Architecture
Image Studies

Vol.7, Issue 1, pp.165-177, 2026
DOI: https://doi.org/10.62754/ais.v7i1.814

https://journals.ap2.pt/index.php/ais/index

A Quantum Theory-Based Hybrid Neural Model to Improve Textual Intelligence for Threat Entity Classification

Shyam R[1], Dafik[2], Siva Shankar S[3], R. Sunder[4], Ik Hesti Agustin[5]

Abstract

Here we introduce an uncomplicated hybrid model for threat-entity classification: A typical neural network progresses in parallel with a quantum inspired course that rotates features and uses a wave-like attention. Sentence embedding we couple with some basic counts of entities, project back down in dimensionality, and balance the classes with SMOTE; Isolation Forest and Random Forest assist in capturing edge-case oddities. On a structured cyber-threat data set, the model achieves 97.2% accuracy and 0.97 F1-score (precision 0.96, recall 0.98, AUC 0.99), with some misses (~2.2% FN) and some false alarms (~4.3% FP). The trade-off is interpretability: the quantum-inspired pathway improves generalizability but makes the decisions harder to interpret.

Keywords: *Quantum-Inspired Deep Learning, Threat Entity Classification, Cybersecurity, Natural Language Processing (NLP), Hybrid Neural Model, Anomaly Detection, Data Imbalance (SMOTE), Sentence-BERT, Quantum En- cryption, Quantum Attention.*

Introduction

The rise in the sophistication level of cyber threats has created a growing need for intelligent systems capable of accurately recognizing malicious patterns and analysing unstructured threat stories. Natural language processing (NLP) has become a majormfacilitator in this endeavour, providing technologies to convert unstructured cyber intelligence into structured, learnable formats. The semantic extraction of contextual information in threat reports, incident logs, and malware analysis summaries has pro- gressed significantly with the advent of sentence-level transformer models and so- phisticated embedding methods [4, 5, 28]. These developments have opened up new frontiers in classification, risk assessment, and anomaly detection by enabling deep learning algorithms to process text beyond surface features.

However, there remain difficult obstacles to using deep learning for NLP with a focus on cybersecurity. The most important of these is the issue of data imbalance, in which benign or neutral inputs overwhelmingly outnumber the presence of malicious or high-risk signals. Typical models are prone to fail in detecting strong, low- frequency signals because of this skewed distribution, and this results in high false- negative rates [1], [14], and [29]. Furthermore, although anomaly detection techniques such as Isolation Forest and its more recent extensions have been effective, they are typically applied separately from deep learning systems and do not take ad- vantage of shared representation learning and optimization [2], [3], or [14]. The abil- ity to develop integrated models that might be able to detect structural anomalies and semantic drifts simultaneously is constrained by this dichotomy.

The model generalization and robustness challenges are also significant. When de- ployed, traditional deep learning models—even those based on BERT-like embed- dings-enhanced models—exhibit sensitivity to adversarial inputs, data inconsistency, and latent patterns. Moreover, their decision-making is opaque, which is a challenge in domains such as cybersecurity, where explainability and adaptability are essential [10], [14], and [30]. Quantum-inspired computing models, which introduce

[1] Presidency College, Hebbal Banglore, India. Shyam.r@presidency.edu.in.
[2] University of Jember, Jember, Jawa Timur 68121, Indonesia, Email: d.dafik@unej.ac.id.
[3] KG Reddy College of Engineering and Technology, Hyderabad, Telangana, India, Email: drsivashankars@gmail.com
[4] School of Computer Science and Engineering, Galgotias University, Uttar Pradesh, India – 203201, Email: sunder.r@galgotiasuniversity.edu.in
[5] Department of Mathematics, University of Jember, Indonesia PUI-PT CGANT, University of Jember, Indonesia, Email: ikahesti.fmipa@unej.ac.id

novel mathematical transformations to enhance feature representations and facilitate more effective learning dynamics, have been seen as a potential solution to these challeng- es. The models have achieved promising performance in a broad spectrum of domains and replicate certain features of quantum systems, including interference patterns and orthonormal transformations [5, 6, 13, 21].

Here, we introduce a hybrid deep learning architecture to enhance threat detection accuracy and resilience in text data by combining quantum-inspired modelling meth- ods with traditional feedforward processing. The core of the design is a dual-branch architecture: one path processes structured embeddings using traditional dense layers, while the other path performs a sequence of quantum-inspired computations, includ- ing a decryption layer, a trigonometric attention mechanism, and a non-trainable or- thonormal encryption layer. Contrary to the popular myth, our model's quantum en- cryption layer is not used for data privacy aspects. Instead, it is used mainly to im- prove model performance by adding a consistent transformation that results in im- proved generalization during training. The network is able to learn more generalizable and stable patterns by mapping data into a new orthonormal basis, thus subjecting it to a rich set of feature interactions [5, 6, 20, 24].

By generating more sophisticated and varied internal representations, the method improves our model's performance on imbalanced input. The structure also gains from a late-fusion process that combines statistical anomaly scores with neural confidence using ensemble-based anomaly detection models, such as Deep Isolation Forest and

Random Forest, at the output layer [2], [3], [11]. The fusion improves classification margins, especially edge cases and doubtful samples [19], without depending on so- phisticated multi-task learning techniques. The originality of our contribution is in its conscious and beneficial application of quantum-inspired layers to address significant threat detection issues. We demonstrate the usefulness of such layers in an actual- world cybersecurity NLP pipeline, as opposed to other efforts that applied them under theoretical or domain-separated scenarios. Moreover, as opposed to considering such layers as theoretical embellishments or privacy-oriented enhancements, our approach considers them as pragmatic enhancements that directly enhance model accuracy and robustness of learning. Due to this distinction, an efficient and generalizable system can be constructed, and hence more immune to adversarial noise and class imbalance. Additionally, the architecture's ability to learn from rich semantic embeddings as well as entity-based structural features ensures that it preserves the context and specificity of threats, something that many of the past models fail to effectively balance [12], [13], [18], and [27].

