Navigating the Precision-Recall Trade-off in Network Intrusion Detection with a Hybrid Quantum-Classical LSTM

1st Phuc Hao Do
Department of Software Engineering
Danang Architecture University
Da Nang, Viet Nam
haodp@dau.edu.vn

2nd Tran Duc Le
Department of Mathematics, Statistics & Computer Science
University of Wisconsin-Stout
Menomonie, WI, USA
let@uwstout.edu

3rd Truong Duy Dinh
Faculty of Information Security
Posts and Telecommunications Institute of Technology
Ha Noi, Viet Nam
duydt@ptit.edu.vn

4th Van Dai Pham
Swinburne Vietnam
FPT University
Ha Noi, Viet Nam
daipv11@fe.edu.vn

***Abstract*—As network attacks grow in sophistication, Intrusion Detection Systems (IDS) face a critical operational challenge: maximizing detection rates without generating an overwhelming number of false alarms. While classical Long Short-Term Memory (LSTM) networks are effective for analyzing sequential traffic data, achieving this balance is difficult. This paper investigates the potential of Quantum Machine Learning (QML) to address this problem by introducing and evaluating a hybrid Quantum-LSTM (QLSTM) model. Our architecture uses a classical LSTM for temporal feature extraction and a Parameterized Quantum Circuit (PQC) for classification. Evaluating our model on the large-scale CIC-IDS2017 dataset, we uncover a critical performance trade-off. Our central finding is that while a purely classical LSTM baseline achieves a slightly higher F1-Score driven by superior recall, the hybrid QLSTM model consistently delivers a significant advantage in precision, substantially reducing the false positive rate. The primary contribution of this work is the identification and practical contextualization of this valuable trade-off, demonstrating that QML offers a compelling pathway toward high-fidelity IDS where alert reliability is paramount.**

***Index Terms*—Quantum Machine Learning, Quantum LSTM, Network Intrusion Detection, Cybersecurity, Hybrid Quantum-Classical Models, Deep Learning.**

I. Introduction

Network security has become a cornerstone of modern digital infrastructure, yet it faces an ever-evolving threat landscape. Adversaries are deploying increasingly sophisticated attack vectors, such as Advanced Persistent Threats (APTs) [1] and zero-day exploits, which are designed to bypass traditional security measures. Conventional Intrusion Detection Systems (IDS), which often rely on predefined signatures of known attacks, are proving insufficient in detecting these novel and stealthy threats, creating a critical security gap in enterprise and governmental networks.

To address the limitations of signature-based methods, the cybersecurity community has shifted towards data-driven approaches using machine learning (ML) and deep learning (DL). These techniques enable an IDS to learn the complex patterns of normal network behavior and identify deviations indicative of an attack, without prior knowledge of the attack's signature. In particular, Recurrent Neural Networks, such as Long Short-Term Memory (LSTM) [2], have demonstrated significant success. By modeling network traffic as sequential data, LSTMs can capture temporal dependencies and context, making them well-suited for detecting multi-stage attacks that unfold over time.

Despite their success, classical DL models are not a panacea. The expressive power of a classical neural network is ultimately bounded by its architecture and the non-linear activation functions it employs. As attack patterns become more intricate and blend seamlessly with benign traffic [3], it is plausible that classical models like LSTMs may reach a performance plateau, struggling to learn the highly complex decision boundaries required to distinguish subtle anomalies. This limitation motivates the exploration of more powerful computational paradigms.

Quantum Machine Learning (QML) has emerged as a promising new frontier in this regard [4]. By leveraging the principles of quantum mechanics, QML models operate in a vastly larger computational space

Corresponding author: duydt@ptit.edu.vn

- the exponential-sized Hilbert space. Concepts such as superposition and entanglement allow quantum circuits to represent and process information in ways that are classically intractable. A key component of modern QML is the Parameterized Quantum Circuit (PQC), a quantum circuit with trainable parameters analogous to a classical neural network layer [5]. Also known as variational quantum circuits, PQCs are well-suited for integration into hybrid training frameworks, where they can be combined with classical deep learning models.

