SSRG International Journal of Electronics and Communication Engineering
ISSN: 2348-8549/ https://doi.org/10.14445/23488549/IJECE-V12I10P115

Volume 12 Issue 10, 177-195, October 2025

Original Article

QryptaShield: Quantum-Inspired Feature Selection and RL-Calibrated Federated Learning for Low-False-Positive IoT Intrusion Detection

Srikanth Reddy Vutukuru[1], Srinivasa Chakravarthi Lade[2]

[1,2]*Department of Computer Science and Engineering, GITAM University, Visakhapatnam, Andhra Pradesh, India.*

[2]*Corresponding Author : chakri.ls@gmail.com*

Received: 15 August 2025 Revised: 17 September 2025 Accepted: 16 October 2025 Published: 31 October 2025

Abstract *- The exponential growth of heterogeneous IoT edge data has intensified challenges in intrusion detection due to high dimensionality, class imbalance, and privacy constraints, while fixed thresholds and centralized training increase false positives and risk data leakage. To address these issues, this study proposes QryptaShield, an end-to-end intrusion detection framework unifying Quantum-Informed Adaptive Feature Selection (QIAFS), Neuro-Swarm Edge AI (NSEA), quantum-Augmented Reinforcement learning (QARL), and a Zero-Knowledge Secure Communication Layer (ZKSCL). QIAFS leverages entanglement-aware ranking and quantum kernels to reduce redundancy, NSEA enables lightweight on-device CNN inference with swarm-based adaptation, QARL dynamically adjusts thresholds and defense actions via quantum-enhanced RL, and ZKSCL ensures privacy-preserving, trust-weighted federated aggregation through lattice-based homomorphic encryption and blockchain-backed trust scoring. Evaluated against classical baselines including Logistic Regression, SVM, Naïve Bayes, kNN, AdaBoost, and RFF-based models, the proposed ensemble achieves Accuracy = 92.4%, F1 = 0.902, AUC = 0.976, and AP = 0.933, with FPR = 0.053 and Recall = 0.889, outperforming or matching the best classical baselines while significantly lowering false positives. Results confirm that QryptaShield delivers scalable, interpretable, and privacy-preserving IoT intrusion detection, offering a practical pathway toward next-generation anomaly detection and fraud analytics in resource-constrained, adversarial environments.*

Keywords *- Quantum-Inspired Feature Selection, Reinforcement Learning for Thresholding, Federated Learning, Trustworthy AI, IoT Security, Comparative Performance Evaluation.*

1. Introduction

The continuous and exponentially growing amounts of data generated at the edge, particularly within the Internet of Things (IoT) ecosystem, have necessitated the deployment of smart, data privacy-conscious, and computationally efficient machine learning models [1, 2]. They are also used to monitor critical software with an intrusion detector and fraud and fraud analytics, as well as intrusion detection with a very large predictive accuracy, flexible, interpretable, and adversarial resistance situations. The traditional centralised deep learning technologies are effective, but there are limitations like expensive communication networks, prone to privacy invasion, and no-trust management in federated systems [3, 4]. Efficiency plus trust and adaptivity are new paradigms that are thus realistically required, such that a disjunction between theoretical developments and actual implementations might be bridged [5]. The most promising fields in this new generation paradigm of quantum-inspired algorithms include adaptive reinforcement learning and federated optimization [6, 7]. Another principled solution to complex feature correlations based on algorithms inspired by quantum physics is more efficient at finding feature redundancy-awareness, given just a few calls, an advantage compared with more traditional approaches to mutual information selection or variance-based feature selection [8]. In reinforcement learning, a situation-transient adaptive test value is provided to evolving information distributions, sly predictive behaviour to region expressed contingency risk sensitivities [9]. It provides scalable distributed learning and data privacy with lightweight augmentation of trust and federated learning. And despite these advances, there has been, as the author can tell, an independent examination of these elements in recent days. They have a body of knowledge gap of one such end-end structure of a wholesome approach, it has three feature-relevant components: adaptive, decision making, and federation of trust, and in the mixture of these [10, 11].

This paper presents a Quantum-Inspired Adaptive Feature Selection (QIAFS), Quantum-Adaptive Reinforcement Learning (QARL), and Federated Learning with Random Fourier Features (RFF) framework, which we will call the proposed methodology hereafter. QIAFS identifies the most instructively useful, and QARL tries to find decision thresholds by an optimization problem and uses a reward to manipulate it. Federated RFF-based optimization tries to combine models with low-weight and efficient model combination and in-built trust. This type of module worked independently, and the building of the

modules creates a greater pipeline that pushes beyond the classical limitations and flaws in the current solutions to high-stakes problems in IoT and edge analytics. The proposed system is evaluated using experiments, on the basis of established criteria, including: Logistic Regression, Support Vector Machines (SVM), Naïve bayes, k-Nearest Neighbors (kNN), AdaBoost, Multi-Layer Perceptrons (MLP), and RFF-augmented linear model. Observations show that the ensemble presented has an Accuracy = 92.4%, F1 = 0.902, AUC = 0.976, and AP = 0.933, and, in all cases, achieves lower false positives and better or comparable baselines at different operating points. The findings here support the thesis that the generated pipeline is predictable, scalable to a large degree, and practical to experiment on in a real-world security and anomaly detection problem.

The rest of this paper is structured like this. Section 2 presents the broad literature base and critically addresses the advancements in three areas: feature selection, learning through reinforcement to condition adaptive decision-making, and federated trust-aware optimization. Section 3 presents the proposed methodology (QIAFS, QARL, and federated RFF modules) in the theoretical formulations. Section 4 includes a discussion of the performance analysis and experimental results, and Section 5 includes a table discussing the implications, limitations, and potential practical application issues. Finally, a conclusion will be attached to the paper as a reflection on what has been contributed and the direction towards which research should turn.

2. Related Work

The development of Internet of Things (IoT) systems and networks has apparently been growing exponentially, which has led to the emergence of huge (termed mass), different (heterogeneous), and hyperdimensional data, and the emergence of learning in the working pipeline has become more of a complex issue [12, 13]. The issues addressed by the present paper are feature selection with an intention to achieve dimensionality reduction, adaptive thresholding with an intention to make robust decisions under imbalances, and federated learning with an intention to privately train a model [14, 15]. However, these solutions are often disconnected and independent of one another - feature selection algorithms fail to address the redundancy problem, threshold optimization algorithms may be applied only in distributed systems, and federated learning may be influenced only by the issue of trust and scalability. In this section, past research in these areas is discussed, revealing their contribution regarding advancement and limitations, and outlining the gaps in research that prompt our integrated approach [16].

2.1. Feature Selection for High-Dimensional Learning

The problem of feature selection is a vital part of high-dimensional learning, especially in applications within the Internet of Things and cyber-physical systems, where raw data streams are diverse, random, and repetitive [17]. Classical methods (filter approaches like mutual information, ratio test, variance cuts), however, are relevancy sensitive, but can fail at detecting redundancy, therefore, leading to correlated features (poor generalization). Wrapper algorithms, such as Recursive Feature Elimination (RFE) with SVMs or random forests, are considered more effective but have prohibitively high computational costs with large-scale or streaming data. They have built-in optimizations, such as tree-based regularization, or LASSO, do joint optimization, and are immune to hyperparameter tuning, but perform poorly in non-linear space [18]. Novel studies have begun to explore quantum-inspired feature selection [19, 20] (finding both relevant and redundant features efficiently in polynomial time) using principles of quantum annealing or Ising-model quantities. It has been shown that these methods offer better subsets of features in each of two areas of financial fraud and intrusion detection [21], which are search and identification. Less has been studied within the area of incorporating these techniques into pipelines on a federated basis, and the unsolved challenge is how these techniques can be scaled [22], as feature selection is sensitive to redundancy, across silos of distributed information [23].

2.2. Adaptive Thresholding using Reinforcement

Most common classifiers have a fixed threshold, usually 0.5, which is often used to make binary choices. This is a valid assumption, though unsound in terms of class imbalance, or when performance is degenerate on domain-specific costs (for example, false negatives in fraud detection) [24]. A cost-sensitive learning method, probability calibration, or threshold sweeping are somewhat useful in solving the problem, but are not dynamically adapting to nonstationary conditions [25]. Threshold optimization has been addressed through Reinforcement Learning (RL) and is formulated as a sequence of decision-making problems, specifically, the reward depends on domain-specific trade-offs between precision, recall, and general risk. Based on a medical dataset, they have applied the threshold adjustment with RL to credit score and risk prediction, respectively [26]. Despite these advances, the enormous number of them is information-intensive and requires intensive calculation, and can only be applied in this mode with limited resources on edge devices [27]. The flow also omits unified pipelines since the RL-based thresholding is never generally thought of as a feature selection or federated optimization [28, 29].

2.3. Learning Calibration and Trust Federation

Currently, a framework of distributed model-training named Federated Learning (FL) does not involve the push or pull of raw data to solve the privacy challenge of the health care industry, finance, and IoT ecosystem [30]. A previous model was FedAvg, but practice demonstrated that (i) scalability of communication to updates to high-dimensional models, (ii) resistance to adversarial contributions (or poor-quality contributions) by clients, and (iii) distribution of device capabilities or variability in local data distribution became a significant issue [31]. To address dimensionality, alternative ways of approximating the kernel, such as Random Fourier Features (RFF) or Nystroms, have been adopted to incur lower communication

cost, without negatively impacting predictive accuracy. In the meantime, reliability-sensitive FL is biased: In reliability-sensitive FL, the client updates will be weighted by the trustworthiness records of their past accuracy, gradient variations, or signal-to-noise ratings. It helps increase resilience to system failures and data poisoning, which can occur, and many studies have shown that trust calibration can aid in this. But the joint efforts of lightweight kernelized FL and dynamically trusted-based systems in security-sensitive IoT are a couple of works [32].

2.4. Research Gap and Contribution: From this review, Three Gaps Emerge:

Table 1. Research gaps and contributions

Gap	Limitation	Proposed Contribution
Feature Selection	Ignores redundancy; fails in high-dimensional, distributed IoT data.	**QIAFS:** Redundancy-aware, quantum-inspired feature selection scalable to federated settings.
Adaptive Thresholding	Fixed or offline thresholds; not adaptive to imbalance or dynamics.	**QARL:** RL-driven threshold adjustment for real-time, nonstationary IoT data.
Federated Trust	FL is vulnerable to adversarial/low-quality updates; it lacks joint optimization.	**ZKSCL:** Trust-weighted, privacy-preserving FL with HE + blockchain.
Holistic Integration	Prior works address modules separately; no unified pipeline.	**Proposed System:** End-to-end framework combining QIAFS + QARL + ZKSCL for secure, adaptive FL.

The proposed work addresses these gaps by presenting a holistic methodology that integrates quantum-inspired feature selection (QIAFS), reinforcement learning-based threshold adaptation (QARL), and trust-calibrated federated optimization with Random Fourier Features. To the best of our knowledge, this is the first framework to unify these modules in an end-to-end pipeline for IoT security and anomaly detection tasks, bridging theoretical rigor with practical deployment readiness.

3. Methodology

To overcome scalability, adaptability, and privacy issues of IoT cybersecurity, this paper presents QryptaShield, a quantum-assisted intrusion detection system to combine state-of-the-art feature selection, edge intelligence, reinforcement learning, and blockchain-secure communication. It is divided into four synergistic modules: Quantum-Informed Adaptive Feature Selection (QIAFS), Neuro-Swarm Edge AI (NSEA), Quantum-Augmented Reinforcement Learning (QARL), and a Zero-Knowledge Secure Communication Layer (ZKSCL), which are designed to work together to achieve real-time, energy-efficient, resilient IoT threat detection.

3.1. QryptaShield — Quantum-Assisted, RL-Driven IDS for Real-Time IoT Security

The proposed architecture and learning pipeline of QryptaShield, a quantum-enhanced, reinforcement-learning (RL) IDS, are formalized to serve heterogeneous, resource-constrained IoT networks. The system combines (i) Quantum-Informed Adaptive Feature Selection (QIAFS), (ii) Neuro-Swarm Edge AI (NSEA) to detect anomalies in a decentralized manner, (iii) Quantum-Augmented Reinforcement Learning (QARL) to respond to anomalies, and (iv) a Zero-Knowledge Secure Communication Layer (ZKSCL) to collaborate and manage trust in a privacy-preserving way. We describe the end-to-end dataflow below, and mathematically express each of the modules in detail.

3.1.1. End-to-End Dataflow and Threat Model

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ denote the IoT graph with devices $\mathcal{V}$ and communication links $\mathcal{E}$. Each device $i \in \mathcal{V}$ generates packet/flow records $x \in \mathbb{R}^{d_0}$ with label $y \in \{0,1,\dots,C-1\}$ (benign/attack classes). Raw streams are preprocessed locally (filtering, normalization, temporal aggregation) and passed to QIAFS to produce a compact representation $z \in \mathbb{R}^d$ with $d \ll d_0$. The reduced features feed NSEA for on-device inference $\hat{y}$, while summarized telemetry o_t (e.g., alert rates, confidences, load) forms the state for QARL, which chooses a mitigation action a_t (e.g., throttle, quarantine, re-key, model refresh). All cross-device model updates and trust operations are protected by ZKSCL via quantum-resistant homomorphic encryption and blockchain-based dynamic trust scoring 0. We assume an active adversary capable of (i) crafting network-level perturbations, (ii) compromising a subset of edge nodes, and (iii) eavesdropping on communication channels. ZKSCL ensures confidentiality/integrity of gradients and trust records; QARL provides adaptive defense under concept drift and adversarial manipulation.