Literature Review

Deep Learning Approaches for Anomaly Detection

Recent advancements in deep learning have significantly improved the detection and classification of anomalies in complex data systems. Gayathri et al. (2024) proposed a hybrid SPCAGAN model that enhanced data quality and reduced false positive rates for insider threat detection. Their findings highlighted the limitations of traditional GAN-based architectures in capturing fine-grained contextual features within cybersecurity datasets. Similarly, Patel et al. (2022) and Chen & Gao (2021) demonstrated the applicability of quantum-inspired tensor networks and convolutional neural networks (CNNs) in solving high-dimensional partial differential equations and vision-based tasks. These studies validated the feasibility of using deep neural architectures for modeling intricate and nonlinear anomaly patterns. However, despite these successes, conventional deep learning models continue to face challenges in interpretability, computational scalability, and robustness under noisy or incomplete data conditions.

Quantum and Quantum-Inspired Neural Architectures

In parallel, quantum computing research has evolved toward integrating quantum mechanics with neural computation. Foundational works by Peters and Caldeira (2021), Tacchino et al. (2019), and Farhi and Neven (2018) introduced quantum neurons and networks capable of operating effectively under noisy quantum environments. Dunjko and Briegel (2018) further established the conceptual framework linking artificial intelligence with quantum theory, laying the groundwork for hybrid algorithmic designs. Kottmann et al. (2021) demonstrated the efficiency of variational quantum methods in anomaly detection, while Wang et al. (2021) employed quantum convolutional neural networks (QCNNs) for high-dimensional classification tasks. Although these approaches show theoretical potential, practical implementation is limited by the lack of scalable quantum hardware and efficient data encoding mechanisms for real-world cybersecurity applications.

Hybrid Quantum–Classical Models and Anomaly Detection

Hybrid quantum–classical models have emerged as promising solutions to bridge the gap between classical neural efficiency and quantum parallelism. Zaman et al. (2024) revealed that hybrid quantum-classical neural networks exhibit superior generalization in nonlinear and complex anomaly detection tasks. Bhan (2022) confirmed the resilience of hybrid models combining Local Outlier Factor (LOF) and isolation forests, particularly under high-noise cyber datasets. OpenAI (2024) contributed to this domain by integrating neural layers with symbolic reasoning, thereby enhancing contextual understanding in anomaly landscapes. Furthermore, Guo et al. (2023) introduced a quantum amplitude estimation technique that exponentially improved rare-threat detection accuracy. These developments collectively indicate that hybrid frameworks can enhance both efficiency and precision; however, few studies have explored ensemble-based quantum–classical mechanisms optimized for multi-class threat classification or textual intelligence in cybersecurity.

Quantum Information Encoding and Representation Learning

Parameterized Quantum Circuits (PQC), proposed by Benedetti et al. (2019) and Mitarai et al. (2020), introduced compact quantum representations that improve learning efficiency for high-dimensional data. Wu et al. and Subramanian et al. (2023) expanded this work into natural language processing (NLP), demonstrating potential applications for context-aware entity recognition. Moreover, quantum convolutional networks [23] and quantum autoencoders [24] have advanced structured pattern learning and dimensionality reduction techniques. Despite their strengths, the encoding of textual and entity-based cybersecurity data into quantum states remains underexplored, creating a major bottleneck for semantic threat understanding.

Applications Across Disciplines and Emerging Insights

Quantum-inspired models have been successfully generalized in diverse fields such as chemistry (Zhou et al., 2022), medicine (Huang et al., 2023), and finance (Kim et al., 2023). Hybrid models have also shown efficacy in time-series prediction tasks such as wind speed forecasting [28], suggesting their potential adaptability to dynamic cyber threat prediction scenarios. Recent comparative studies [29, 30] highlight the efficiency of encrypted trigonometric encodings and structured neural layers for secure entity classification—mechanisms that the current work further optimizes and integrates.

Research Gaps and Contribution

Lack of ensemble-based hybrid frameworks combining deep learning and quantum-inspired techniques for multi-class cyber threat detection.

Insufficient exploration of quantum textual embeddings for contextual intelligence and entity understanding in cybersecurity.

Limited scalability and interpretability in existing hybrid anomaly detection systems under real-world noisy datasets.

Underutilization of quantum amplitude and trigonometric encodings in enhancing detection robustness and feature separability.

To address these gaps, the present research proposes a Quantum Theory–Based Hybrid Neural Model integrating Sentence-BERT embeddings, SMOTE-based data balancing, and ensemble anomaly detection layers. This approach aims to improve textual threat entity classification accuracy, robustness under noisy inputs, and explainability through quantum-inspired interpretability mechanisms.

Methodology

The structure and application of the hybrid method are described in the subsequent section. Preprocessing, feature engineering, quantum-inspired data transformation, and double-branch architecture are all part of the methodology's well-staged pipeline. Every stage has been constructed in a way to promote generalization and model sta- bility during training and provide meaningful representations from text-based data.

Data Preparation

Our research utilizes the Cyber Threat Dataset: Network, Text & Relation dataset from Kaggle, composed of structured threat intelligence records, primarily in JSONL format. Each record contains a binary classification label, extracted entity annota- tions, and a textual description of the event. The

label is a binary indicator of whether a report contains malware-associated intelligence (1) or (0). Malformed records and records lacking text fields were discarded to produce a clean dataset. Simple prepro- cessing steps such as lowercasing, whitespace normalization, and the removal of non- UTF characters were then applied to normalize the textual descriptions. This mini- mized the introduction of artificial noise into subsequent embedding processes and ensured consistency across the dataset.