While hybrid quantum-classical RNNs have shown promise in other sequential data domains such as natural language processing and finance, their application to cybersecurity remains nascent. This paper investigates a critical question in this domain: can the enhanced representational capacity of quantum computing be leveraged not just to improve aggregate performance metrics, but to achieve a more favorable and practical performance profile for real-world network intrusion detection? To this end, we propose and evaluate a novel hybrid Quantum-LSTM (QLSTM) model [6], using it as a tool to explore the trade-offs inherent in this complex classification task. Our primary contributions are threefold:

- We design a novel hybrid architecture that synergistically combines a classical LSTM network for temporal feature extraction with a PQC for classification, building upon recent work in hybrid quantum-classical RNNs.
- We conduct a comprehensive empirical evaluation of our model and its variants on the large-scale, realistic CIC-IDS2017 dataset [7], establishing a solid benchmark for this task.
- We uncover and analyze a critical performance trade-off: while a strong classical baseline excels in recall, our QLSTM model consistently achieves superior precision. This finding is of significant practical importance, as it demonstrates a path toward high-fidelity intrusion detection systems that can substantially reduce the costly burden of false alarms in operational security environments.

The remainder of this paper is organized as follows. Section II reviews related work on deep learning for intrusion detection and quantum machine learning. Section III details our proposed hybrid QLSTM model and the data preprocessing pipeline. Section IV presents the experimental setup and baseline models. Section V discusses the empirical results and provides a comparative analysis. Finally, Section VI concludes the paper and outlines future research directions.

II. Background and Related Work

Our research is situated at the intersection of deep learning for cybersecurity and the emerging field of quantum machine learning. This section briefly reviews the key concepts from both domains that form the foundation of our work.

A. Deep Learning for Intrusion Detection

Traditional IDS [8] often struggle against novel and sophisticated threats. Consequently, DL has become a primary paradigm for developing data-driven, adaptive IDS. The ability of DL models to automatically learn hierarchical features from complex data makes them highly effective. Given the sequential nature of network traffic, Recurrent Neural Network (RNN) architectures, particularly LSTM and Gated Recurrent Units (GRU) [9], have proven especially successful. These models can capture temporal dependencies in network flow or packet sequences, enabling the detection of multi-stage attacks. Numerous studies have validated the effectiveness of RNN-based approaches on benchmark datasets like NSL-KDD, UNSW-NB15, and CIC-IDS2017 [7]. However, challenges such as severe class imbalance and the potentially limited expressive power of classical non-linearities motivate the exploration of alternative computational models.

B. Hybrid Quantum Machine Learning

Quantum Machine Learning is an interdisciplinary field that explores the interplay between quantum computing and machine learning [10]. It aims to devise quantum algorithms that can perform learning tasks faster or more effectively than their classical counterparts. The potential advantages of QML stem from the fundamental principles of quantum mechanics.

1) Quantum Bits (Qubits): The basic unit of quantum information is the qubit. Unlike a classical bit, which can only be in a state of 0 or 1, a qubit can exist in a *superposition* of both states simultaneously. A single qubit state, denoted $|\psi\rangle$, is represented as a linear combination of the basis states $|0\rangle$ and $|1\rangle$:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \tag{1}$$

where α and β are complex numbers known as probability amplitudes, satisfying the normalization condition $|\alpha|^2 + |\beta|^2 = 1$. This property allows a register of N qubits to represent 2^N classical states at once, providing an exponentially large computational space. Furthermore, multiple qubits can be *entangled*, creating strong, non-classical correlations between them that are a key resource for quantum computation.

2) Parameterized Quantum Circuits: In the current era of Noisy Intermediate-Scale Quantum (NISQ) devices, one of the most promising QML paradigms is the PQC [11], also commonly referred to as a variational quantum circuit. A PQC is a quantum circuit composed of a sequence of quantum gates, where some of these gates have tunable parameters (e.g., rotation angles, θ). This structure serves as the quantum equivalent of a

classical neural network layer. The circuit, represented by a unitary transformation $U(\theta)$, typically consists of three parts: an encoding layer $U_{enc}(x)$ that maps classical data x into a quantum state, a variational ansatz $U_{var}(\theta)$ with trainable parameters, and a measurement of an observable $\hat{M}$ to extract classical information from the final quantum state.