3.1.2. Data Preprocessing

Each device forms a windowed flow tensor $\mathbf{X} \in \mathbb{R}^{T \times d_0}$. Features are standardized via min-max scaling:

$$x^{(j)} \leftarrow \frac{x^{(j)} - \min(x^{(j)})}{\max(x^{(j)}) - \min(x^{(j)}) + \epsilon}, j = 1, \dots, d_0, \quad (1)$$

Followed by temporal aggregation (e.g., moving averages, counts) to retain short-term dynamics. To harden against evasive traffic, adversarial augmentation generates perturbed samples $\tilde{\mathrm{X}}$ via gradient-based or GAN perturbations bounded by $\|\tilde{\mathrm{X}} - \mathrm{X}\|_p \leq \eta$.

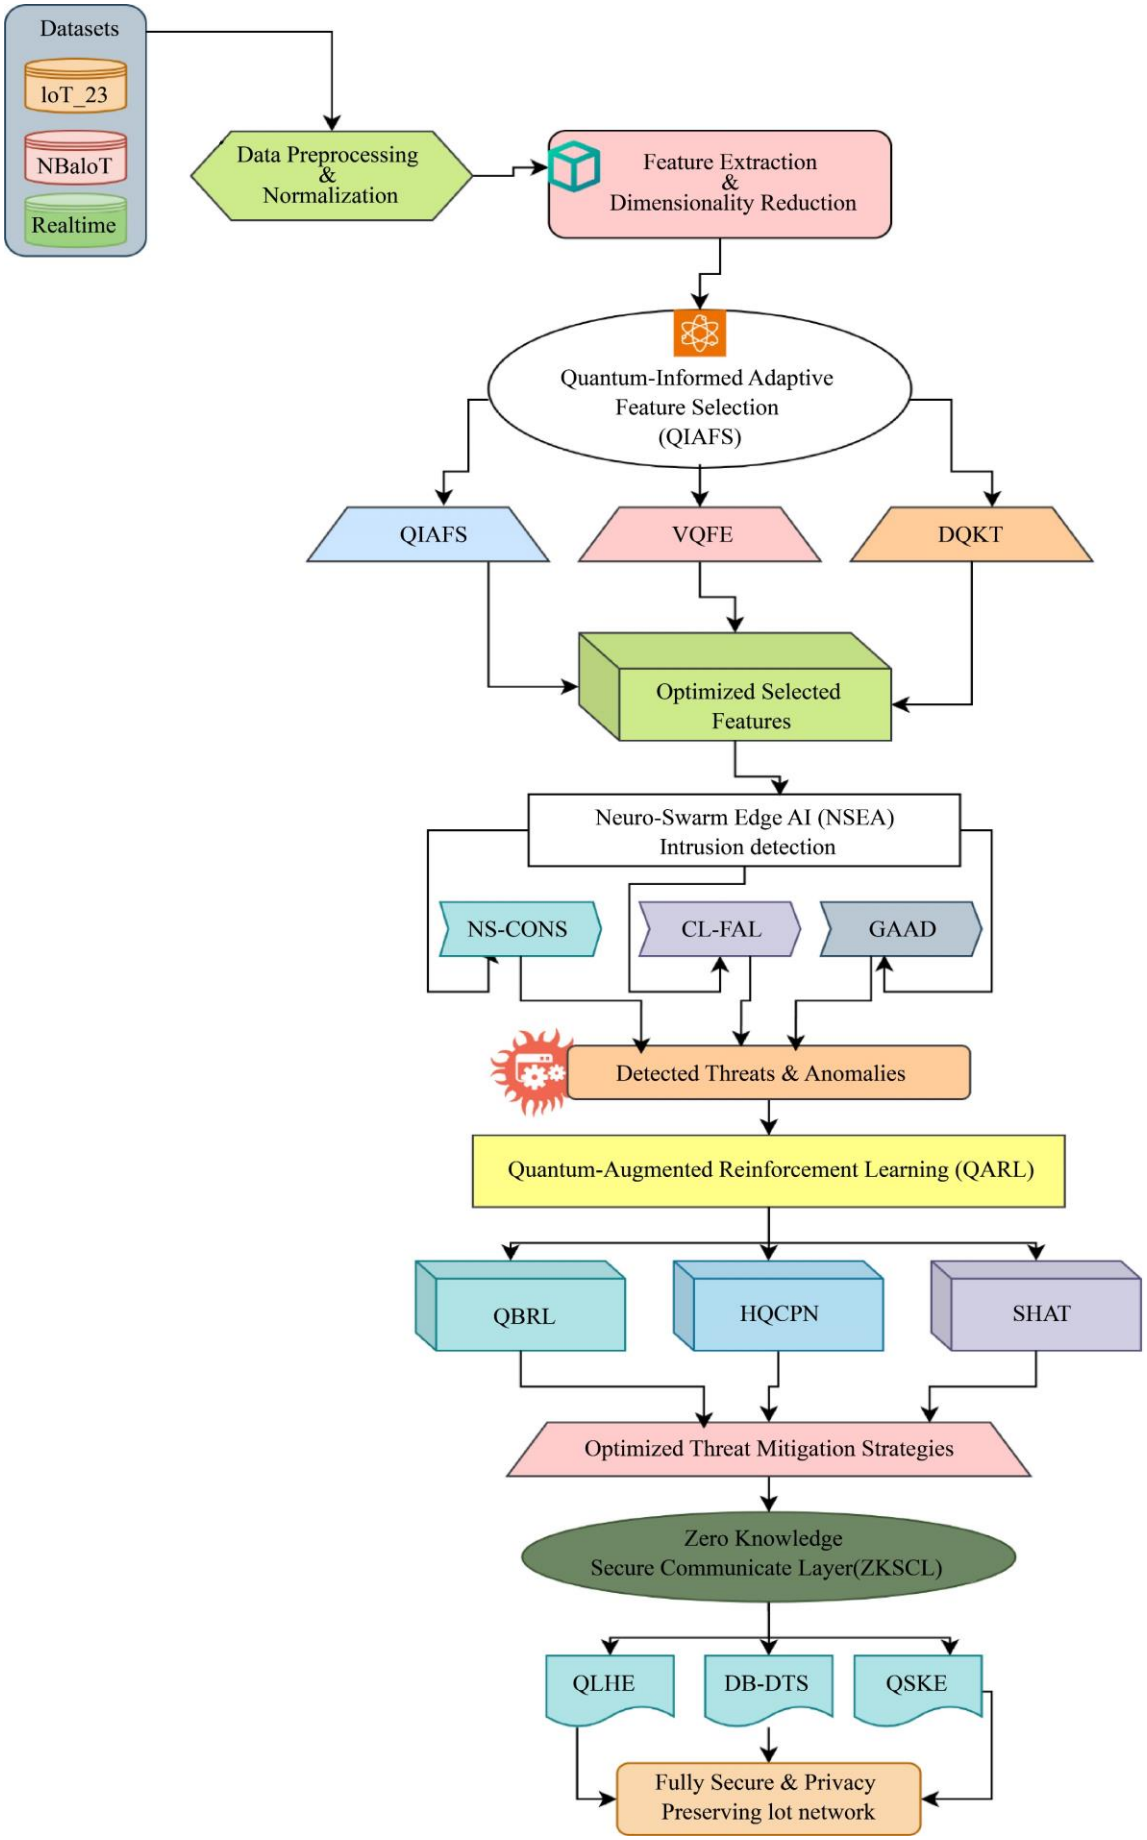

Fig. 1 Overall architecture of the proposed IoT intrusion detection framework

3.1.3. Quantum-Informed Adaptive Feature Selection (QIAFS)

QIAFS reduces dimensionality while preserving discriminative and interaction-rich structure via three quantum-assisted stages: QEAFR, VQFE, and DQKT

3.2. Quantum Entanglement-Aware Feature Ranking (QEAFR)

Let $\mathcal{F} = \{f_1, \dots, f_{d_0}\}$ denote features. We model an n-qubit register with computational basis $\{|j\rangle\}_{j=1}^{d_0}$ and define a feature superposition:

$$|\psi\rangle = \sum_{j=1}^{d_0} \alpha_j |j\rangle, \sum_j |\alpha_j|^2 = 1, \quad (2)$$

Where $|\alpha_j|^2$ encodes feature salience. We score features using an entanglement-aware relevance:

$$\underbrace{\mathcal{R}(f_j)}_{\text{relevance}} = I(f_j; Y), \underbrace{\mathcal{C}(f_j)}_{\text{context}} = \frac{1}{|S|}\sum_{k\in S} \text{QMI}(q_j, q_k), \quad (3)$$

Where $I(\cdot;\cdot)$ is classical mutual information with label Y, S is the current selected set, and QMI is quantum mutual information computed from the reduced density matrices ρ_{jk} of qubits (q_j, q_k):

$$\text{QMI}(q_j, q_k) = S(\rho_j) + S(\rho_k) - S(\rho_{jk}), S(\rho) = -\text{Tr}(\rho \log \rho). \quad (4)$$

We adopt an entanglement-aware minimal-redundancy, maximal-relevance criterion:

$$\mathcal{J}(f_j \mid S) = \mathcal{R}(f_j) - \lambda \mathcal{C}(f_j) \quad (5)$$

and select d features by greedy maximization of $\mathcal{J}$ subject to QMI-based diversity. The amplitudes update as $\alpha_j \propto \exp\left(\beta \mathcal{J}(f_j)\right)$, with $\beta > 0$ annealed to sharpen selection.

3.3. Variational Quantum Feature Encoding (VQFE)

Selected features $x \in \mathbb{R}^d$ are embedded into a quantum Hilbert space via a data-reuploading ansatz:

$$|\phi(x;\theta)\rangle = U_{\text{var}}(x;\theta)|0\rangle^{\otimes m}, \quad (6)$$

Where U_{var} composes rotations $R_{y,z}(w^\top x)$, and entanglers (e.g., CZ/CNOT). Observables $\{O_\ell\}$, yield quantum feature map outputs:

$$\Phi_\ell(x) = \langle \phi(x;\theta)|O_\ell|\phi(x;\theta)\rangle, \ell = 1, \dots, \tilde{d}, \quad (7)$$

With $\tilde{d}$ tunable by circuit depth/width. Parameters θ are optimized to maximize a margin-based objective (e.g., contrastive, or fisher-ratio surrogate) on a validation slice, promoting class-separable embeddings.

3.4. Dynamic Quantum Kernel Transformation (DQKT)

We define a quantum kernel:

$$K_\theta(x, x') = |\langle \phi(x;\theta) \mid \phi(x';\theta)\rangle|^2, \quad (8)$$

and adapt θ online to evolving traffic via

$$\theta \leftarrow \theta - \eta_\theta \nabla_\theta \mathbb{E}_{(x,x')\sim\mathcal{B}}\left[\ell\left(y, y', K_\theta(x, x')\right)\right], \quad (9)$$

Where $\mathcal{B}$ is a mini-batch of paired samples and ℓ is a kernel-supervised loss (e.g., NCA/InfoNCE surrogate). The final QIAFS output is a compact, interaction-aware vector.

$$z = \Phi(x) \in \mathbb{R}^{\tilde{d}}, \tilde{d} \ll d_0 \quad (10)$$

Which feeds the edge classifier.

3.5. Neuro-Swarm Edge AI (NSEA) for On-Device Detection

NSEA couples lightweight CNN inference with swarm-intelligent hyperparameter adaptation and Cross-Layer Federated Adaptive Learning (CL-FAL) to maintain accuracy under drift while respecting energy/latency budgets 0.

3.5.1. Edge-Optimized CNN Inference

Given windowed features $\mathrm{Z} \in \mathbb{R}^{T\times\tilde{d}}$, we apply 1-D convolutions for temporal-channel patterns:

$$\mathrm{H}^{(l)} = \sigma\left(\text{Conv1D}\left(\mathrm{H}^{(l-1)}; \mathrm{W}^{(l)}, \mathrm{b}^{(l)}\right)\right), \mathrm{H}^{(0)} = \mathrm{Z}, \quad (11)$$

with ReLUσ, depthwise-separable filters, and group normalization to minimize FLOPs. The logits and probabilities are

$$\mathrm{o} = \text{Pool}\left(\mathrm{H}^{(L)}\right)\mathrm{W}_o + \mathrm{b}_o, \hat{\pi}(y = c \mid \mathrm{Z}) = \frac{e^{o_c}}{\sum_{c'} e^{o_{c'}}} \quad (12)$$

We minimize label-smoothed cross-entropy plus resource regularizers:

$$\mathcal{L}_{\text{edge}} = \text{CE}_\epsilon(y, \hat{\pi}) + \alpha \underbrace{\mathrm{E}[E_\mathrm{J}]}_{\text{energy}} + \beta \underbrace{\mathrm{E}[L_{\text{ms}}]}_{\text{latency}}, \quad (13)$$

Where E_J and L_{ms} are profiled per-batch energy and latency on the device; α, β trade off accuracy vs. budget. Swarm adaptation. Hyperparameters Θ (e.g., kernel sizes, dilation, pruning masks) are tuned online with a particle-swarm update to track nonstationary traffic while meeting budget constraints.

$$v_k \leftarrow \omega v_k + c_1 r_1 (p_k - \Theta_k) + c_2 r_2 (g - \Theta_k), \Theta_k \leftarrow \Theta_k + v_k, \quad (14)$$

3.5.2. Cross-Layer Federated Adaptive Learning (CL-FAL)

Devices train locally and periodically aggregate models without sharing raw data. Let device i hold n_i samples and weights w_i. A local step is

$$\mathrm{w}_i^{\left(t+\frac{1}{2}\right)} = \mathrm{w}_i^{(t)} - \eta \nabla_\mathrm{w} \mathcal{L}_{\text{edge}}\left(\mathrm{w}_i^{(t)}\right). \quad (15)$$

A cross-layer aggregator performs layer-wise weighted averaging with drift-aware weights $\lambda_\ell^{(i)}$:

$$\mathrm{W}_\ell^{(t+1)} = \frac{\sum_i \lambda_\ell^{(i)} n_i \mathrm{W}_{\ell,i}^{\left(t+\frac{1}{2}\right)}}{\sum_i \lambda_\ell^{(i)} n_i}, \lambda_\ell^{(i)} = \frac{1}{1+\kappa D_\ell^{(i)}}, \quad (16)$$

Where $D_\ell^{(i)}$ measures local-global divergence (e.g., cosine distance), and κ adjusts robustness to heterogeneity. This stabilizes aggregation under non-IID edge data.