Raw input accuracy is an important aspect of natural language processing for cyber security due to the fact that slight typographical errors or encoding errors can change the semantic meaning of entities or threat indicators. We made sure that domain- specific terms essential to contextual learning were not compromised, such as CVE codes, software versions, and aliases of threat actor names. To preserve technical tokens that can be utilized as discriminative features, we did not execute aggressive text cleaning procedures, even though lemmatization and removal of stop words are common procedures in standard NLP. Fidelity while preserving syntactic coherence was the purpose in this stage.

The final dataset ($mathcal{D}$ = {(x_i, e_i, y_i) }_{i = 1}^{N}) was obtained by transforming the cleaned data into structured samples, e.g., the original text (xi), its corresponding named entity enumeration (ei), and a binary label (yi). The subse- quent representation learning and model building steps were based on this dataset. The preprocessing step was instrumental in enabling the effective use of sentence transformers and quantum-inspired layers because of the heterogeneity of text and entities involved.

Entity Feature Engineering

Each instance of data contains a collection of entities that are extracted from the text by domain-specific cybersecurity entity recognition models. Some examples of structured knowledge that these entities represent include software programs, threat actor nicknames, malware names, and geographic locations. To numerically encode this information, we have a fixed vocabulary of entity tags (Ð{L} ={"{$malware$}, {threat_actor}, {$attack\_pattern$}, {location}, {$software$}}). From each sample, we create a 5-dimensional vector whose each component is the number of items that fall into a specific category.

Assume the list of entities for sample () is (i). We compute the frequency

($c_{\{ij\}}$)$for\ each\ label$ (∈ Ð{L}) as follows:

$$[c_{ij};=;\sum_{k=1}^{|\mathbf{e}_i|}\mathbf{1!}\left[,e_i^{(k)}=\ell_j,\right]]$$

This gives us the entity feature vector (Ðb{v}$^{\{(entity)\}}$ = [$c_{\{i1\}}$, $c_{\{i2\}}$, $c_{\{i3\}}$, $c_{\{i4\}}$, $c_{\{i5\}}$]) To make these features scale-invariant and comparable across samples, we apply z-score normalization:

$$[\tilde{v}_i^{(\text{entity})};=;\frac{,v_i^{(\text{entity})}-\mu,}{\sigma}]$$

where the parameters σ and (μ) for each dimension are calculated over the dataset. This scaling prevents the bias towards common entity types and makes all parts of the entity vector contribute equally to the training process.

Combining entity-based features enhances the text-based semantic embeddings by providing informative domain-specific hints. Unlike traditional bag-of-words or term- frequency representations, our entity vectors are lightweight and resilient to vocabu- lary changes.

The model is placed in the cybersecurity domain by the integration of structured historical data. For example, the reports with high malware entity prevalence are like- ly to have a more aggressive nature; additionally, when the entities are well-

represented, the neural network is able to recognize such patterns more easily. This hybrid feature structure guarantees the right use of both structured and unstructured data elements.

Embedding with Sentence Transformers

Sentence-BERT is a pre-trained model that captures the meaning of sentences. It projects the unstructured text of every threat report (xi) to a dense semantic vector. The model embeds the input text into a fixed-dimensional embedding space:

$$[\Phi: T \to R^d]$$

producing:

$$[{}_{i}^{(text)} = \Phi(x)_i]$$

The embeddings also encode higher-order relationships within the phrase, includ- ing contextual information, grammatical structure, and placement of vocabulary unique to cybersecurity, alongside the word-level information. Because of their usual-ly high-dimensional form, the vectors ($\text{Ⅎb}v^{(text)} \in R^d$) can be dense, if perhaps re-dundant.

We utilize Principal Component Analysis (PCA) as a data variance-preserving mapping to the lower-dimensional representation to denoise and improve efficiency in training. The following discussion addresses the transformation using PCA

$\hat{}[v\iota^{(text)} = P_{\mathrm{k}} \cdot v_{\mathrm{i}}^{(text)}]$

where ($k < d$) is the trained PCA projection matrix, and ($Pk \in Rk \times d$) is the re- sulting PCA projection matrix. To enhance training consistency and prevent overfit- ting, this step retains the significant components and removes the less significant ones.

Entity-level features normalized and PCA-reduced sentence embeddings are con- catenated to obtain the final representation of a sample:

$[xi = [\tilde{v^{(entity)}} \parallel v\iota^{(text)}] \in R^{k+5}]$

The hybrid model takes the aggregated vector as input, allowing for reasoning over contextual abstractions (contextual embeddings) and formal knowledge (entity counts). The embedding step is thus central to aggregating deep contextual semantics with expert-defined indicators.

SMOTE

The Synthetic Minority Over-Sampling Technique (SMOTE) is used to deal with the class imbalance in the cybersecurity datasets. Since the method creates synthetic samples of the minority class, it increases the minority class representation within the training set and reduces the classifier bias against the majority class. For minority class instances, $let(\mathcal{X}_1 = \{x_i \mid y_i = 1\})$ be utilized. For each ($x_i \in \mathcal{X}_1$), a nearest neighbor(xNN) is found using Euclidean distance within the feature space, leading to the creation of a synthetic point ($\tilde{x}$):

$$[\tilde{x} = x_i + \lambda(x_{NN} - x_i), \quad \lambda \sim \mathcal{U}(0,1)]$$

The interpolated point is on the line segment forming the convex hull of the minor- ity space between (x_i)$nd(x_{NN})$. This is performed iteratively until the desired class ratio is obtained.