3) Hybrid Quantum-Classical Training: PQCs are trained using a hybrid quantum-classical loop, which is fully compatible with standard deep learning frameworks like PyTorch or TensorFlow. The output of the hybrid model is the expectation value of the observable, $f(x, \theta) = \langle 0|U^\dagger(\theta, x)\hat{M}U(\theta, x)|0\rangle$. The training workflow is as follows:

- **Forward Pass:** A batch of classical input data x is fed into the PQC. This data is used to set the parameters of the encoding gates. The circuit is then executed with the current set of trainable weights θ. The final quantum state is measured, yielding the classical expectation value $f(x, \theta)$.
- **Loss Calculation:** This classical output is used in a classical loss function $\mathcal{L}$ (e.g., cross-entropy) to compare against the true labels.
- **Backward Pass:** The gradient of the loss function with respect to the trainable parameters, $\nabla_\theta \mathcal{L}$, is calculated. This can be done efficiently on quantum hardware or simulators using techniques like the parameter-shift rule.
- **Parameter Update:** A classical optimizer (e.g., Adam) uses these gradients to update the PQC's parameters: $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}$.

This hybrid approach allows us to treat the PQC as a differentiable layer that can be seamlessly integrated into any classical deep learning architecture [12], as we demonstrate with our proposed QLSTM model. Theoretically, such hybrid QML models can be interpreted as powerful kernel methods [13]. By implicitly mapping classical data to a high-dimensional quantum feature space (the Hilbert space), the PQC can transform complex, non-linear classification problems into ones that are more easily separable, providing a basis for learning more robust and precise decision boundaries.

III. Proposed Methodology

This section details our proposed hybrid quantum-classical framework for network intrusion detection. We first describe the data preparation pipeline, then formulate the problem mathematically and present the detailed architecture of our hybrid QLSTM model [1].

A. Data Preprocessing and Sequencing

We utilize the CIC-IDS2017 dataset, which contains over 2.8 million network flows described by 78 features. Our preprocessing pipeline involves four main steps: 1) **Data Cleaning** to remove corrupted entries (NaN/Infinity values); 2) **Label Transformation**, where we convert the multi-class problem into a binary classification task by consolidating all 14 attack types into a single 'ATTACK' class (1) against the 'BENIGN' class (0); 3) **Feature Scaling**, where we apply Min-Max normalization to scale all features to the range [0, 1], fitting the scaler only on the training set; and 4) **Sequencing**, where the tabular data is transformed into time-series data. We employ a sliding window of length $L = 20$ to group consecutive flows, creating sequences of shape (N, L, D), where N is the number of sequences, $L = 20$ is the number of timesteps, and $D = 78$ is the number of features. This value was chosen as it balances the need to capture sufficient temporal context from the network flow sequences against the computational overhead of processing longer sequences.

[1] https://github.com/ailabteam/qml-ids

B. Hybrid QLSTM Architecture

Given a dataset of sequences $\mathcal{D} = \{(S_i, y_i)\}_{i=1}^N$, where $S_i \in \mathbb{R}^{L \times D}$ and $y_i \in \{0, 1\}$, our goal is to learn a hybrid function $h(S_i; \theta_c, \theta_q)$ parameterized by classical (θ_c) and quantum (θ_q) weights. The model architecture,

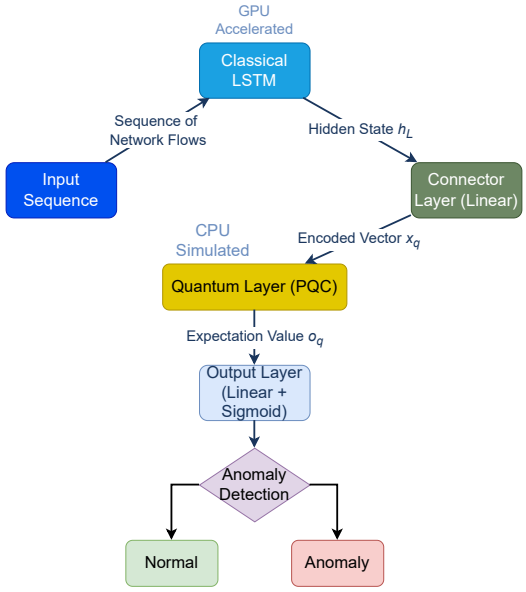

Fig. 1. The proposed hybrid QLSTM architecture. A classical LSTM extracts temporal features, which are then processed by a PQC for classification.