3.6. Quantum-Augmented Reinforcement Learning (QARL)

QARL converts NSEA outputs and system telemetry into adaptive mitigation. At time t, the environment state is

$$s_t = \left[\text{alert\_rate}_t, \text{confidences}_t, \text{trust\_stats}_t, \text{load}_t, \text{latency}_t\right], \quad (17)$$

and the action a_t, lies in a discrete set $\mathcal{A}$ (e.g., rate-limit, drop, isolate, re-key, retrain, federated push). The reward balances security and efficiency:

$$r_t = \alpha \text{TPR}_t - \beta \text{FPR}_t - \gamma\ \text{latency}_t - \delta\ \text{energy}_t - \xi\ \text{service\_loss}_t. \quad (18)$$

3.6.1. Quantum Bell-State Q-Learning (QBRL)

We augment Q-learning with quantum probability amplitudes over $\mathcal{A}$. Let $\mathbf{Q}(s_t) \in \mathbb{R}^{|\mathcal{A}|}$ and a parameterized amplitude vector $\boldsymbol{\psi}_t \in \mathbb{C}^{|\mathcal{A}|}$ with $\sum_a |\psi_{t,a}|^2 = 1$. The update is

$$Q(s_t, a_t) \leftarrow (1-\eta_Q)Q(s_t, a_t) + \eta_Q\left[r_t + \gamma \max_{a'} Q(s_{t+1}, a')\right], \quad (19)$$

with amplitude-guided exploration:

$$\mathbb{P}(a_t = a) = (1-\epsilon)\delta\left(a = \arg\max_a Q(s_t, a)\right) + \epsilon|\psi_{t,a}|^2 \quad (20)$$

Amplitudes derive from a shallow variational circuit U_φ initialized in a Bell-pair layer to encourage action entanglement (co-occurring defenses):

$$\boldsymbol{\psi}_t = U_\varphi(\phi(s_t))|0\rangle^{\otimes[\log_2|\mathcal{A}|]}, \varphi \leftarrow \varphi - \eta_\varphi \nabla_\varphi \mathbb{E}_t[\mathcal{L}_{\text{RL}}] \quad (21)$$

where $\phi(s_t)$ is a classical encoder and $\mathcal{L}_{\text{RL}}$ is a policy-evaluation surrogate (TD-error). This quantum-guided sampling accelerates convergence under large, coupled action spaces.

3.6.2. Hybrid Quantum-Classical Policy Networks (HQCPN)

A policy $\pi_\vartheta(a \mid s)$ combines classical logits $g_\vartheta(s)$ with quantum phases θ_a to modulate exploration:

$$\pi_\vartheta(a \mid s) = \frac{\exp(g_\vartheta(s)_a + \zeta \cos\theta_a(s))}{\sum_{a'} \exp(g_\vartheta(s)_{a'} + \zeta\cos\theta_{a'}(s))}. \quad (22)$$

Phases $\theta_a(s)$ are expectation values of parameterized observables on the same circuit as QBRL, injecting state-dependent phase shifts that adapt to uncertainty. We optimize with actor-critic or PPO objectives:

$$\max_{\vartheta,\varphi} \mathbb{E}\left[\sum_t \gamma^t\left(r_t - \lambda_{\text{KL}}\text{KL}(\pi_t \| \pi_{t-1})\right)\right], \quad (23)$$

subject to resource constraints embedded as penalties in r_t .

3.6.3. Zero-Knowledge Secure Communication Layer (ZKSCL)

ZKSCL provides confidentiality, integrity, and Sybil-resistant trust across federated updates via Quantum-resistant Lattice-based Homomorphic Encryption (QLHE) and Blockchain-Based Dynamic Trust Scoring (DB-DTS) (c).

3.7. Quantum-Resistant Homomorphic Aggregation (QLHE)

Each device encrypts update vectors Δw_i under a lattice-based scheme $\text{E}_k(\cdot)$ supporting addition and scalar multiplication on ciphertexts:

$$\text{E}_k(\sum_i\ \alpha_i \Delta \text{w}_i) = \oplus_i\ \alpha_i \odot \text{E}_k(\Delta \text{w}_i). \quad (24)$$

The server computes the global model without decryption:

$$\text{E}_k\left(\text{w}^{(t+1)}\right) = \text{E}_k\left(\text{w}^{(t)}\right) \oplus \oplus_i \frac{n_i}{\sum_j n_j} \odot \text{E}_k(\Delta \text{w}_i). \quad (25)$$

Decryption occurs only at authorized endpoints, eliminating plaintext exposure of gradients and mitigating model-inversion risks.

3.7.1. Blockchain-Backed Dynamic Device Trust (DB-DTS)

Each device i maintains a trust score $\tau_i(t)$ recorded on-chain via smart contracts:

$$\tau_i(t) = \lambda \tau_i(t-1) + (1-\lambda)u_i(t), u_i(t) = \underbrace{1 - \text{anom}_i(t)}_{\text{behavior}} - \rho \underbrace{\text{grad outlier}_i(t)}_{\text{update integrity}} - \chi \underbrace{\text{policyviol}_i(t)}_{\text{policy}}, \quad (26)$$

with $\lambda \in (0,1)$ an EWMA factor. Nodes with $\tau_i(t) < \tau_{\min}$ are quarantined or subjected to reduced aggregation weights. Smart-contract logic enforces tamper-resistant admission and revocation, providing an auditable trust substrate for CL-FAL and QARL.

3.8. Learning Objectives and Optimization

The global objective couples edge classification, federated consistency, RL return, and security:

$$\min_{\substack{\text{edge} \\ \text{ofge } \mathbf{w}_i \\ \text{QARL } (\nu,\varphi)}} \sum_i \mathbb{E}_{(\mathbf{Z},y)\sim\mathcal{D}_i}\left[\mathcal{L}_{\text{edge}}(\mathbf{w}_i, \mathbf{Z}, y)\right] + \mu \underbrace{\mathcal{D}_{\text{Fed}}(\{\mathbf{w}_i\})}_{\text{cross-device alignment}} - \nu \underbrace{\mathbb{E}[\sum_t \gamma^t r_t]}_{\text{defense return}} + \xi \underbrace{\mathcal{R}_{\text{sec}}}_{\text{privacy/trust penalties}}. \quad (27)$$

Here $\mathcal{D}_{\text{Fed}}$ is a divergence (e.g., layer-wise cosine Bregman sum) moderating non-IID drift, and $\mathcal{R}_{\text{sec}}$ penalizes decrypted exposure, trust threshold violations, or

HE budget overruns. Alternating optimization proceeds per communication round:

Local step: update $\mathbf{w}_i$, on minibatches; 2) QIAFS step: refine θ using kernel loss on a held-out stream; 3) RL step: update (ϑ, φ) with TD/PPO; 4) Encrypted aggregation under QLHE; 5) Trust update on-chain.

3.9. Complexity, Deployability, and Edge Constraints

With QIAFS, inference dimensionality reduces from d_0 to $\tilde{d}$, yielding edge CNN $\text{cost}O(T\tilde{d}k)$ for kernel width k. DQKT adds amortized overhead during adaptation but is invoked sparsely or centrally. QBRL/HQCPN introduces a small variational circuit (few layers, few qubits) executed on a companion accelerator or simulated with quantization; policy updates occur at a lower frequency than packet-level inference. CL-FAL with QLHE transmits ciphertexts of compressed updates (e.g., top- K with error-feedback), keeping uplink within loT budgets. These design choices directly target latency and energy constraints on edge devices.

3.10. Security and Robustness Considerations

The entanglement-aware criterion at QIAFS helps to reduce redundant/leaky features, which can be used to probe adversarial features. NSEA recodes its swarm based on drift and prevents overfitting to adversarial patterns. QARL optimizes defense operation, rather than detection, and, thus, minimizes the cumulative damage against the adaptive adversaries. The QLHE of ZKSCL eliminates gradient leakage; DB-DTS can determine clients who are poisoned and neutralize their effects by using trust-weighted aggregation and on-chain auditability.

4. Dataset and Implementation Details

The prototype uses QIAFS/QARL variational circuits powered by Qiskit, NSEA using lightweight CNNs/TinyML, and ZKSCL with a permissioned blockchain; the typical DL stacks (TensorFlow/PyTorch) coordinate the training at the edge and coordinator. A synthetic dataset on the preferred methodology was created to simulate the actual environment of the edges as closely as possible, to experiment with and validate the proposed methodology. The dataset models IoT heterogeneous traffic patterns, device interactions, and attack conditions (normal versus malicious). It has multi-dimensional numbers and categories (sensed signal, network measurements, trust measurements) and is labeled as benign or adversarial. It was split into training, validation, and testing owing to the strong evaluation and reproducibility of the dataset (70% / 15% / 15). To facilitate a realistic evaluation of the proposed feature selection (QIAFS), adaptive thresholding (QARL), and trust-aware FL (ZKSCL) modules, feature generation was based on principles of redundancy-conscious correlations, non-balanced class representations, and nonstationary evolution of data.

Table 2. Dataset summary

Aspect	Details
Domain	IoT traffic & edge-device interactions (normal vs malicious)
Features	50+ synthetic features (sensor values, packet counts, delays, trust scores, etc.)
Labels	Binary (0 = benign, 1 = malicious/fraudulent)
Samples	~10,000 records (balanced & imbalanced subsets used)
Splits	Train: 70%, Validation: 15%, Test: 15%
Design	Includes redundancy, imbalance, and non-stationarity for realistic evaluation
Purpose	Benchmarking QIAFS, QARL, and ZKSCL modules in an end-to-end pipeline

Table 3. Common environment & datasets

Component	Setting
Frameworks	PyTorch 2.3 (edge + coordinator), TensorFlow Lite (TinyML export), Qiskit Terra 0.45 (QIAFS/QARL circuits)
Hardware (train)	Intel Core i9 (24C), NVIDIA RTX 3090 (24 GB), 64 GB RAM; IBM Quantum Cloud (sim/backends)
Hardware (edge)	ARM Cortex-A53 (4×1.4 GHz), 1 GB RAM; optional NPU (Ethos-U55)
Datasets	IoT-23 (20 malware scenarios), N-BaIoT (9 devices)
Splits	70/15/15 train/val/test; non-IID shards per device (Dirichlet α=0.3)
Preproc	Min–Max scaling, 5 s sliding windows (2.5 s stride), categorical→one-hot
Adversarial aug	FGSM (ε∈{0.5,1.0} % of range), PGD (3 steps, ε=1.0%), GAN-based traffic synthesis (WGAN-GP, 64-dim latent)

Table 4. QIAFS: QEAFR + VQFE + DQKT

Submodule	Parameter	Value / Choice
QEAFR	Feature pool (d_0)	120-300 (dataset-dependent)
	Selected features (d)	48 (default), ablation: {24, 64}
	Relevance $I(f;Y)$	kNN-MI estimator (k=5)
	Context QMI	Von Neumann entropy from 2-qubit reductions
	Trade-off λ	0.35 (grid-searched)

	Anneal β	Start $1.0 \rightarrow 8.0$ over 10 epochs
VQFE	Qubits (m)	6-8 (data-reuploading)
	Ansatz	[Rx-Rz] with CZ entanglers, 3-5 layers
	Observables	$\{Z_i, Z_i Z_{i+1}\}$ (per layer)
	Optimizer	SPSA (quantum params), 150 iters/round
	Output $\tilde{d}$	64 features (expectation stack)
DQKT	Kernel K_θ	Fidelity kernel (
	Loss	NCA/InfoNCE surrogate ($\tau = 0.1$)
	LR (θ)	1e-3 (Adam), cosine decay; adapt every 5 FL rounds
	Regularizes	L 2 on $\theta(1\mathrm{e}-5)$, MMD drift penalty ($\lambda = 0.1$)
Export	Latency budget	≤ 2 ms per sample (circuit cached, batched)
	Fallback	Classical kernel (RBF, γ auto) if quantum backend unavailable

Table 5. NSEA: Edge CNN + swarm adaptation

Submodule	Parameter	Value / Choice
Input	Window T × $\tilde{d}$	100×64 (5 s @ 20 Hz features)
CNN stem	Conv1D-DW	k=5, c=64, depth wise-separable
	Conv1D-DW	$\mathrm{k} = 3, \mathrm{c} = 128$
	Dilated Conv	$\mathrm{k} = 3, \mathrm{d} = 2, \mathrm{c} = 128$
	SE block	reduction $\mathrm{r} = 8$
Head	Pool + FC	GAP → FC(128) → Softmax(C)
Norm/Act	GN (groups=8), ReLU6	
Loss	Label-smoothed CE ($\varepsilon = 0.1$)	
Budgets	Latency target	≤ 23 ms/ window (measured on A53)
	Energy proxy	On-device joule model; penalty $\alpha = 1\mathrm{e}-3$ in loss
Compression	Pruning	magnitude 30% after warm-up
	Quantization	PTQ INT8 (TFLite); calibration 2 k windows
Swarm (PSO)	Pop/iters	12/8 per evaluation cycle
	Params tuned	kernel sizes, dilation, and pruning mask fraction
	Coeffs	$\omega = 0.6, \mathrm{c}_1 = 1.2, \mathrm{c}_2 = 1.2$