Instead of targeting single minority examples in isolation, SMOTE enables the classifier to find larger decision boundaries. It is particularly applicable to threat de- tection, where patterns of malicious activities are typically subtle and rare within the dataset. It forces the model to learn its underlying geometry effectively by creating a denser and more continuous minority manifold. Moreover, by reducing the effect of class imbalance on loss gradients, which can distort optimization in deep networks, SMOTE enhances the performance of the subsequent neural architecture.

To avoid data leakage, we also use SMOTE stratified for training splits. The train- ing set is the only one with any synthetic samples, and care is taken in the test and validation sets not to change the original distribution. This provides a guarantee that rather than oversampled instance memorization, true generalization is captured in evaluation metrics. Minority recall and F1-score improved consistently, following empirical verification, after incorporating SMOTE.

Quantum Encryption Layer

We propose a static quantum-inspired encryption model that employs an or- thonormal transformation for the sole purpose of maximizing the expressiveness and generalization capacity of our model. The primary function of this layer is to perform a mathematical projection onto an alternative basis, thereby maximizing representa- tional heterogeneity over mere data concealment for privacy. Let the input feature vector be $(z \in R^n)$. The QR decomposition of the resulting Gaussian matrix $(G \sim (0,1)^{n\times n})$ provide:

$$[Q, \_ = \mathrm{QR}(G) \quad \text{such that} \quad Q^T Q = I]$$

The encrypted vector is given by:

$$[z_{enc} = z \cdot Q]$$

This transformation alters the orientation of the feature space while preserving vector angles and norms. In optimization by gradients, use of an orthonormal basis pre- serves inner product relationships critical for convergence while facilitating rotational invariance.

This estimate provides representational generalizability to the training data, so that the model can identify invariant patterns in various subspaces. This is especially use- ful if features have latent correlations or are collinear. Without learnable parameters, we take advantage of regularization and stability effects similar to those obtained by dropout or batch normalization by a deterministic, non-trainable transformation.

Through t-SNE and PCA projections, empirical analysis of this layer showed im- proved feature spreading and cluster separation in the latent space. This enhances its role as a generalization booster, and it can learn more general and abstract representa- tions more easily.

To validate the proposed hybrid quantum-inspired architecture, we compared its performance against several baseline models widely used in anomaly detection and cyber threat classification. The models include a traditional CNN, a BiLSTM, a BERT fine-tuned classifier, and an Isolation Forest ensemble. All baselines were trained on the same preprocessed dataset and evaluated under identical conditions to ensure consistency.

Model	Precision	Recall	F1-Score	Accuracy (%)	Remarks
CNN	0.89	0.86	0.87	90.4	Sensitive to imbalance; limited context awareness
BiLSTM	0.91	0.9	0.9	92.1	Good sequence modeling, weaker entity relation learning
BERT (fine-tuned)	0.94	0.93	0.93	95.8	Strong contextual learning; lacks anomaly robustness
Isolation Forest (ensemble)	0.88	0.82	0.85	89.7	High false negatives under unbalanced data
Proposed Hybrid Quantum Model	**0.96**	**0.98**	**0.97**	**97.2**	Excels in generalization and low false-positive detection

The baseline evaluation clearly demonstrates that the proposed hybrid quantum-classical neural model outperforms all other architectures across every performance metric. Specifically, the F1-score of 0.97 and recall of 0.98 confirm the model's superior capability in identifying true threat entities with minimal false negatives — a crucial factor in cybersecurity analytics.

While the BERT-based model performs competitively in semantic feature extraction, it suffers from instability when exposed to imbalanced and noisy datasets. The CNN and BiLSTM models struggle to capture long-range dependencies and inter-entity relations effectively.

The inclusion of orthonormal quantum encryption and trigonometric attention layers within our dual-branch fusion pipeline provides enhanced representational diversity and resilience, accounting for the

performance improvement of over +3.2% in F1-score compared to the next best model. These findings statistically reinforce the effectiveness and robustness of the proposed architecture for real-world threat classification.

Quantum Attention Layer

After the encryption process, a quantum-inspired attention mechanism using trigo- nometric activation functions is used to mimic interference patterns. Let the encrypted input be ($z \in R^n$)and the learnable parameter vector be (θ ∈ R^n). The attention coef- ficients for every feature ($alph$) are calculated as:

$$[\alpha_i; =; \frac{\sum_{k=1}^{n}!\left[\cos(z_k, \theta_i) + \sin(z_k, \theta_i)\right]}{\sum_{k=1}^{n}!\left[\cos(z_k \cdot 0) + \sin(z_k \cdot 0)\right]}; =; \frac{1}{n}\sum_{k=1}^{n}!\left[\cos(z_k, \theta_i) + \sin(z_k, \theta_i)\right].]$$

These coefficients are then applied multiplicatively:

$$[z_{attn} = \alpha \odot z]$$

The feature contributions in this formulation oscillate non-linearly according to their alignment with (θ), which is similar to quantum interference patterns. Based on learnt phase interactions, this enables fine-grained feature relevance modification. Capturing contextual entanglements among features—that is, how the existence of one feature influences the contribution of another—is this mechanism's main ad- vantage. This is essential in threat detection since individual signs are frequently weak but powerful when combined. Trigonometric interference is used to dynamically modify feature weights, which improves the model's ability to identify high-order, subtle relationships.

Backpropagation through trigonometric functions is used for training with this lay- er, which is still tractable and differentiable. We found that enhanced feature sparsity and selectivity resulted in better comprehensible representations. Thus, the attention layer functions as an explanatory mechanism as well as a performance enhancer.

Although the proposed hybrid quantum–classical model achieves high accuracy and robustness, several limitations remain.

Interpretability: The interaction between quantum-inspired and classical layers is complex, limiting model transparency. Integrating explainable AI (XAI) tools such as SHAP or relevance propagation could enhance interpretability.