1) Classical LSTM Feature Extractor: An LSTM layer processes the input sequence S_i and outputs its final hidden state, $\mathbf{h}_L \in \mathbb{R}^{d_h}$, which serves as a compressed representation of the sequence's temporal features.

$$\mathbf{h}_L = \text{LSTM}(S_i; \theta_{\text{lstm}}). \quad (2)$$

2) Classical-to-Quantum Connector: A fully-connected layer maps the hidden state $\mathbf{h}_L$ to a vector $\mathbf{x}_q \in \mathbb{R}^{n_q}$, where n_q is the number of qubits. This serves as the input for the quantum circuit.

$$\mathbf{x}_q = W_c \mathbf{h}_L + \mathbf{b}_c. \tag{3}$$

3) Parameterized Quantum Circuit Classifier: The PQC acts as the quantum classifier. Its structure consists of three stages:

- **State Preparation (Embedding):** The classical vector $\mathbf{x}_q \in \mathbb{R}^{n_q}$ is encoded into the initial quantum state using angle embedding. Each component $x_{q,j}$ is used to rotate the corresponding qubit $|0\rangle$ around the Y-axis: $U_{enc}(\mathbf{x}_q) = \bigotimes_{j=0}^{n_q-1} R_Y(x_{q,j})$.
- **Variational Ansatz:** A layered ansatz, $U_{var}(\theta_q)$, with trainable parameters θ_q is applied. Each layer consists of single-qubit rotations ($R_Y(\theta)$ and $R_Z(\theta)$ on all qubits) followed by a linear chain of CNOT gates for entanglement (i.e., CNOT(i, i+1) for all adjacent qubits). This two-part structure is repeated d_q times, where d_q is the circuit depth.
- **Measurement:** Finally, the expectation value of the Pauli-Z observable is measured on the first qubit ($\hat{Z}_0 = \hat{Z} \otimes \hat{I} \otimes \cdots \otimes \hat{I}$). This is a standard method for extracting a classical prediction from a quantum circuit, as it maps the final quantum state to a scalar value $o_q \in [-1, 1]$ that is suitable for subsequent use in a binary classification task.

$$o_q = f_q(\mathbf{x}_q; \theta_q) = \langle 0 | U_{enc}^\dagger(\mathbf{x}_q) U_{var}^\dagger(\theta_q) \hat{Z}_0 U_{var}(\theta_q) U_{enc}(\mathbf{x}_q) | 0 \rangle. \tag{4}$$

4) Classical Output Layer: The scalar output o_q from the PQC is passed through a final linear layer with a sigmoid activation function σ to yield the final probability prediction.

$$\hat{y} = \sigma(W_o o_q + b_o). \tag{5}$$

The entire model is differentiable and is trained end-to-end.

IV. Experimental Setup

To rigorously evaluate the performance of our proposed QLSTM model, we designed a controlled experimental setup. This section details the dataset, implementation parameters, the models used for comparison, and the evaluation metrics.

A. Dataset and Implementation

Our experiments are conducted on the CIC-IDS2017 dataset, a widely recognized benchmark for IDS evaluation. After the preprocessing pipeline described in Section III-A, our final dataset consists of 2,827,876 sequential samples. We formulate the task as a binary classification problem ('BENIGN' vs. 'ATTACK'). The dataset is highly imbalanced, with the 'ATTACK' class comprising only 19.7% of the samples. We use a stratified 80/20 split for training and testing.

Our models are implemented in PyTorch and PennyLane. All experiments are run for 5 epochs with a batch size of 256, using the Adam optimizer with a learning rate of $\eta = 0.001$. Training was performed on a system with an NVIDIA GeForce RTX 4090 GPU for classical components and a multi-core CPU for quantum circuit simulation.