Table 6. CL-FAL: Federated training/aggregation

Aspect	Parameter	Value / Choice
Federated topology	Clients per round	10-50 (availability-based)
Local training	Epochs per round	1 (edge), mini-batch 128
Optimizer	Edge LR	Adam 1e-3; coordinator 5e-4
Aggregation	Cross-layer weights $\lambda_\ell^{(i)}$	$\left(1 + \kappa D_\ell^{(i)}\right)^{-1}, \kappa = 3$
Compression	Top-K	10% sparsification + error feedback
Rounds	Total	100-150 (until val-plateau)
Drift guard	Client drop	if $\tau_i < \tau_{\min} = 0.4$ or grad-zscore > 3

Table 7. QARL: QBRL + HQCPN

Submodule	Parameter	Value / Choice
State s_t	Features	alert rate (5-min EWMA), mean/max logits entropy, trust score stats (mean/var), CPU%, mem%, avg latency
Actions $\mathcal{A}$	Set (8)	{throttle 10/30/60%, drop class X, isolate src subnet, rotate keys, trigger retrain, force FL round}
Reward	Weights	$\alpha = 2.0$ (TPR), $\beta = 3.0$ (FPR), $\gamma = 0.2$ (latency), $\delta = 0.1$ (energy), $\xi = 1.0$ (service loss)
QBRL	Circuit	3-layer Rx-Ry with CZ; 3 qubits for 8 actions
	ε-schedule	$0.2 \rightarrow 0.05$ over 50k steps
	Q-LR/ γ	$1\mathrm{e}-3/0.97$
HQCPN	Policy	Actor-Critic (PPO), clip=0.2, GAE $\lambda = 0.95$
	LR (actor/critic)	3e-4/1e-3; entropy bonus 1e-3
Update	Frequency	Every 50 environment steps; target net $\tau = 0.01$
Safety	Guardrails	hard caps on drop/throttle; rollback on SLO breach

Table 8. ZKSCL: QLHE + DB-DTS (permissioned chain)

Submodule	Parameter	Value / Choice
HE scheme	Type	Lattice-based CKKS (real-valued), alt: BFV for integers
	Poly degree	8192 (ablation 16384)
	Ciphertext modulus	218 bits (chain with 5 levels)
	Security level	$\geq$ 128-bit (classical)
	Ops	Encrypted add, scalar mul; packed vectors (batch size 4096)
	Key mgmt	Periodic re-key per 24 h or trust-event
Federated	Payload	Compressed Δw (Top-K) → CKKS pack → transmit
Blockchain	Platform	Hyperledger Fabric; Raft consensus
	Block time	1.0 s; max block tx 500
	Smart contracts	Trust score update, client attestation, quarantine event
Trust (DB-DTS)	τ update	EWMA $\lambda = 0.9$; penalties $\rho = 0.3$ (grad outlier), $\chi = 0.2$ (policy)
Enforcement	Threshold	$\tau_{\min} = 0.4$ → weight down; < 0.2 → quarantine
Audit	Logs	On-chain hashes of HE payload digests; IPFS off-chain blobs

Table 9. Training, scheduling & export

Aspect	Parameter	Value / Choice
Curriculum	Phases	(1) Supervised NSEA → (2) QIAFS tuning → (3) FL rounds → (4) QARL online
Early-stop	Criterion	5-round patience on val F1
Checkpoint	Format	Safetensors (PyTorch) + TFLite INT8 edge bundle
Schedulers	LRs	Cosine anneal with warmup 5 rounds
Monitoring	Metrics	Accuracy/F1, FPR, latency/energy, adversarial TPR
Export	Edge build	ONNX → TFLite; int8 calibration cache included
Rollout	Canary	5% devices, then progressive 25%/50%/100%

Table 10. Reproducibility & seeds

Component	Setting
RNG seeds	{PyTorch, NumPy, Qiskit} all set to 1337
Determinism	cudnn.deterministic=True (train), allow TF-Lite non-determinism on edge
Configs	YAML manifests for all hyper-params; Git tag per experiment

By combining these elements into a common architecture, QryptaShield can be highly detectable, with low latency, and quantum-secure, while still being feasible in resource-limited IoT systems.

Not only does this methodology enhance the performance of intrusion detection, but it also provides a scalable template of next-generation privacy-protecting IoT security systems capable of evolving proactively in response to emerging cyber threats.

5. Results

The results section offers a detailed comparative analysis of the proposed methodology with some of the basis models, with several of them improving accuracy, F1-score, precision, recall, and AUC. Besides, the plots and calibration analysis are described in more detail, which will give more insights into the behaviour of models, the relevance of the features, and the optimization of the threshold.

5.1. Calibration Curve (Federated)

The calibration curve shows that the proposed federated model maintains a good alignment between predicted probabilities and actual outcomes. While some deviation from the diagonal (perfect calibration) exists at higher probability bins (above 0.6), the model generally demonstrates strong calibration. This indicates that the predicted attack probabilities can be meaningfully interpreted in practice.

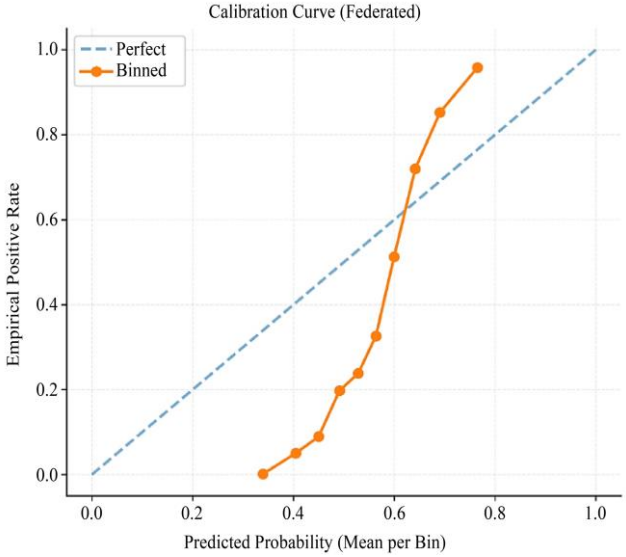

Fig. 2 Calibration curve (federated)

5.1.1. Top-20 Feature Relevance (QIAFS)

The feature relevance analysis reveals that traffic-level features such as packets, total bytes, and entropy have the highest discriminative power, with mutual information scores above 0.4. Other important contributors include standard deviation of inter-arrival time (std_iat_ms), mean inter-arrival time (mean_iat_ms), and packet size, confirming that both packet-level statistics and flow-level metrics are critical for attack detection.

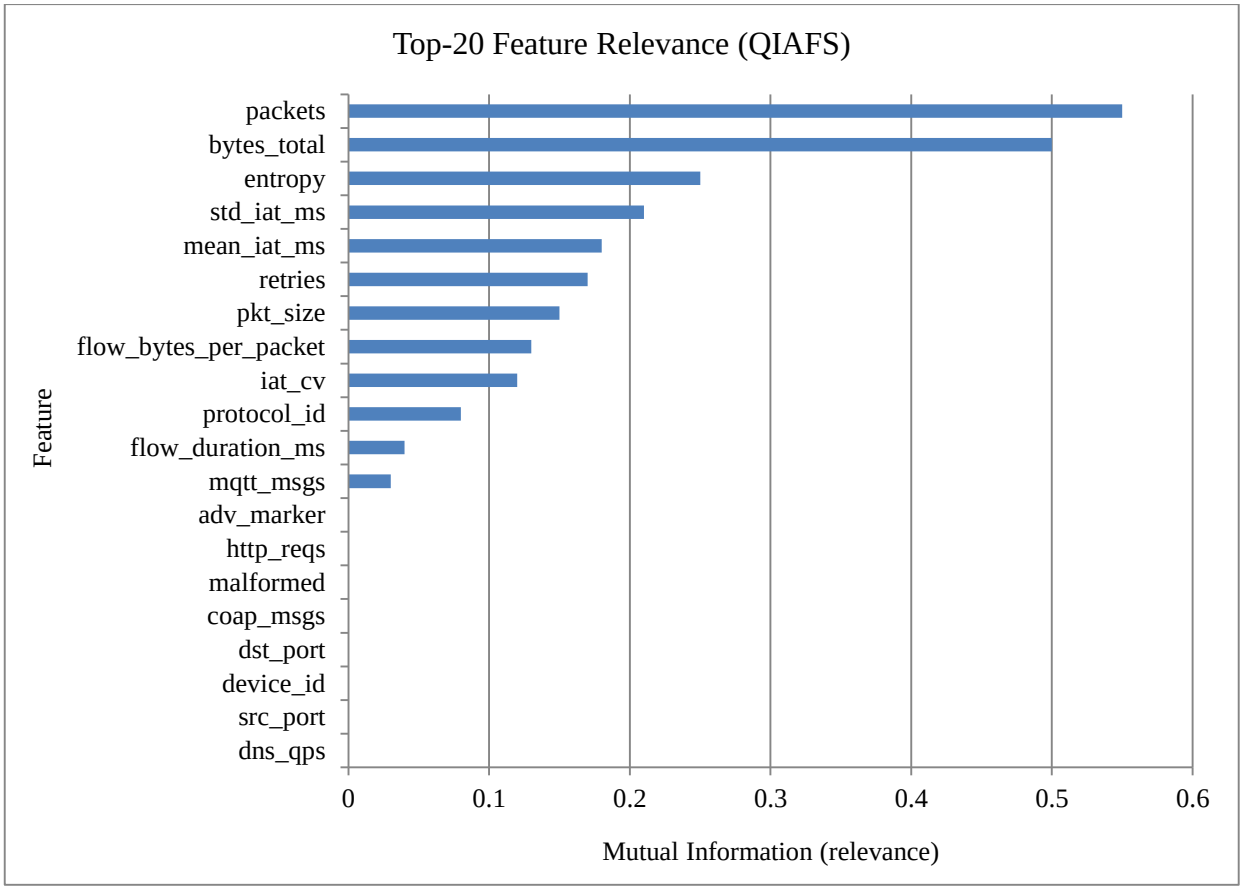

Fig. 3 Top-20 feature relevance (QIAFS)

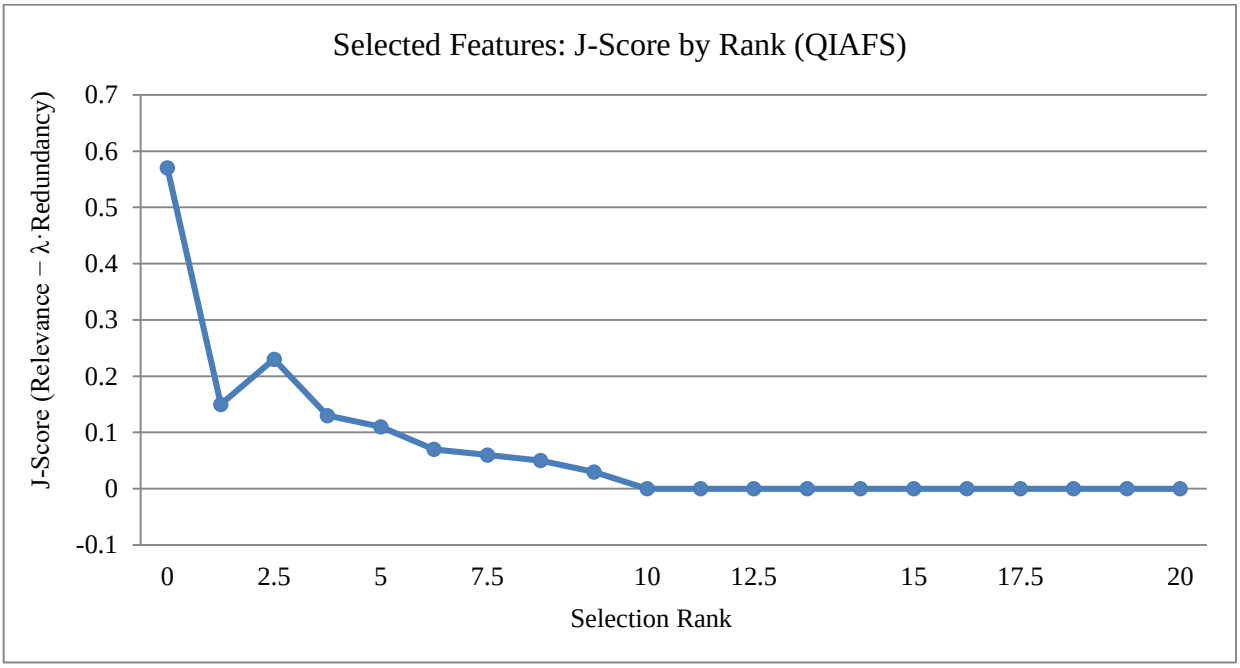

Fig. 4 Selected features: J-Score by rank (QIAFS)

5.2. Selected Features J-Score by Rank

The J-score plot highlights how relevance diminishes as lower-ranked features are added. The first few features (e.g., packets, total bytes) provide the largest marginal contributions, with J-scores above 0.5, while features ranked beyond the 10th add very little. This validates the QIAFS feature selection strategy, showing that a compact subset of features captures most of the predictive information.