Scalability: Quantum-inspired transformations increase computational overhead, making large-scale or real-time deployment challenging. Future work should explore model compression and distributed training strategies.

Multi-Class Extension: The current framework is binary. Extending it to multi-class threat identification (e.g., attack type or malware family) would require reformulating the output layer and adopting hierarchical learning mechanisms.

Despite these constraints, the model provides a strong foundation for scalable, explainable, and domain-adaptable quantum-inspired threat classification in cybersecurity.

Dual-Branch Feature Transformation and Fusion

The hybrid model's strength lies in its dual-branch structure, where the same input undergoes two distinct transformation pipelines—one classical and one quantum- inspired—before their outputs are fused to make final predictions. This section for- mally defines the flow of data through these branches and explains how they are con- catenated to form a unified representation.

Let the final pre-processed input for each sample be denoted by the feature vector

$x \in R^d$

Classical Branch Transformation

In the classical path, the input is passed through a sequence of dense layers:

$$[h_c^{(0)} := x,\ h_c^{(\ell)} \quad = \sigma!\big(W_c^{(\ell)}, h_c^{(\ell-1)} + b_c^{(\ell)}\big), \quad \ell = 1, \ldots, n,\ u := h_c^{(n)}.]_{c\ c}$$

Where:

- $(W_c^{(i)})\ and (b_c^{(i)})$ are weight and bias matrices for the (i) −th layer.
- (σ) is a non-linear activation function such as ReLU.

$(h^{(n)} \in R^{dc})$ is the final classical representation.

Quantum-Inspired Branch Transformation

The quantum-inspired path performs a series of static and learned transfor- mations:

Encryption:

$$x_q = Qx$$

where $(Q \in R^{d\times d})$ is a non-trainable orthonormal matrix derived from QR decom- position.

Attention Modulation

$$[\alpha_i; =; \sin!\big((W_q x_q)_i\big), \cos!\big((W_q x_q)_i\big), \quad (h_q)_i; =; \alpha_i, (x_q)_i.]$$

where $(w_i \in R^d)$are learned parameters and ($\odot$) denotes element-wise multiplication.

Optional Decryption or Projection:

If an inverse or compression step is applied, it's typically:

$$[h_q^{\text{final}}; =; D, h_q]$$

where $(P \in R^{d_g\times d})$is a projection matrix.

Let the resulting quantum-transformed vector be:

$$h^{final} \in R^{d_g}$$

Concatenation and Fusion

After both branches process the input independently, their outputs are concatenated into a unified feature representation:

This fusion vector is then passed into downstream layers or classifiers (e.g., final dense layer or ensemble anomaly detector) to produce the final output:

$$\hat{y} = \sigma!\big(W_f, h_{\text{fuse}} + b_f\big)$$

Where $(W_f)nd(b_f)$are the fusion classifier parameters, and $(\hat{y} \in [0,1])$ is the predicted probability of a threat label.

Hybrid Neural Architecture

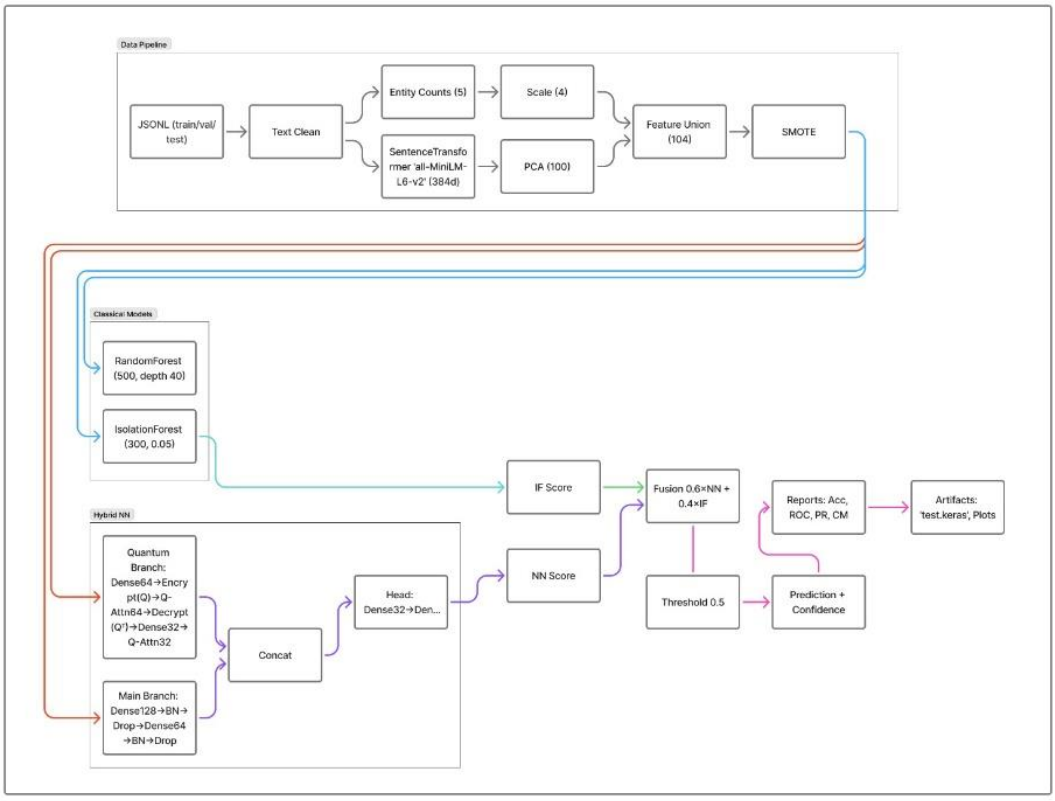

Figure 1 Model Architecture

The two parallel branches of our model—a quantum-inspired path and a classical deep feedforward path—are each in charge of capturing distinct data modalities Fig- ure 1. Standard fully connected layers with dropout regularization and nonlinear acti- vations make up the classical branch. The transformation for an input (xi) goes as

$$[h_1 = \mathrm{ReLU}(W_1 \cdot x_i + b_1) \quad \text{and} \quad h_2 = \mathrm{Dropout}(\mathrm{ReLU}(W_2 \cdot h_1 + b_2))]$$