B. Models for Comparison

To isolate the contribution of the quantum layer and explore the impact of key hyperparameters, we evaluate the following four models:

- **Classical LSTM (Baseline):** This model shares an identical architecture with our QLSTM, but the PQC classifier is replaced by a classical Multi-Layer Perceptron (MLP). The MLP head, which takes the 32-dimensional LSTM hidden state as input, consists of two fully-connected layers (32x8 and 8x1) with a ReLU activation and a Dropout layer. It contains 273 trainable parameters.
- **QLSTM-4Q-32H (Default):** Our default hybrid model, featuring an LSTM with a hidden size of $d_h = 32$ and a PQC classifier using $n_q = 4$ qubits and a depth of $d_q = 2$.
- **QLSTM-4Q-64H:** A variant designed to test the impact of classical representational power, using a larger LSTM hidden size of $d_h = 64$ with a 4-qubit PQC ($d_q = 2$).
- **QLSTM-8Q-32H:** A variant designed to test the impact of quantum representational power, using a standard LSTM ($d_h = 32$) with a larger 8-qubit PQC ($n_q = 8, d_q = 2$).

The classical baseline was designed to ensure a fair comparison of classification power. Its MLP head (273 parameters) is parametrically comparable to the entire quantum classification head (connector layer plus PQC ansatz) of the main QLSTM model. For instance, the QLSTM-8Q-32H's head contains a total of 296 trainable parameters ($(32 \times 8) + 8$ in the connector and $8 \times 2 \times 2 = 32$ in the PQC), validating the fairness of the architectural comparison.

V. Results and Discussion

In this section, we present and analyze the empirical results of our experiments. We provide a unified comparative analysis based on a comprehensive results table, examine classification errors using a confusion matrix, and discuss the broader implications of our findings.

A. Performance Analysis and Hyperparameter Impact

The definitive results of our comparative evaluation are presented in Table I. This table provides a complete overview of all model configurations evaluated on the full CIC-IDS2017 test set, including precision, recall, F1-score, and overall accuracy. The chart in Fig. 2 serves as a visual summary of the core precision-recall trade-off detailed in the table.

TABLE I
COMPREHENSIVE PERFORMANCE COMPARISON OF ALL MODEL CONFIGURATIONS. THE CLASSICAL MODEL ACHIEVES THE HIGHEST RECALL AND F1-SCORE, WHILE ALL QLSTM VARIANTS DEMONSTRATE SUPERIOR PRECISION. BEST SCORES FOR EACH METRIC ARE IN **BOLD**.

Model Config.	Acc. (%)	Prec.	Rec.	F1
Classical LSTM	97.21	0.9117	**0.9825**	**0.9457**
QLSTM-4Q-32H (Default)	97.35	0.9250	0.9420	0.9334
QLSTM-4Q-64H	**97.58**	0.9181	0.9624	0.9397
QLSTM-8Q-32H	97.43	**0.9312**	0.9396	0.9354

The results reveal a clear and nuanced performance trade-off. The purely classical LSTM baseline achieves the highest F1-Score (0.9457), which is driven by an exceptionally high recall of 0.9825. This indicates the classical model is highly effective at identifying the vast majority of true attack instances. However, this high sensitivity comes at the cost of the lowest precision (0.9117), implying a greater propensity for false alarms.

In stark contrast, all QLSTM configurations consistently outperform the classical baseline in terms of precision. The QLSTM-8Q-32H model, in particular, reaches the highest precision of 0.9312. This finding is central to our work, as it suggests that the quantum classifier is more adept at learning the fine-grained characteristics that distinguish true attacks from benign traffic, leading to more reliable and confident predictions. This higher precision is a valuable asset in real-world Security Operations Centers (SOCs) where alert fatigue from false positives is a major operational concern.

Table I also illuminates the impact of hyperparameters. Increasing the classical LSTM's hidden size from 32 to 64 (QLSTM-4Q-64H) yielded the best F1-Score (0.9397) and accuracy (97.58%) among the hybrid models, primarily by boosting recall. Conversely, increasing the quantum capacity by moving from 4 to 8 qubits (QLSTM-8Q-32H) produced the highest precision. This demonstrates a key trade-off between classical feature representation and quantum processing capacity, suggesting that a larger quantum space can refine the decision boundary to reduce false positives, but may require a correspondingly richer classical input to maintain high recall.