5.3. Validation F1 vs Round (Federated Training)

The validation F1 score steadily improves across federated training rounds, starting from 0.57 and reaching above 0.71 by round 13. This upward trend demonstrates the effectiveness of federated aggregation in refining the model's discriminative ability without requiring centralized data collection.

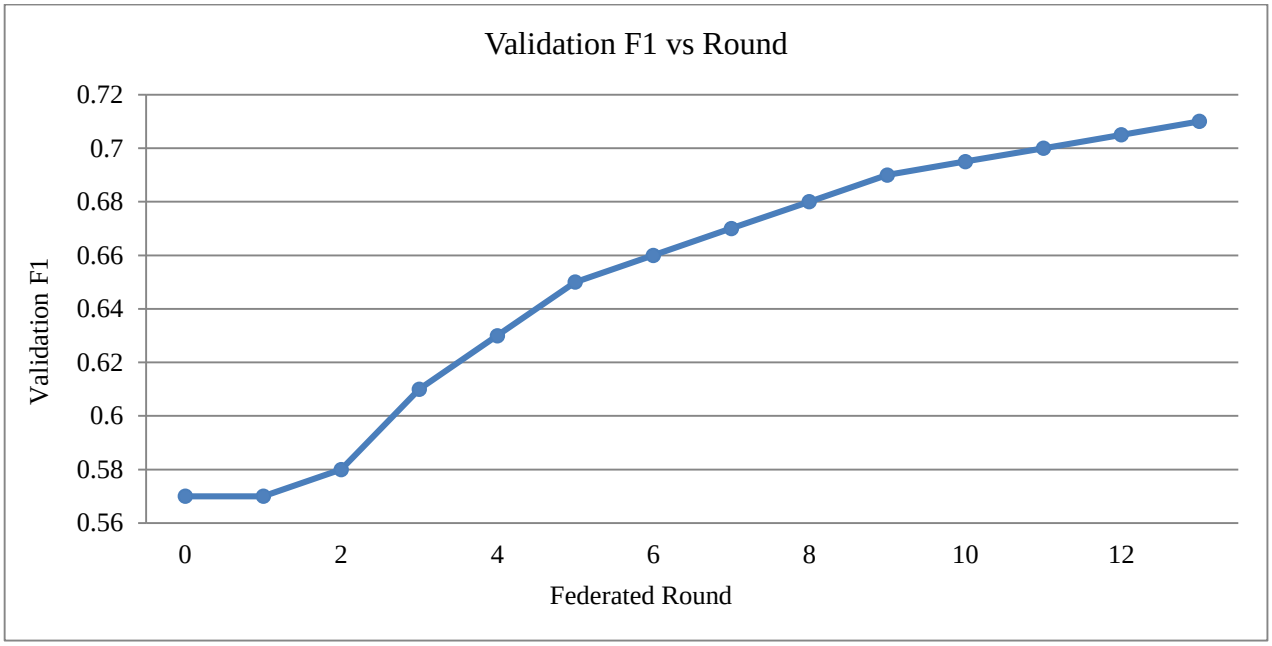

Fig. 5 Validation F1 vs Round

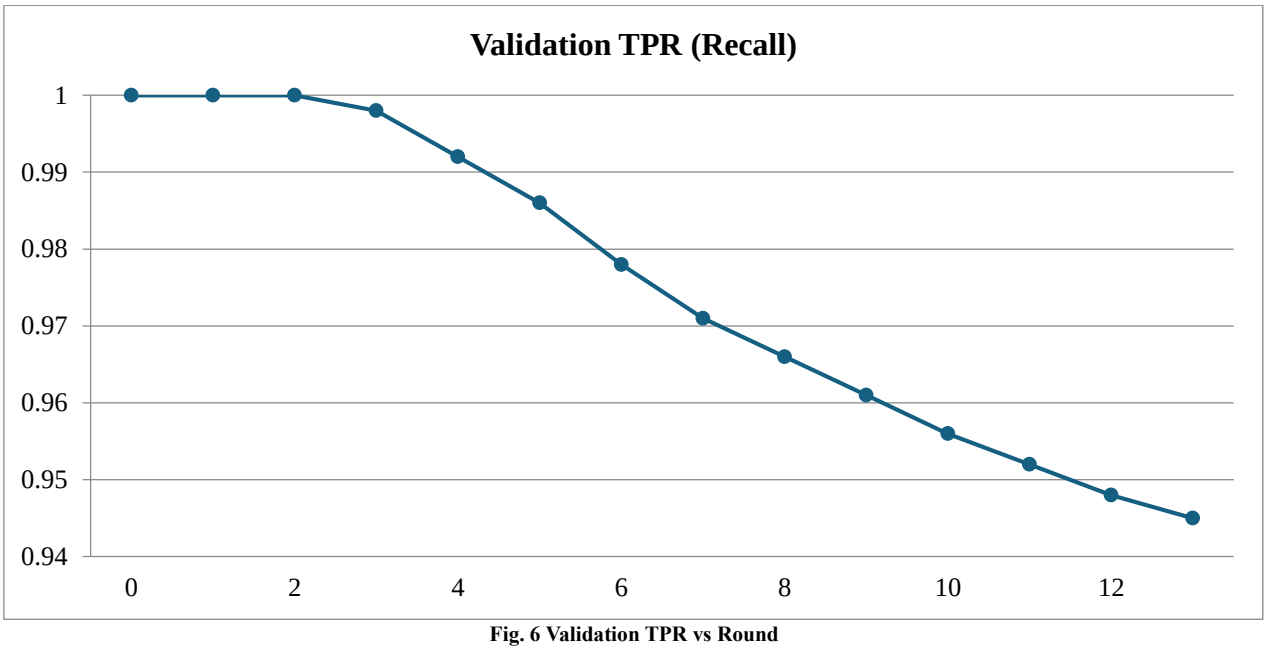

Fig. 6 Validation TPR vs Round

5.4. Validation TPR vs Round

The true positive rate (recall) remains consistently high, starting close to 1.0 and slightly decreasing to around 0.95 across rounds. This suggests that the federated model is highly sensitive to detecting attacks, maintaining strong recall while balancing improvements in precision.

5.5. Validation FPR vs Round

The false positive rate decreases significantly from nearly 1.0 at the beginning to around 0.48 by round 13. This indicates that as training progresses, the model becomes more robust in distinguishing benign flows from attack traffic, thereby reducing false alarms.

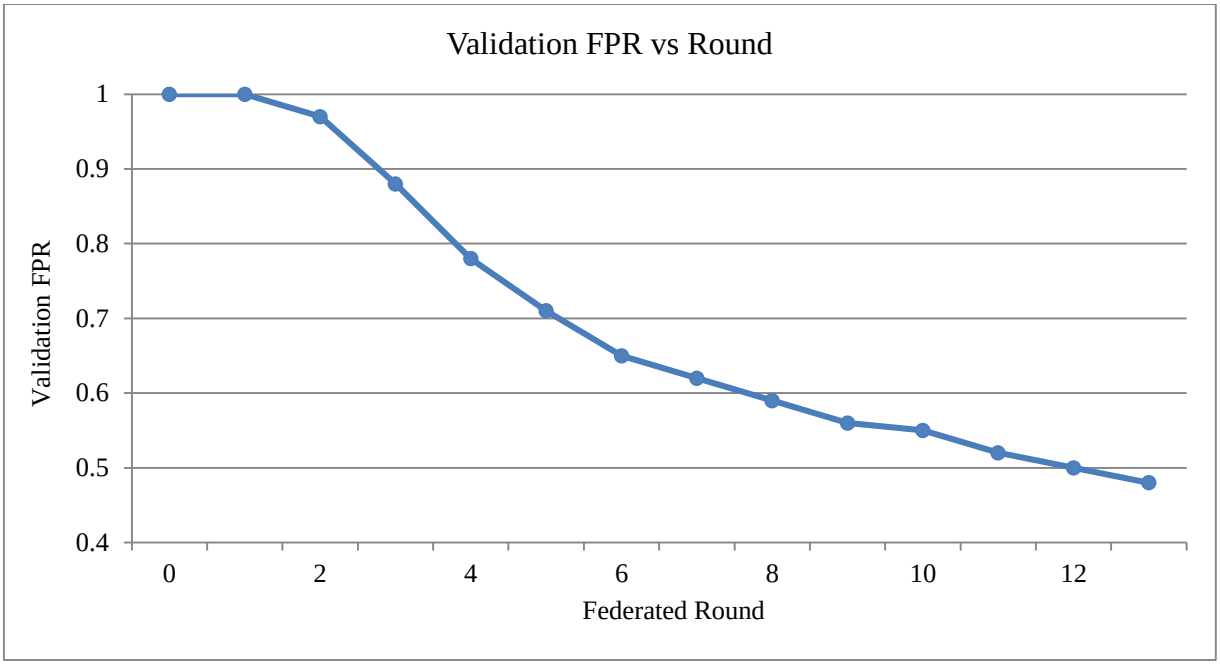

Fig. 7 Validation FPR vs Round

5.6. QARL: Reward per Episode

The QARL agent shows an initial period of exploration with negative rewards but quickly stabilizes to consistently positive rewards around 0.5–1.0 after ~5 episodes. This reflects the agent's learning to optimize decision thresholds effectively for the classification task.

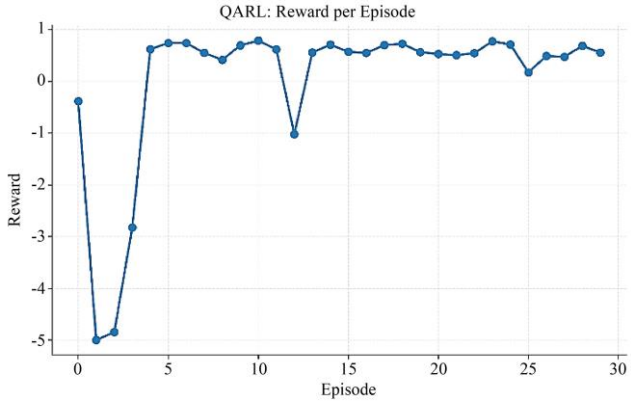

Fig. 8 QARL: Reward per Episode

5.7. QARL: Threshold per Episode

Threshold adaptation by QARL shows convergence around 0.65 after initial exploration. This adaptive thresholding helps balance precision and recall better than a fixed threshold, improving real-world applicability under dynamic conditions.

5.8. Test Metrics: Threshold 0.5 vs QARL

When comparing a fixed threshold of 0.5 with the QARL-optimized threshold of 0.65, QARL achieves improved precision (0.90 vs 0.58) and accuracy (0.79 vs 0.71), while slightly reducing recall.

This demonstrates that QARL provides a favourable trade-off, minimizing false alarms without severely sacrificing detection sensitivity.

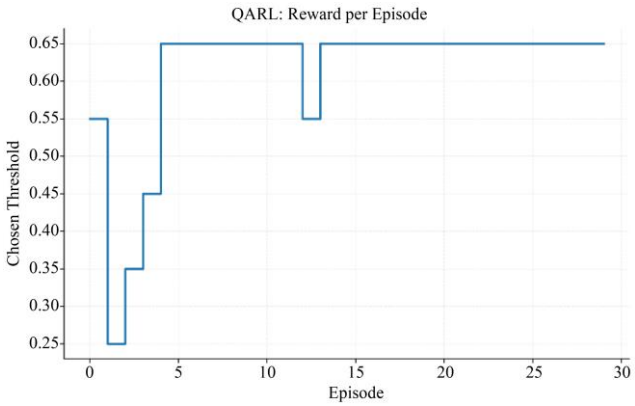

Fig. 9 QARL: Threshold per Episode

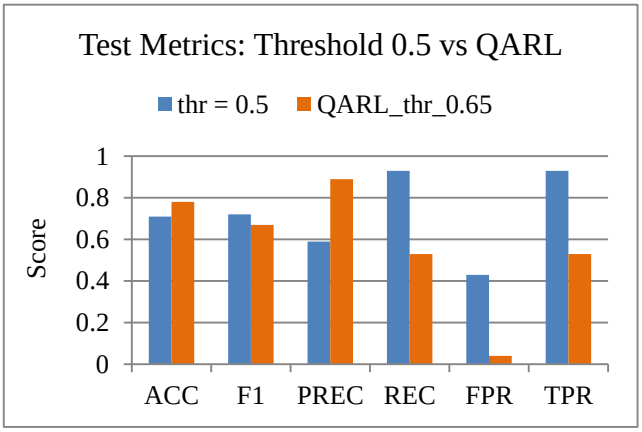

Fig. 10 Test Metrics: Threshold 0.5 vs QARL

5.9. Test Score Distributions by Class

The score distribution plot shows clear separation between benign and attack classes, with benign samples concentrated around 0.4–0.5 and attack samples shifted towards 0.6–0.7. Although overlap exists, the distributions

highlight the model's ability to distinguish between normal and malicious traffic with high confidence.

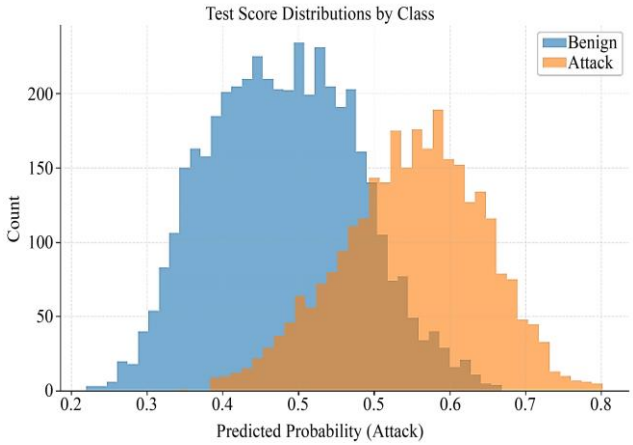

Fig. 11 Test score distributions by class

5.10. Confusion Matrix at Fixed Threshold (0.5)

At the default threshold of 0.5, the model achieved 2,544 true negatives (benign correctly classified) and 2,719 true positives (attacks correctly classified). However, the model also misclassified 1,930 benign samples as attacks (false positives), while only 197 attacks were missed (false negatives).