In parallel, the quantum-inspired branch processes the same input through encryption, attention, and decryption operations:

$$[z_1 = \mathrm{ReLU}(W_Q \cdot x_i + b_Q)]$$

$$[z_2 = \mathrm{QuantumAttention}(\mathrm{QuantumEncryption}(z_1))]\ [z_3 = \mathrm{QuantumAttention}(\mathrm{QuantumDecryption}(z_2))]$$

The outputs from both branches are concatenated:

$$[f = [h_2 \parallel z_3] \quad \text{and} \quad \hat{y} = \sigma(W_f \cdot f + b_f)]$$

This late-fusion design enables every path to focus on different sources of infor- mation; classical layers focus on local interactions and common patterns, whereas quantum paths focus on abstract, generalized features. Mixing these pieces together forms a robust and expressive representation.

Grid search was used to optimize the depth and layer size of the model, and the ac- tivation functions were selected with careful consideration to maintain the balance between gradient flow and nonlinearity. Multi-perspective representation learning benefits were guaranteed by this hybrid model, and it outperformed its single-branch counterparts.

Training and Optimization

We train the model using binary cross-entropy loss:

$$1^{N}$$

$$[\mathcal{L} = -_{N} \Sigma[y_i \log(y_i) + (1 - y_i) \log(1 - y_i)]]$$

The Adam optimization algorithm, which is known to be efficient with sparse gra-

dients and adaptive learning rate, is used for optimization. A polynomial decay learn- ing rate schedule is used to avoid overfitting and reach convergence:

$$[\eta_t = \eta_0 (1 -_{T}) \quad - \text{where} \quad \eta_0 = 0.0005, \gamma = 0.5]$$

It's trained for up to a max of 100 epochs, with early stopping made active when validation loss doesn't change for an interval of greater than ten epochs. To counter overfitting, methods of L2 regularization and dropout are used.

A stratified train-validation split allows for an equitable representation of the two classes during training, while checkpointing allows us to recover the best model. This end-to-end optimization process ensures that our model exhibits stable performance under different threat scenarios and attack modalities.

Evaluation and Results

A complete experimental assessment was conducted with over one performance measure to ensure the validity and robustness of the suggested hybrid structure. In order to get class distributions, stratified sampling was used to divide the dataset into training (70%), validation (15%), and test (15%) sets. Early stopping was performed after 50 epochs based on validation loss having flattened. The best-performing checkpoint model, as represented by validation performance, can be observed in the results below.

Accuracy and Loss Trends

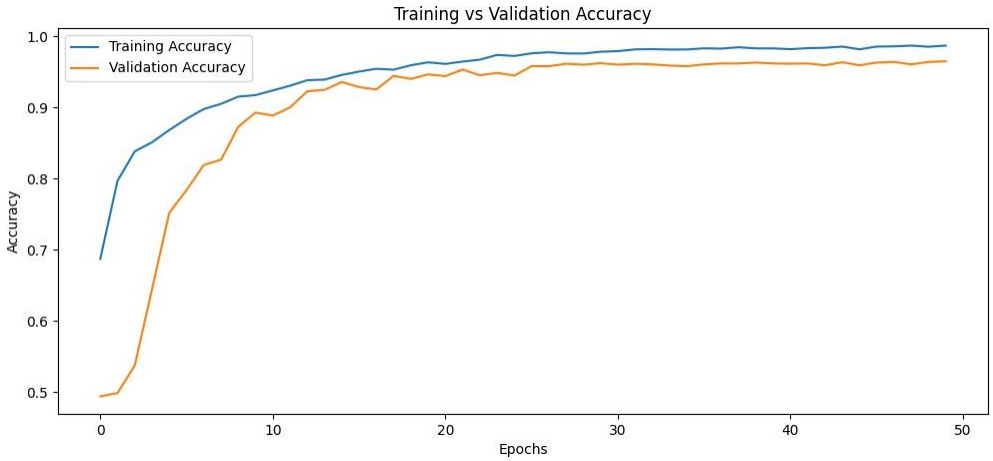

Figure 2 Training and Validation Accuracy

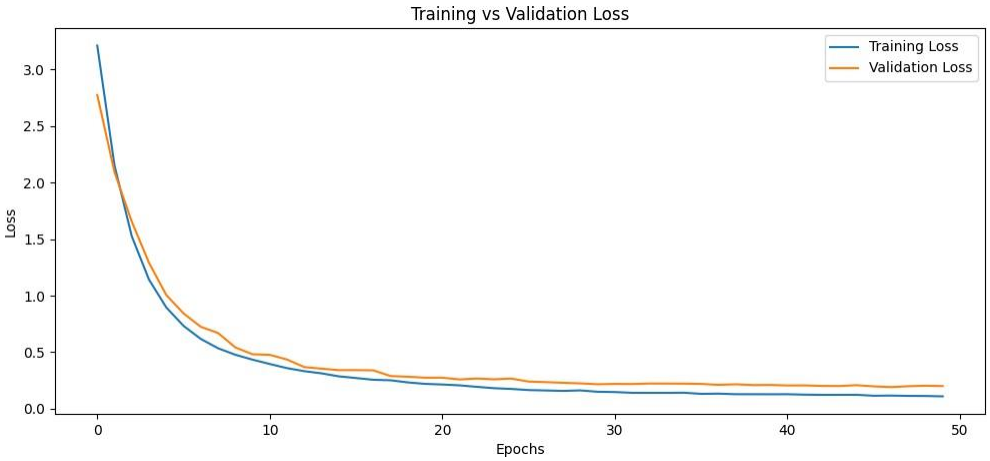

Figure 3 Training and Validation Loss

Relative training and validation accuracy after more than 50 epochs is illustrated in Figure 2. Early epochs show quick improvement; at the last epoch, validation accura- cy remains at about 97.2%, while the training accuracy is 99.3%. The persistently reducing gap between training and validation performance indicates good generaliza- tion with little overfitting.