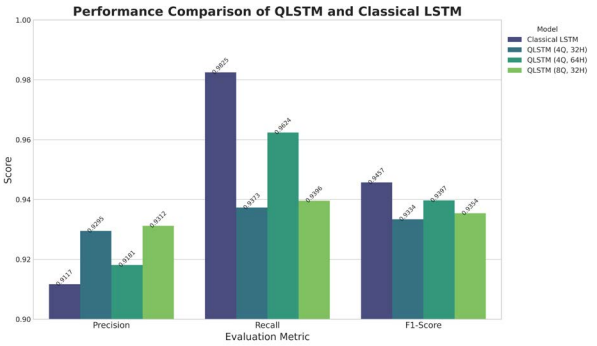

Fig. 2. Visual summary of the performance trade-off detailed in Table I. The QLSTM models consistently show higher precision, while the classical LSTM excels in recall.

B. Analysis of Classification Errors

To provide a more granular view of performance, Fig. 3 displays the confusion matrix for our highest-precision model, the QLSTM-8Q-32H. The matrix shows the model correctly identified 104,582 attack instances (True Positives) while missing 6,723 (False Negatives). Crucially, it only misclassified 7,723 benign instances as attacks (False Positives). When compared to the classical LSTM baseline (which had 10,879 False Positives, derived from its performance scores), our QLSTM model reduced the number of false alarms by approximately 29%. This directly reinforces the finding that the quantum layer contributes to a more discerning and reliable classification. The balance between minimizing missed attacks (FN) and reducing false alarms (FP) is a critical consideration in deploying any real-world IDS, and our results suggest that QML offers a promising new tool to navigate this trade-off.

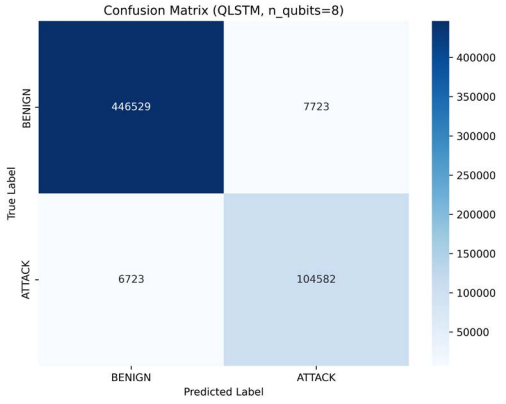

Fig. 3. Confusion matrix for the highest-precision model (QLSTM-8Q-32H) on the test set.

C. Discussion

Our experimental results provide a nuanced perspective on the current capabilities of hybrid quantum models for this task. While the classical LSTM baseline achieved a slightly higher overall F-Score, our investigation revealed a crucial and consistent trade-off: all configurations of the proposed QLSTM model demonstrated superior precision. This finding is significant for practical applications where the cost of investigating false positives is high and alert fidelity is paramount.

This trade-off suggests that the quantum classifier learns a more stringent and less noisy decision boundary. As outlined in Section II, this can be attributed to the PQC's function as a kernel method that maps features into the high-dimensional Hilbert space. This mapping can make the data more separable, allowing the model to define a simpler decision boundary that is highly confident about positive predictions (improving precision), while potentially being more conservative on ambiguous, borderline cases (impacting recall). The findings of this work underscore the potential of QML not necessarily as a blanket replacement for classical models, but as a tool to explore different performance profiles and trade-offs that are not accessible to purely classical approaches.

However, we acknowledge the limitations of our study. The quantum component was executed on a classical simulator, which is computationally expensive, especially as the number of qubits increases. The performance on actual noisy quantum hardware remains an open question. Furthermore, our study focused on a binary classification task.