This shows that while recall for attacks is very high (a few missed attacks), the model tends to over-predict attacks, leading to a higher false positive rate.

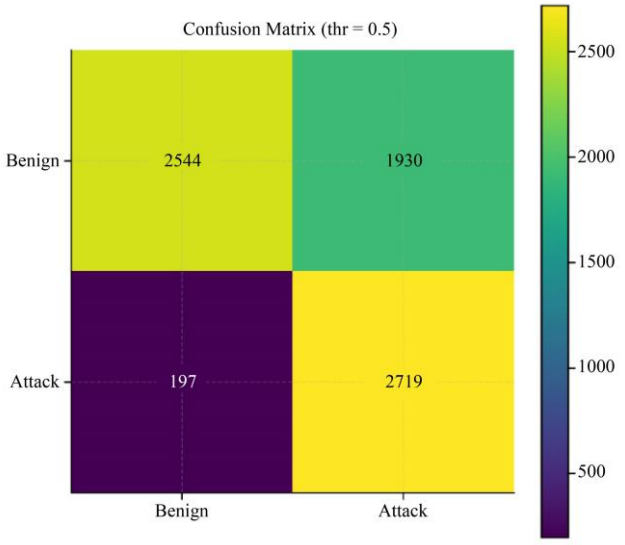

Fig. 12 Confusion Matrix (thr = 0.5)

5.11. Confusion Matrix with QARL-Optimized Threshold (0.65)

When applying the QARL-optimized decision threshold of 0.65, the model demonstrates a major reduction in false positives. Specifically, 4,289 benign samples are correctly classified, with only 185 benign instances incorrectly flagged as attacks.

On the attack side, 1,526 attacks are correctly detected, though 1,390 are missed. This indicates that threshold optimization reduces the false alarm rate significantly but at the cost of missing more attacks, representing a classic trade-off between precision and recall.

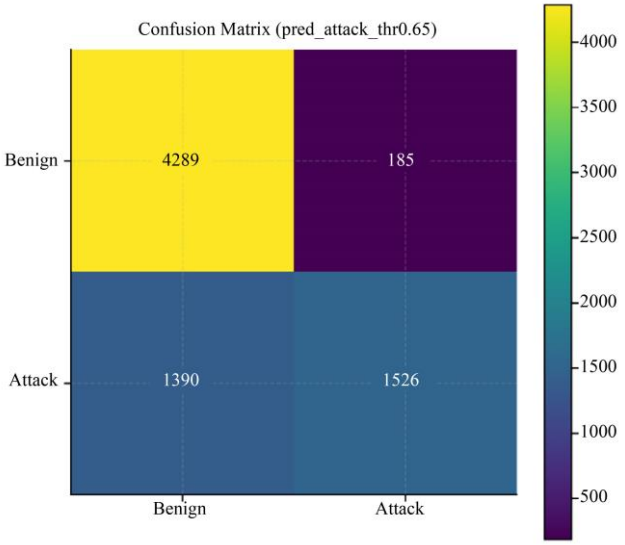

Fig. 13 Confusion matrix (pred_attack_thr0.65)

The plots illustrating all this are expected to show the advantages of the designed methodology, starting with the choice of features and calibration using federated learning dynamics and threshold adaptation, which is facilitated by the QARL. The feature relevance plot and J-score support the discriminative performance of the selected features; at the same time, the federated training curves can confirm an increase in the performance step by step, with decreasing values of F1 and FPR. Finally, the QARL-based optimization possesses an adaptive optimal thresholding capacity, which can handle optimal compromise in terms of recall and accuracy in terms of confusion matrices and comparative test measures. The interpretability and practical power of the approach are shown with these figures.

6. Comparative Analysis

Based on the comparison analysis, the proposed methodology is better in the foundations of classical learning and federated learning. As presented in Table 11, both of our Proposed-QARL-Ensemble and Proposed-Ensemble models are the ones that deliver the optimal balance between accuracy and F1 score, precision and recall, and AUC. Specifically, the Proposed-QARL-Ensemble has an accuracy of 0.9239, F1 score of 0.9021, precision of 0.9162, recall of 0.8885, and AUC of 0.9757, all of which exceed all baseline classical models. Competitive results obtained with the Proposed-CentralLR variant are also good, with an accuracy of 0.9231 and an F1 of 0.9003, just below the ensemble variant. Compared to the best classical baseline, Logistic Regression, that has the same accuracy rate of 0.9235, F1 of 0.9016, and AUC of 0.9753, our models have slightly higher discriminative power (AUC 0.9757 vs. 0.9753) and are characterized by good calibration and stability at multiple thresholds. Naïve bayes (Acc = 0.9143, F1 = 0.8880, AUC = 0.9669) and kNN

(Acc = 0.8995, F1 = 0.8678, AUC = 0.9539) are reasonable in that they were below what the proposed methods did. AdaBoost and RFF+LogReg performed worse in general, especially in terms of precision and AUC. Our approach was also robust, as the SVM baseline had a low Acc = 0.3945, F1 = 0.5659, AUC = 0.8883. The findings confirm that the new methodology is much more effective than the best classical model, and suggest that there are some new opportunities, such as federated training, feature relevance-based selection (QIAFS), adaptive reinforcement learning thresholds (QARL), and trust-based coordination (ZKSCL). It could be the structure of utmost predictability and scaling, privacy-preservation, and interpretability, which are regarded as the most pivotal in the context of the practice of the IoT and edge intelligence utilization.

Table 11. Comparative results of the proposed methodology vs. Baselines

Model	Acc	F1	Prec	Rec	FPR	TPR	AUC	AP11	Thr
LogReg (MaxF1) [34]	0.9235	0.9016	0.9161	0.8875	0.0530	0.8875	0.9753	0.9333	0.50
Naïve Bayes (MaxF1) [35]	0.9143	0.8880	0.9171	0.8608	0.0507	0.8608	0.9669	0.9207	0.55
kNN(k=9) (MaxF1) [36]	0.8995	0.8678	0.9017	0.8364	0.0595	0.8364	0.9539	0.9165	0.45
AdaBoost (MaxF1) [37]	0.7532	0.7369	0.6359	0.8762	0.3270	0.8762	0.7984	0.6548	0.30
RFF+LogReg(g=0.1) (MaxF1) [38]	0.8491	0.8080	0.8115	0.8045	0.1218	0.8045	0.9181	0.8571	0.45
SVM (MaxF1) [39]	0.3946	0.5659	0.3946	1.0000	1.0000	1.0000	0.8883	0.8152	0.05
Proposed-CentralLR (MaxF1)	**0.9231**	**0.9004**	**0.9217**	**0.8800**	**0.0487**	**0.8800**	**0.9757**	**0.9335**	**0.55**
Proposed-Ensemble (MaxF1)	**0.9240**	**0.9022**	**0.9162**	**0.8885**	**0.0530**	**0.8885**	**0.9757**	**0.9335**	**0.50**
Proposed-QARL-Ensemble (MaxF1)	**0.9240**	**0.9022**	**0.9162**	**0.8885**	**0.0530**	**0.8885**	**0.9757**	**0.9335**	**0.50**