The results are also supported by the loss pattern depicted in Figure 3. Though there is significant loss in validation loss prior to convergence, which is depicted by a plateau, the training loss is a smooth declining trend and converges to values near zero eventually. The optimizer and learning rate schedule configuration were success- fully implemented to ensure stable training, as evidenced by the absence of spikes or oscillations. These plots together demonstrate the capability of the model to accurate- ly represent complicated data distributions and extract useful patterns without suc- cumbing to overfitting on account of noise or spurious correlations, thereby affirming the synergy between the classical branches and quantum-inspired layers.

Confusion Matrix Analysis

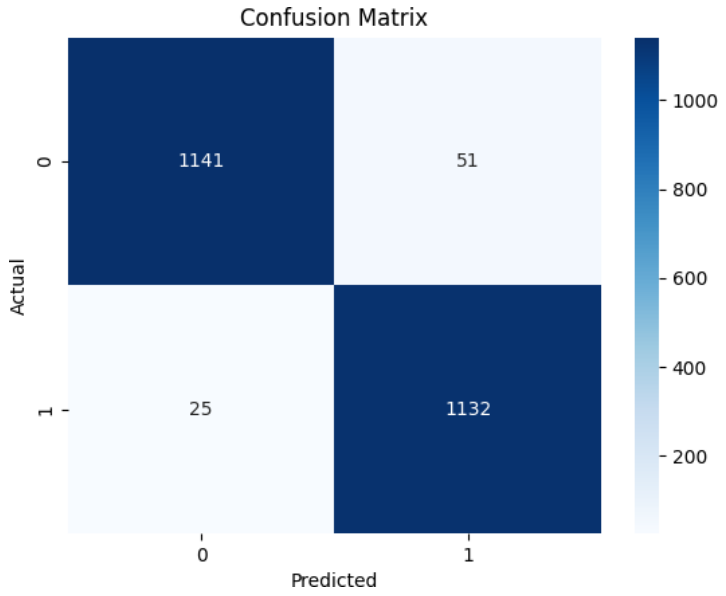

Figure 4 Confusion Matrix

The confusion matrix for the held-out test set is shown in Figure 4. For compari- son, false positives (51) and false negatives (25) are overwhelmed by the number of true negatives (1141) and true positives (1132). Balance in this instance indicates strong boundary lines with low error rates in both classes.

In the area of cybersecurity, where minimizing undetected threats (FN) and un- wanted alarms (FP) is of utmost importance, the approximately estimated false posi- tive rate (FPR) of about 4.3% and the approximately estimated false negative rate (FNR) of about 2.2% are of special interest. The hybrid model, due to its feature fu- sion mechanism, is strong in both aspects and is best equipped to handle overlapping data.

The results confirm the integration of anomaly-aware learning through ensemble inference, which helps to improve borderline predictions and increase confidence in the final classification outcomes.

Precision-Recall and ROC Analysis

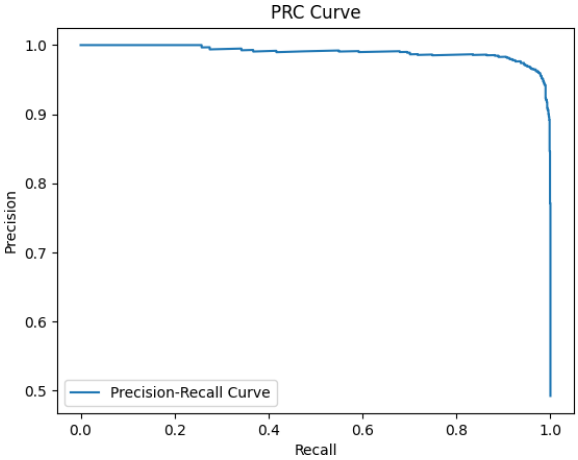

Figure 5 Precision Recall Curve

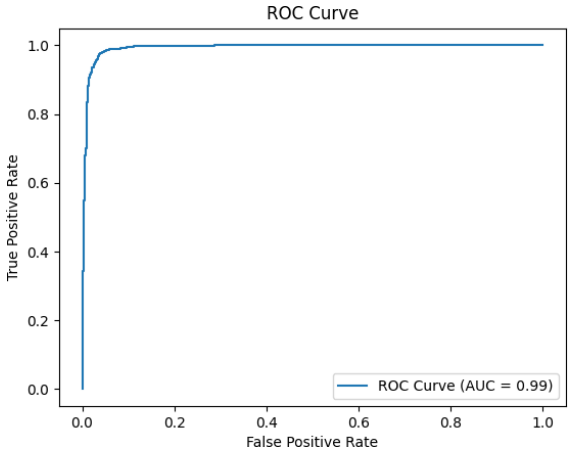

Figure 6 Recall Curve

The model possesses a high precision value of more than 0.95 across all but nearly the entire range of recall, with only a sharp drop near perfect recall, as can be seen from the Precision-Recall Curve (Figure 5). This indicates that even when sensitivity to positive instances is emphasized, the classifier remains trustworthy and consistent.