VI. Conclusion

This paper investigates the practical application of hybrid quantum-classical deep learning for network intrusion detection, revealing a critical performance trade-off. Our central finding, derived from evaluating the QLSTM model against a strong classical baseline on the CIC-IDS2017 dataset, is that the quantum-enhanced model consistently achieves superior precision. While the classical model excels in recall, the QLSTM's ability to significantly reduce false alarms demonstrates the potential of QML as a tool for engineering high-fidelity cybersecurity systems where alert reliability is paramount. The key contribution of this work is therefore not the proposal of a model that universally outperforms classical methods, but the discovery and practical contextualization of a valuable trade-off enabled by the hybrid QML approach.

Future work will proceed in three main directions. First, we will explore more advanced quantum circuit designs (ansatze) and data encoding strategies to improve the model's recall without sacrificing its high precision. Second, we will investigate the model's robustness against adversarial attacks. Finally, we aim to apply and validate the hybrid QLSTM framework on other sequential cybersecurity tasks, such as malware analysis or malicious URL detection.

Code Availability

The source code used in this study are publicly available on GitHub: https://github.com/ailabteam/qml-ids.

Acknowledgements

This work has been sponsored by the scientific research from Posts and Telecommunications Institute of Technology, Vietnam

References

[1] I. Stellios, P. Kotzanikolaou, and M. Psarakis, *Advanced Persistent Threats and Zero-Day Exploits in Industrial Internet of Things*. Springer International Publishing, 2019, p. 47–68.

[2] P. S. Muhuri, P. Chatterjee, X. Yuan, K. Roy, and A. Esterline, "Using a long short-term memory recurrent neural network (lstm-rnn) to classify network attacks," *Information*, vol. 11, no. 5, p. 243, May 2020.

[3] B. Lampe and W. Meng, "Intrusion detection in the automotive domain: A comprehensive review," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 4, p. 2356–2426, 2023.

[4] S. M. N. Islam, P. Kinger, N. Fatima, and A. Kumar, *Quantum Machine Learning: Bridging Classical and Quantum Frontiers*. Springer Nature Singapore, 2024, p. 325–347.

[5] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, "Parameterized quantum circuits as machine learning models," *Quantum Science and Technology*, vol. 4, no. 4, p. 043001, Nov. 2019.

[6] S. Tripathi, H. Upadhyay, and J. Soni, "A quantum lstm-based approach to cyber threat detection in virtual environment," *The Journal of Supercomputing*, vol. 81, no. 1, Nov. 2024.

[7] I. Sharafaldin, A. Habibi Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," in *Proceedings of the 4th International Conference on Information Systems Security and Privacy*. SCITEPRESS - Science and Technology Publications, 2018.

[8] P. H. Do, T. D. Le, V. Vishnevsky, A. Berezkin, and R. Kirichek, "A horizontal federated learning approach to iot malware traffic detection: An empirical evaluation with n-baiot dataset," in *2024 26th International Conference on Advanced Communications Technology (ICACT)*. IEEE, Feb. 2024, p. 1494–1506.

[9] A. F. M. Agarap, "A neural network architecture combining gated recurrent unit (gru) and support vector machine (svm) for intrusion detection in network traffic data," in *Proceedings of the 2018 10th International Conference on Machine Learning and Computing*, ser. ICMLC 2018. ACM, Feb. 2018, p. 26–30.

[10] J. Qi, C.-H. H. Yang, S. Y.-C. Chen, and P.-Y. Chen, "Quantum machine learning: An interplay between quantum computing and machine learning," in *2025 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, May 2025, p. 1–5.

[11] F. Chen, L. Jiang, H. Müller, P. Richerme, C. Chu, Z. Fu, and M. Yang, "Nisq quantum computing: A security-centric tutorial and survey [feature]," *IEEE Circuits and Systems Magazine*, vol. 24, no. 1, p. 14–32, 2024.

[12] A. Sharma and S. Rani, "Post-quantum cryptography (pqc) for iot-consumer electronics devices integrated with deep learning," *IEEE Transactions on Consumer Electronics*, vol. 71, no. 2, p. 4925–4933, May 2025.

[13] S. Jerbi, L. J. Fiderer, H. Poulsen Nautrup, J. M. Kübler, H. J. Briegel, and V. Dunjko, "Quantum machine learning beyond kernel methods," *Nature Communications*, vol. 14, no. 1, Jan. 2023.