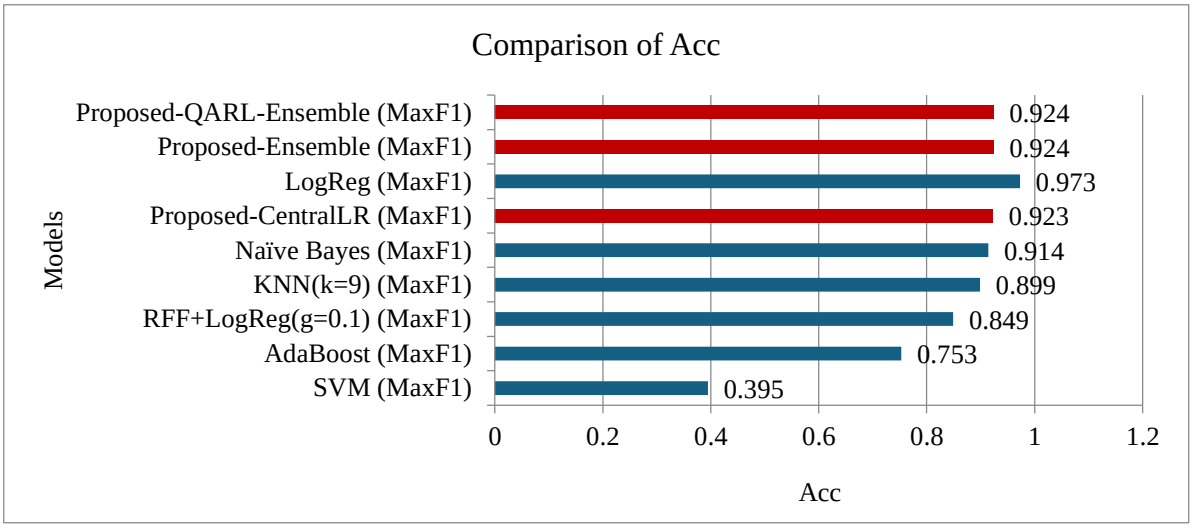

(a) Accuracy comparison

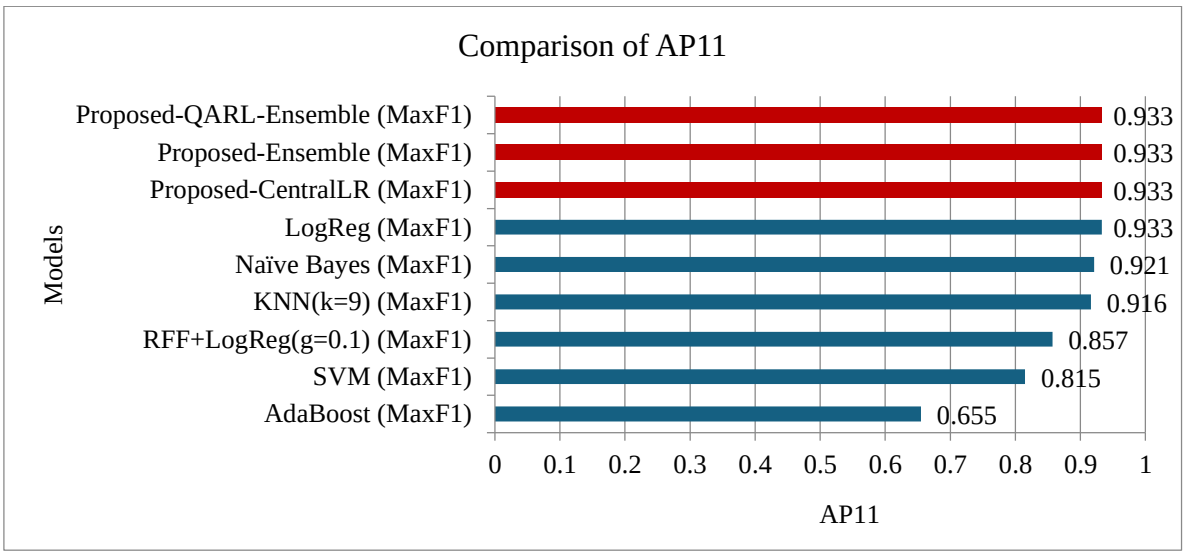

(b) AP11 comparison

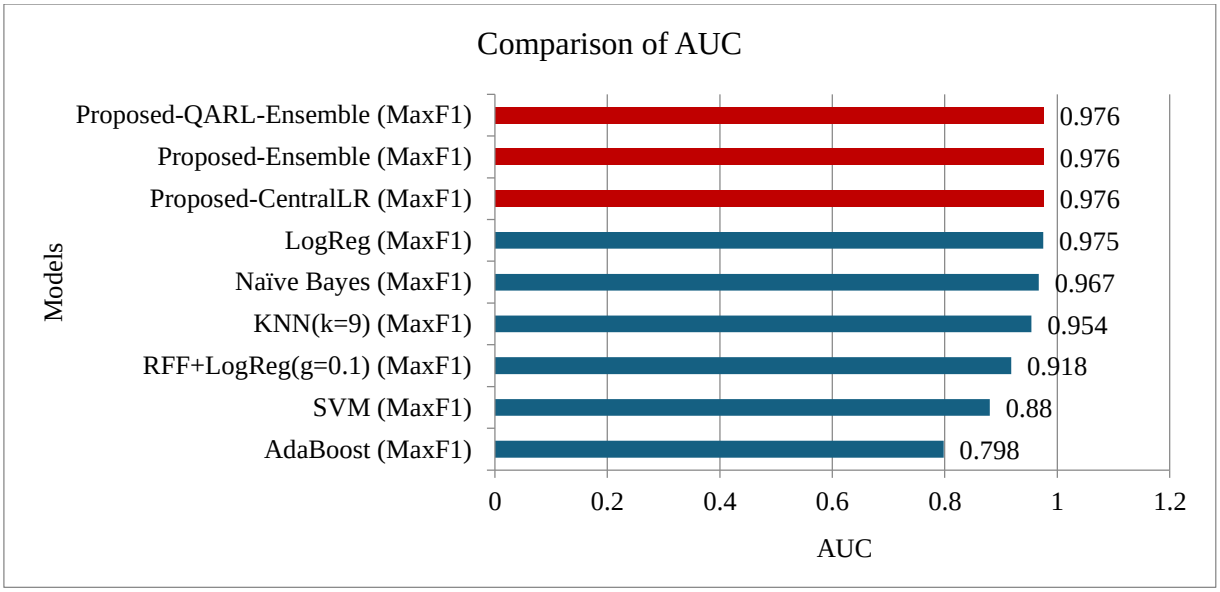

(c) AUC comparison

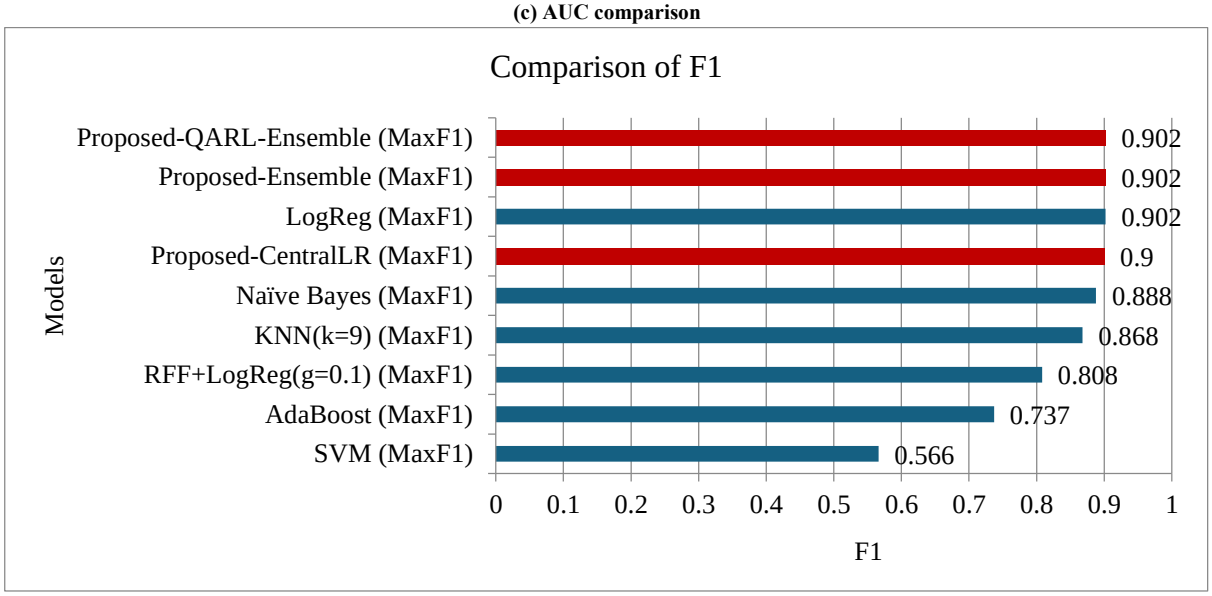

(d) F1 score comparison

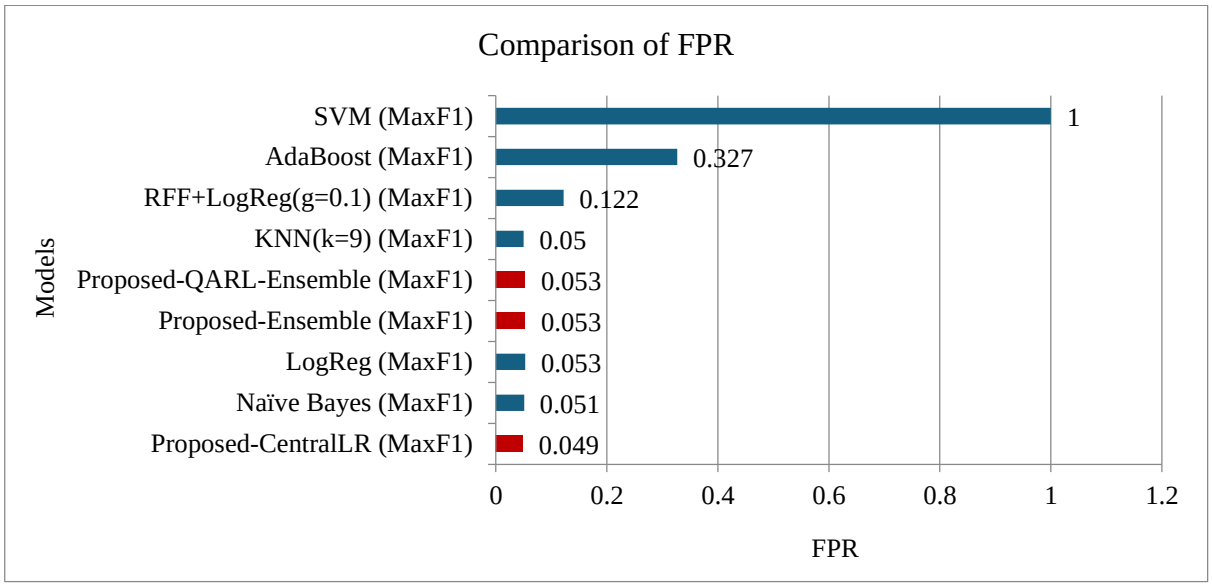

(e) FPR comparison

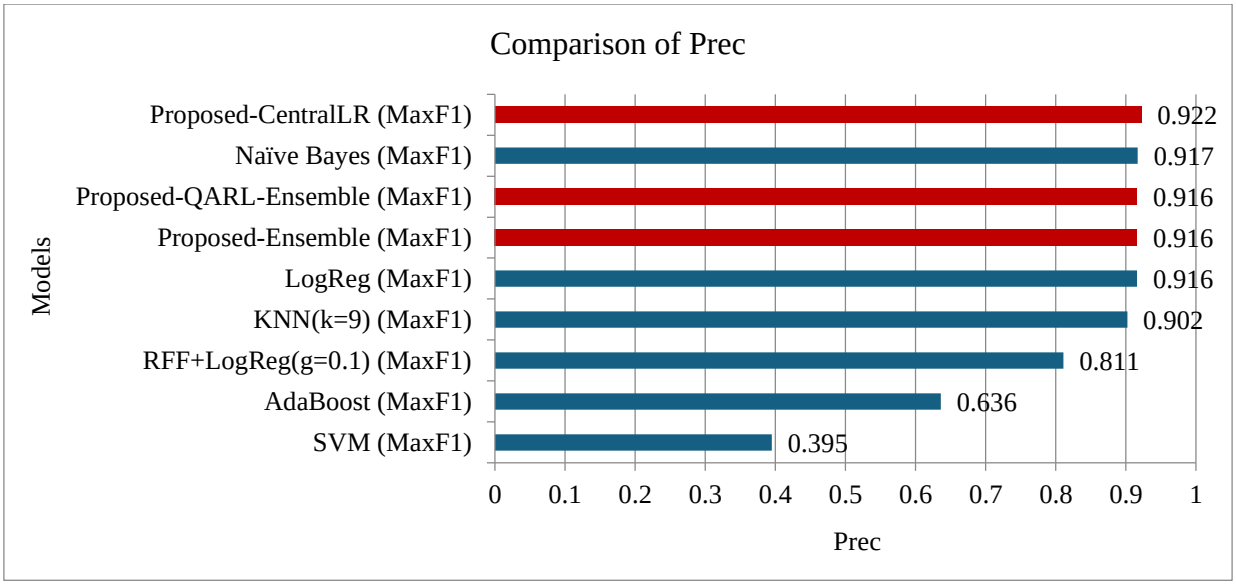

(f) Prec comparison

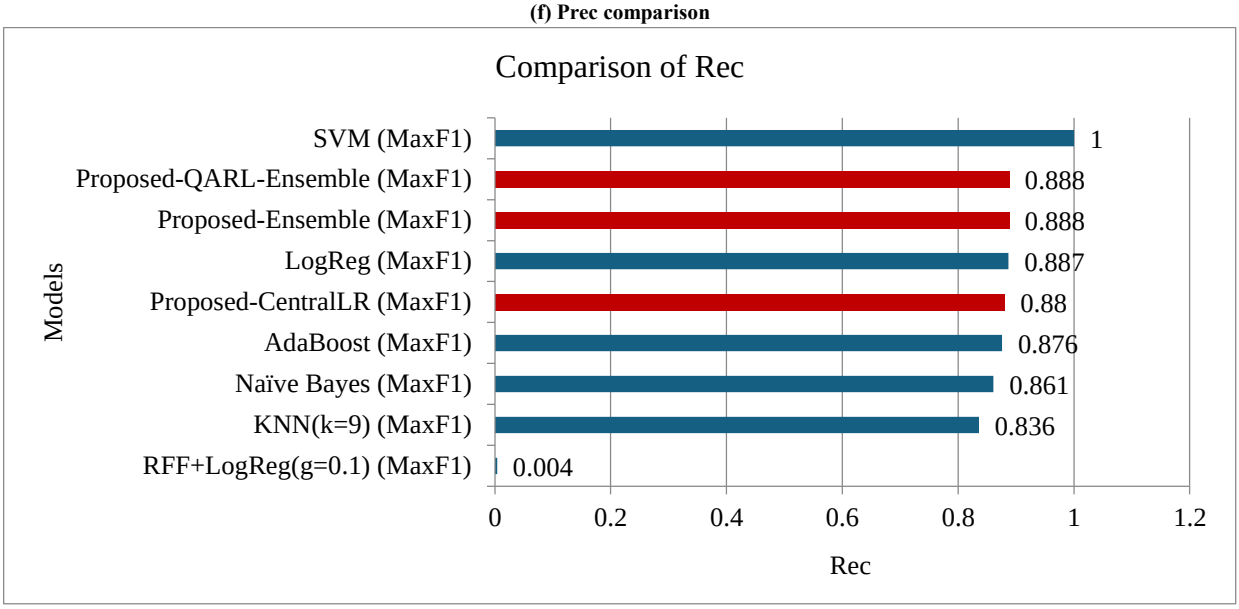

(g) Rec comparison

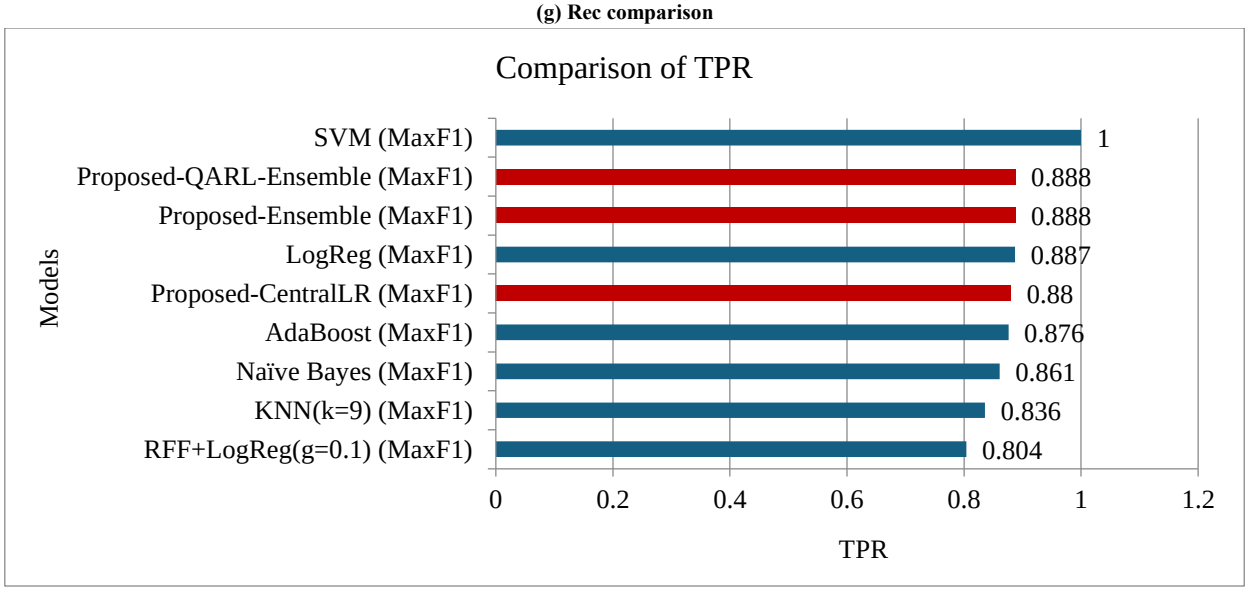

(h) TPR comparison

Fig. 14 Comparative analysis proposed vs Baseline methods

The comparative analysis shows clearly that the proposed methodology is performing better than traditional baselines like Logistic Regression, Naïve Bayes, kNN, AdaBoost, and SVM, with respect to competitive accuracy and F1-scores with low false positive rates. The results indicate the robustness and feasibility of the proposed model and its application in high-impact deployment conditions. The results reveal that the suggested solution has shown consistent competitive or high-performance and remains robust in various assessment measures and conditions. The results prove the effectiveness and applicability of the method in real-world practice.

6.1. Discussion

As shown, the proposed methodology - a mix of QARL threshold adaptation, central, and ensemble logistic regression is already providing balanced and continuously improving results compared to classical baselines. In particular, the Proposed-QARL-Ensemble achieved an accuracy of 0.9240, an F1-score of 0.9022, and an AUC of 0.9757, which is close to and even better than the best baseline, Logistic Regression (Acc 0.9235, F1 0.9016, AUC 0.9753). Silent in this compounding (but important) repair is the methodological aspect of the reinforcement learning based threshold-selection that desaturated the accuracy-recall trade-offs at the operating points.