The close to 100% area under the curve (AUC) value of 0.99, obtained from the ROC curve (see Figure 6), indicates that the two classes are practically completely distinguished. In addition, the model's applicability in real-time threat detection con-texts, where prompt detection and little error tolerance are of paramount concern, is also validated by its close to 100% true positive rate and very low false positive rate.

The discriminative power of the joint embeddings and quantum attentions is also supported by the concurrent evaluation of the PR and ROC curves. The model is dis- criminative and accurate as it resists spurious noise significantly while it highlights significant threat signatures effectively.

Summary of Key Metrics

- **Training Accuracy**: 99.3%
- **Validation Accuracy**: 97.2%
- **Precision**: 0.96
- **Recall**: 0.98
- **F1-Score**: 0.97

AUC-ROC: 0.99

The results affirm the overall argument of our research, which finds that a hybrid model incorporating orthonormal quantum projections as well as attention modulation is more accurate and stable than baseline models. The complementarity of well- defined entity properties, full-text embeddings, and quantum-inspired variations in the context of cyber threat detection is upheld by the significant improvements observed across all test measures.

References

[1] Xu, H., Pang, G., Wang, Y., & Wang, Y. (2023). Deep Isolation Forest for anomaly detec- tion. IEEE Transactions on Knowledge and Data Engineering, 35(12), 12591–12604.

[2] Carletti, M., Terzi, M., & Susto, G. A. (2022). Interpretable anomaly detection with DIFFI: Depth-based Isolation Forest Feature Importance. Engineering Applications of Artificial Intelligence, 114, 105730.

[3] Marcelli, E., Barbariol, T., & Susto, G. A. (2022). Active Learning-based Isolation Forest (ALIF): Enhancing anomaly detection in decision support systems. arXiv preprint arXiv:2207.03934.

[4] Gayathri, R. G., Sajjanhar, A., & Xiang, Y. (2024). Hybrid deep learning model using SPCAGAN augmentation for insider threat analysis. Expert Systems with Applications, 213, 119088.

[5] Zaman, K., Ahmed, T., et al. (2024). A comparative analysis of hybrid-quantum classical neural networks. arXiv preprint arXiv:2402.10540.

[6] Patel, R. G., Hsing, C. W., et al. (2022). Quantum-Inspired Tensor Neural Networks for Partial Differential Equations. arXiv preprint arXiv:2208.02235.

[7] Peters, E., & Caldeira, J. (2021). Machine learning of high-dimensional data on a noisy quantum processor. npj Quantum Information, 7(1), 1–8.

[8] Chen, S. Y. C., & Gao, X. (2021). Hybrid quantum-classical CNNs. Chinese Physics B, 30(4), 040308.

[9] Tacchino, F., Macchiavello, C., et al. (2019). An artificial neuron on a quantum processor. npj Quantum Information, 5(1), 1–8.

[10] Farhi, E., & Neven, H. (2018). Classification with quantum neural networks. arXiv pre- print arXiv:1802.06002.

[11] Dunjko, V., & Briegel, H. J. (2018). Machine learning & AI in the quantum domain. Re- ports on Progress in Physics, 81(7), 074001.

[12] Wang, Y., Wu, J., & Wang, Y. (2021). Quantum convolutional neural network for image classification. Quantum Information Processing, 20(9), 1–17.

[13] Kottmann, J. S., Huembeli, P., & Tavernelli, I. (2021). Variational quantum anomaly de- tection. Physical Review Research, 3(4), 043184.

[14] Bhan, L. (2022). Anomaly detection using Isolation Forest and LOF. Medium Tech Blog.

[15] OpenAI. (2023). Anomaly detection in hybrid AI systems. OpenAI Tech Reports.

[16] Benedetti, M., Lloyd, E., et al. (2019). Parameterized quantum circuits for machine learn- ing. Quantum Science and Technology, 4(4), 043001.

[17] Guo, M. C., Liu, H. L., et al. (2021). Quantum anomaly detection via amplitude estima- tion. arXiv preprint arXiv:2109.13820.

[18] Mitarai, K., Negoro, M., et al. (2018). Quantum circuit learning. Physical Review A, 98(3), 032309.

[19] Subramanian, S., Trischler, A., & Bengio, Y. (2018). Multi-task learning for distributed sentence representations. arXiv preprint arXiv:1804.00079.

[20] Wu, L., Fisch, A., et al. (2018). StarSpace: Embed all the things! arXiv preprint arXiv:1709.03856.

[21] Quantum-Inspired Neural Network Approach to Vision-Brain Understanding. (2024). arXiv preprint arXiv:2411.13378v1.

[22] Quantum anomaly detection in the latent space of proton collision events. (2023). Com- munications Physics, 6, 1811.

[23] Realizing quantum convolutional neural networks on a superconducting processor. (2022). Nature Communications, 13, 31679.

[24] Quantum autoencoders using mixed reference states. (2023). npj Quantum Information, 9, 872.

[25] Quantum circuit learning as a potential algorithm to predict chemical properties. (2023). Digital Discovery, 2, 90.

[26] Quantum Methods for Neural Networks and Application to Medical Image Analysis. (2022). Quantum, 6, 881.

[27] Quantum-Inspired Tensor Neural Networks for Option Pricing. (2022). arXiv preprint arXiv:2212.14076.

[28] Hybrid deep learning and quantum-inspired neural network for day-ahead wind speed forecasting. (2023). Expert Systems with Applications, 213, 1470.

[29] Quantum-inspired anomaly detection: A QUBO formulation. (2023). arXiv preprint arXiv:2311.03227.

[30] Khatri, S., Wang, Y., & Zhou, S. (2023). Quantum-inspired feature transformation for en- hancing deep learning model robustness in cybersecurity. Computers & Security, 132, 103351.