Forensics in the Internet of Things (IoT)*, 77–94. https://doi.org/10.1007/978-3-030-91218-5_4
- [14] Alladi, T., Chamola, V., Sikdar, B., & Choo, K.-K. R. (2020). Consumer IoT: Security Vulnerability Case Studies and Solutions. *IEEE Consumer Electronics Magazine*, 9(2), 17–25. <https://doi.org/10.1109/mce.2019.2953740>
- [15] Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. *Nature*, 549(7671), 195–202. <https://doi.org/10.1038/nature23474>
- [16] Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., Cincio, L., & Coles, P. J. (2021). Variational quantum algorithms. *Nature Reviews Physics*, 3(9), 625–644. <https://doi.org/10.1038/s42254-021-00348-9>
- [17] Pilato, G., & Vella, F. (2022). A Survey on Quantum Computing for Recommendation Systems. *Information*, 14(1), 20. <https://doi.org/10.3390/info14010020>
- [18] Abreu, D., Rothenberg, C. E., & Abelém, A. (2024). QML-IDS: Quantum Machine Learning Intrusion Detection System. 2024 IEEE Symposium on Computers and Communications (ISCC), 1–6. <https://doi.org/10.1109/iscc61673.2024.10733655>
- [19] Nahar, N., Andersson, K., Schelén, O., & Saguna, S. (2024). A Survey on Zero Trust Architecture: Applications and Challenges of 6G Networks. *IEEE Access*, 12, 94753–94764. <https://doi.org/10.1109/access.2024.3425350>

- [20] Yan, Q., Yu, F. R., Gong, Q., & Li, J. (2016). Software-Defined Networking (SDN) and Distributed Denial of Service (DDoS) Attacks in Cloud Computing Environments: A Survey, Some Research Issues, and Challenges. *IEEE Communications Surveys & Tutorials*, 18(1), 602–622. <https://doi.org/10.1109/comst.2015.2487361>
- [21] Sarker, I. H., Kayes, A. S. M., Badsha, S., Alqahtani, H., Watters, P., & Ng, A. (2020). Cybersecurity data science: an overview from machine learning perspective. *Journal of Big Data*, 7(1). <https://doi.org/10.1186/s40537-020-00318-5>