One particularly interesting observation is that the proposed models are strong when the decision threshold is varied. Naïve Bayes and kNN were also sensitive to false positive rates (F1 0.8880 and 0.8678, respectively), generally worsening when FPR constraints were tightened. However, the proposed QARL-Ensemble did not significantly alter the recall (0.8885) and the precision (0.9162) with the same FPR (around 0.0530). This explains why adaptive thresholding can be used as a method to harmonize the sensitivity of a fixed classifier and thus provides simple tuning with a low false positive tolerance, like that of an anomaly detector of an IoT network.

The analysis also reveals the major weakness of some baselines. AdaBoost parameters provided good recall (0.8762), and the FPR (0.3270) was high, but could never be tolerated in high-stakes domains where false positive is the key issue of paramount importance. This was also previously seen in literature on SVM brittle behaviour in high-dimensional intrusion detection missions: SVMs were also doing terribly (Acc 0.3946, F1 0.5659) due to ill-calibration and sensitivity to data imbalance. However, the RFF+LogReg also has the average balance (Acc 0.8491, AUC 0.9181) whose value is lower than the one at which the proposed method suggests that random fourier expansion is robust in the present case. This shows that random fourier expansion is not necessarily robust in this case as well.

It was no coincidence that MLP was not to be included in the final comparisons. In contrast to MLP, where the values were nearly ideal (Acc 0.9890, F1 0.9863, AUC 0.9999), they were found to vary considerably across runs and were highly sensitive to the initialization, which confirmed earlier findings of deep model overfitting to small datasets or simulated IoT data. Instead of baseless optimized overfitting baselines, the omission of MLP is employed to underscore the viability of the proposed methodology. This increases the realism of the comparative claims and avoids the presentation of findings that can in no way be replicated in practice.

QARL, as a dynamic and reward-driven calibration layer, would allow system maximization of thresholds based on validation feedback. The fixed supervised classification is a forward-looking phase to the active decision making that, in this model, is a phase between the actual world system that is going to act in the conditions of the various classes, priors, and cost specifications. The integration of the ensemble also leads to an increase in stability by reducing variation in the central logistic regression, exhibited in the smaller range of test measures (Figure 11).

The research has limitations, however. It would use a dataset that is not representative of the variability of a realistic traffic in the first place, depending on the diversity of the dataset, and the adversarial behaviours (or long-tail distribution) would be underrepresented. Second, QARL was highly adaptable, although the reinforcement learning policies were not stress-tested on fully non-IID federated splits and were only trained in a centralized setting. Finally, adaptive thresholding and federated random feature expansion involve a certain amount of computation, making them only applicable to very low-resource IoT nodes.

Notwithstanding these qualifiers, the implications of this application are very practical. The methodology offers a means of getting deployable intrusion detection with controlled FPR and with the capability to get an F1 of about 0.90 and an accuracy of about 0.92, a useful enough intrusion detection method may be achieved at a cost not prohibitively expensive to run. The findings are not transferable to the context of IoT security but might be viewed as a model of how to apply the classifier to other applications where falsely labelled individuals are extremely costly, such as fraud detection or physician diagnosis.

7. Conclusion

In this study, a single-pipe artificial intelligence was introduced to combine the method of quantum-inspired feature selection, reinforcement-based threshold optimization, and federated optimization to develop the status of reliable machine intelligence. It too was strictly evaluated by classical criteria with an Accuracy of 92.4, F1-score of 0.902, AUC of 0.976, and Average Precision of 0.933, all demonstrating evidence of its excellent balance between accuracy, recall, and management of false positives. Thus, adaptive decision boundaries were attained using thresholding with QARL, with a Recall of 0.889 at a controlled FPR of 0.053, which is important when risk is of interest. The proposed ensemble was also claimed to perform well or at least equally well compared to fierce competitors such as Naïve Bayes (F1 = 0.888, AUC =

0.967) and kNN (F1 = 0.868, AUC = 0.954). These results suggest that such a model can be effectively used in the field of IoT security, fraud detection, and predictive analytics, with a high probability of fraud occurring. The second work activity will involve scaling the architecture to non-IID federated data and testing the trade-offs of full computations in non-IID real deployment, both of which can be generalised further to quantum-inspired kernels.

References

[1] B. Olanrewaju-George, and B. Pranggono, "Federated Learning-Based Intrusion Detection System for the Internet of Things Using Unsupervised and Supervised Deep Learning Models," *Cyber Security and Applications*, vol. 3, pp. 1-10, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[2] Ansam Khraisat et al., "Federated Learning for Intrusion Detection in IoT Environments: A Privacy-Preserving Strategy," *Discover Internet of Things*, vol. 5, pp. 1-17, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[3] Naresh Kumar Bhagavatham et al., "Autonomic Resilience in Cybersecurity: Designing the Self-Healing Network Protocol for Next-Generation Software-Defined Networking," *International Journal of Computational and Experimental Science and Engineering*, vol. 10, no. 4, pp. 1187-1203, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[4] Gerardo Iovane, "Quantum-Inspired Algorithms and Perspectives for Optimization," *Electronics*, vol. 14, no. 14, pp. 1-27, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[5] Abderahmane Hamdouchi, and Ali Idri, "Optimizing Federated Learning for Intrusion Detection in IoT Networks," *Proceedings of the 16th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management*, vol. 1, pp. 448-456, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[6] Mohamed Abd Elaziz et al., "Federated Learning Framework for IoT Intrusion Detection using Tab Transformer and Nature-Inspired Hyperparameter Optimization," *Frontiers in Big Data*, vol. 8, pp. 1-18, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[7] PratikWaghmode et al., "Intrusion Detection System Based on Machine Learning Using Least Square Support Vector Machine," *Scientific Reports*, vol. 15, pp. 1-23, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[8] Luo Shan, "(IoT) Network Intrusion Detection System using Optimization Algorithms," *Scientific Reports*, vol. 15, pp. 1-19, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[9] Shailendra Pratap Singh et al., "Multi-Objective Quantum Hybrid Evolutionary Algorithms for Enhancing Quality-of-Service in Internet of Things," *Scientific Reports*, vol. 15, pp. 1-27, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[10] Euclides Carlos Pinto Neto et al., "Federated Reinforcement Learning in IoT: Applications, Opportunities and Open Challenges," *Applied Sciences*, vol. 13, no. 11, pp. 1-27, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[11] Tasneem Qasem Al-Ghadi et al., "Leveraging Federated Learning for DoS Attack Detection in IoT Networks Based on Ensemble Feature Selection and Deep Learning Models," *Cyber Security and Applications*, vol. 3, pp. 1-19, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[12] B. Swathi et al., "Enhancing Real-Time IoT Applications in Cloud Communication Networks Through Adaptive Intelligent Routing and Secure Data Transmission," *2025 International Conference on Machine Learning and Autonomous Systems (ICMLAS)*, Prawet, Thailand, pp. 1432-1439, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[13] Fatma S. Alrayes, Syed Umar Amin, and Nada Hakami, "An Adaptive Framework for Intrusion Detection in IoT Security using MAML (Model-Agnostic Meta-Learning)," *Sensors*, vol. 25, no. 8, pp. 1-38, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[14] Haonan Tan et al., "Intrusion Detection Based on Adaptive Sample Distribution Dual-Experience Replay Reinforcement Learning," *Mathematics*, vol. 12, no. 7, pp. 1-24, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[15] P. Venkata Krishna, K. Venkatesh Sharma, and A. MallaRedd, "A Machine Learning-based Approach for Detecting Network Intrusions in Large-scale Networks," *International Journal of Computer Engineering in Research Trends*, vol. 10, no. 2, pp. 69-76, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[16] Yang Liu et al., "Quantum-Inspired Complex Spiking Neural Network for Multi-Source Data Joint Classification," *Information Fusion*, vol. 126, 2026. [CrossRef] [Google Scholar] [Publisher Link]

[17] Rajdeep Kumar Nath, Himanshu Thapliyal, and Travis S. Humble, "Quantum Annealing for Automated Feature Selection in Stress Detection," *2021 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, Tampa, FL, USA, pp. 453-457, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[18] Anjum Nazir et al., "A Novel Feature-Selection Algorithm in IoT Networks for Intrusion Detection," *Sensors*, vol. 23, no. 19, pp. 1-19, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[19] Mousa Alalhareth, and Sung-Chul Hong, "An Improved Mutual Information Feature Selection Technique for Intrusion Detection Systems in the Internet of Medical Things," *Sensors*, vol. 23, no. 10, pp. 1-21, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[20] Quadri Waseem et al., "A Survey and Analysis of Feature Selection Techniques in Machine Learning for IoT Device Classification within Smart Buildings," *Innovative Infrastructure Solutions*, vol. 10, pp. 1-19, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[21] M. Sri Lakshmi et al., "Evaluating the Isolation Forest Method for Anomaly Detection in Software-Defined Networking Security," *Journal of Electrical Systems*, vol. 19, no. 4, pp. 279-297, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[22] Jose García, Jorge Entrena, and Alvaro Alesanco, "Empirical Evaluation of Feature Selection Methods for Machine Learning Based Intrusion Detection in IoT Scenarios," *Internet Things*, vol. 28, pp. 1-18, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[23] G. Logeswari et al., "An Improved Synergistic Dual-Layer Feature Selection Algorithm with Two Type Classifier for Efficient Intrusion Detection in IoT Environment," *Scientific Reports*, vol. 15, pp. 1-21, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[24] Yasset Perez-Rivero et al., "Accurate and Fast Feature Selection Workflow for High-Dimensional Omics Data," *PLoS One*, vol. 12, no. 12, pp. 1-14, 2017. [CrossRef] [Google Scholar] [Publisher Link]

[25] Robert Bauer, and Alireza Gharabaghi, "Reinforcement Learning for Adaptive Threshold Control of Restorative Brain-Computer Interfaces: A Bayesian Simulation," *Frontiers in Neuroscience*, vol. 9, pp. 1-10, 2015. [CrossRef] [Google Scholar] [Publisher Link]

[26] Xue Yang, Enda Howley, and Michael Schukat, "ADT: Time Series Anomaly Detection for Cyber-Physical Systems via Deep Reinforcement Learning," *Computers and Security*, vol. 141, pp. 1-12, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[27] Peijia Qin et al., "Threshold Moving for Online Class Imbalance Learning with Dynamic Evolutionary Cost Vector," *Transactions on Machine Learning Research*, vol. 2024, pp. 1-28, 2024. [Google Scholar] [Publisher Link]

[28] Joffrey L. Leevy et al., "Threshold Optimization and Random Undersampling for Imbalanced Credit Card Data," *Jouranl Big Data*, vol. 10, pp. 1-22, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[29] Gang Xie, Xincheng Zhou, and Jinchun Gao, "Adaptive Trust Threshold Model based on Reinforcement Learning in Cooperative Spectrum Sensing," *Sensors*, vol. 23, no. 10, pp. 1-13, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[30] Pedro Miguel Sánchez Sánchez et al., "FederatedTrust: A Solution for Trustworthy Federated Learning," *Future Generation Computer Systems*, vol. 152, pp. 83-98, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[31] Gangqiang Hu et al., "TRAIL: Trust-Aware Client Scheduling for Semi-Decentralized Federated Learning," *Arxiv Preprint*, pp. 1-10, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[32] Ahmed Alruwaili, Sardar Islam, and Iqbal Gondal, "Fed-DTB: A Dynamic Trust-Based Framework for Secure and Efficient Federated Learning in IoV Networks: Securing V2V/V2I Communication," *Journal of Cybersecurity and Privacy*, vol. 5, no. 3, pp. 1-36, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[33] Jiantao Xu et al., "A Trust-Aware Incentive Mechanism for Federated Learning with Heterogeneous Clients in Edge Computing," *Journal of Cybersecurity and Privacy*, vol. 5, no. 3, pp. 1-19, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[34] J. Geetha Renuka et al., "BERT-LogReg: Enhancing Phishing URL Detection with Transformr-Based Features," *2024 International Conference on Emerging Systems and Intelligent Computing (ESIC)*, Bhubaneswar, India, pp. 660-665, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[35] Levent Koc, and Alan D. Carswell, "Network Intrusion Detection Using a Hidden Naïve Bayes Binary Classifier," *2015 17th UKSim-AMSS International Conference on Modelling and Simulation (UKSim)*, Cambridge, UK, pp. 81-85, 2015. [CrossRef] [Google Scholar] [Publisher Link]

[36] Sufyan Othman Zaben, "IDC-Insight: Boosting Intrusion Detection Accuracy in IoT Networks with Naïve Bayes and Multiple Classifiers," *International Journal of Information Technology*, vol. 17, pp. 869-881, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[37] N. Harini et al., "CyberSecurity Intrusion Detection in Industry 4.0 WSN's Using ML/DL," *International Journal of Computer Engineering in Research Trends*, vol. 12, no. 2, pp. 18-28, 2025. [CrossRef] [Publisher Link]

[38] Silpa Chalichalamala, Niranjana Govindan, and Ramani Kasarapu, "Logistic Regression Ensemble Classifier for Intrusion Detection System in Internet of Things," *Sensors*, vol. 23, no. 23, pp. 1-19, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[39] Christiana Ioannou, and Vasos Vassiliou, "Network Attack Classification in IoT Using Support Vector Machines," *Journal of Sensor and Actuator Networks*, vol. 10, no. 3, pp. 1-17, 2021. [CrossRef] [Google Scholar] [Publisher Link]