engineering proceedings

MDPI

Proceeding Paper

Adaptive Multi-Embedding Quantum Feature Fusion for Intrusion Detection Systems †

Raid Anis Kerkatou [1,2,*], Hacene Belhadef [3], Aicha Eutamene [4] and Svetlana Stefanova [5]

[1] Department of Computer Science, University of Mujahid Abdelhafid Boussouf—Mila, Mila 43000, Algeria
[2] MISC Laboratory, University of Constantine 2 Abdelhamid Mehri, Constantine 25016, Algeria
[3] Department of Fundamental Computer Science and its Applications, University of Constantine 2 Abdelhamid Mehri, Constantine 25016, Algeria; hacene.belhadef@univ-constantine2.dz
[4] Department of Biochemistry and Cellular and Molecular Biology, University of Constantine 1 Frères Mentouri, Constantine 25000, Algeria; aicha.eutamene@umc.edu.dz
[5] Faculty of Electrical Engineering, Electronics and Automatics, University of Ruse "Angel Kanchev", 7004 Ruse, Bulgaria; svetlaps@ruse.bg
* Correspondence: kerkatouraidanis@centre-univ-mila.dz
† Presented at the International Conference on Electronics, Engineering Physics and Earth Science (EEPES2026), Bandirma, Turkey, 24–27 June 2026.

Abstract

Quantum Machine Learning (QML) has recently emerged as a promising approach for cybersecurity applications due to its ability to encode data into high-dimensional Hilbert spaces. However, most QML-based intrusion detection systems rely on a single quantum embedding strategy, limiting representation diversity. This paper proposes a Multi-Embedding Quantum Ensemble (ME-QE) framework that combines Angle Embedding and IQP Embedding using concatenation and weighted fusion strategies. The proposed approach is evaluated on the NSL-KDD dataset for binary intrusion detection using stratified 5-fold cross-validation. Experimental results show that the best fusion configuration achieves 94.75% accuracy, while the classical Random Forest baseline reaches 97.50%. The analysis further demonstrates that Angle Embedding contributes more effectively to classification performance than IQP Embedding in this context. In addition, the study highlights important scalability limitations of quantum kernel methods due to the computational cost of pairwise kernel evaluation. The proposed framework provides insights into hybrid quantum–classical intrusion detection and establishes a foundation for future scalable quantum cybersecurity architectures.

Keywords: quantum machine learning; intrusion detection system; NSL-KDD; quantum kernel; feature fusion

Check for updates

Academic Editors: Teodor Iliev, Ivaylo Stoyanov, Grigor Mihaylov and Selahattin Kosunalp

Published: 28 August 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Intrusion Detection Systems (IDS) play a critical role in safeguarding modern network infrastructures against increasingly sophisticated cyber threats. With the exponential growth of network traffic and attack complexity, traditional rule-based systems have become insufficient, leading to the widespread adoption of machine learning techniques. Classical models such as Random Forest (RF) and Support Vector Machines (SVM) have demonstrated strong performance on benchmark datasets such as NSL-KDD. However, these models often struggle to capture highly complex nonlinear relationships inherent in network traffic data.

Quantum computing offers a fundamentally different computational paradigm based on quantum mechanical principles such as superposition, entanglement, and interference. Unlike classical bits, quantum bits (qubits) can exist in a superposition of states, enabling the representation of exponentially large state spaces. Entanglement further allows correlations between qubits that cannot be efficiently reproduced classically, potentially enabling richer feature representations and more expressive learning models [1].

Quantum Machine Learning (QML) leverages these principles to enhance data representation and model capacity. In particular, quantum feature maps encode classical data into quantum states in a high-dimensional Hilbert space, where similarities between data points can be evaluated through quantum kernels. This approach has shown promise in classification tasks, including cybersecurity applications [2].

Despite these advantages, most existing QML-based IDS approaches rely on a single embedding strategy, such as Angle Embedding or IQP Embedding [3]. This limitation can restrict the representational capacity of quantum models, especially when dealing with heterogeneous and complex network traffic features.

To address this limitation, we propose a Multi-Embedding Quantum Ensemble (MEQE) framework that combines multiple quantum feature maps through concatenation and weighted fusion strategies. By integrating complementary embeddings, the proposed method aims to enhance feature representation and improve classification robustness.

The main contributions of this paper are as follows:

- A multi-embedding quantum feature fusion framework for intrusion detection;
- A systematic analysis of the fusion parameter α;
- A comparative evaluation against classical baselines;
- An investigation of computational complexity and scalability limitations.

The remainder of this paper is organized as follows. Section 2 reviews related work in classical and quantum intrusion detection systems. Section 3 presents the proposed methodology, including data processing and quantum feature mapping. Section 4 describes the experimental setup and evaluation protocol and analyzes the results. Section 5 discusses limitations and practical implications. Finally, Section 6 concludes the paper.

2. Related Work

Intrusion detection has been extensively studied using classical machine learning techniques. Traditional approaches such as Random Forest, Support Vector Machines, and ensemble methods have demonstrated strong performance due to their ability to handle structured tabular data [4]. Gradient boosting techniques, including XGBoost, further improved detection accuracy through optimized tree-based learning.

Deep learning methods have also been explored, particularly Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) models, which capture temporal dependencies in sequential data [5]. However, their effectiveness on tabular datasets such as NSL-KDD remains limited due to the lack of inherent temporal structure.

Recent studies have explored advanced hybrid learning and representation fusion strategies for intelligent systems. For example, adaptive multi-embedding fusion approaches have recently been investigated for sequential recommendation tasks [6]. Other recent works explored Mamba-based multimodal fusion architectures and hybrid representation learning frameworks for improving feature interaction and representation quality in complex learning problems [7–10]. These studies demonstrate the growing interest in combining multiple embedding spaces and heterogeneous feature representations, which further motivates the proposed multi-embedding quantum fusion framework.

Despite these advancements, existing QML-based IDS approaches predominantly utilize a single embedding strategy. This limitation restricts their ability to capture diverse

https://doi.org/10.3390/engproc2026154006

feature interactions. Moreover, the impact of combining multiple quantum embeddings remains largely unexplored.

In addition, scalability remains a significant challenge in QML. Quantum kernel methods require the computation of pairwise similarities, leading to quadratic complexity with respect to the number of samples [11]. This limitation hinders their application to large-scale datasets.

To address these gaps, this work proposes a multi-embedding quantum fusion framework and provides a systematic analysis of its performance and computational trade-offs.

3. Materials and Methods

The proposed Adaptive Multi-Embedding Quantum Ensemble (AME-QE) framework integrates classical and quantum processing in a hybrid architecture. The overall pipeline, illustrated in Figure 1, consists of data preprocessing, dimensionality reduction, parallel quantum feature mapping, kernel fusion, and classification.

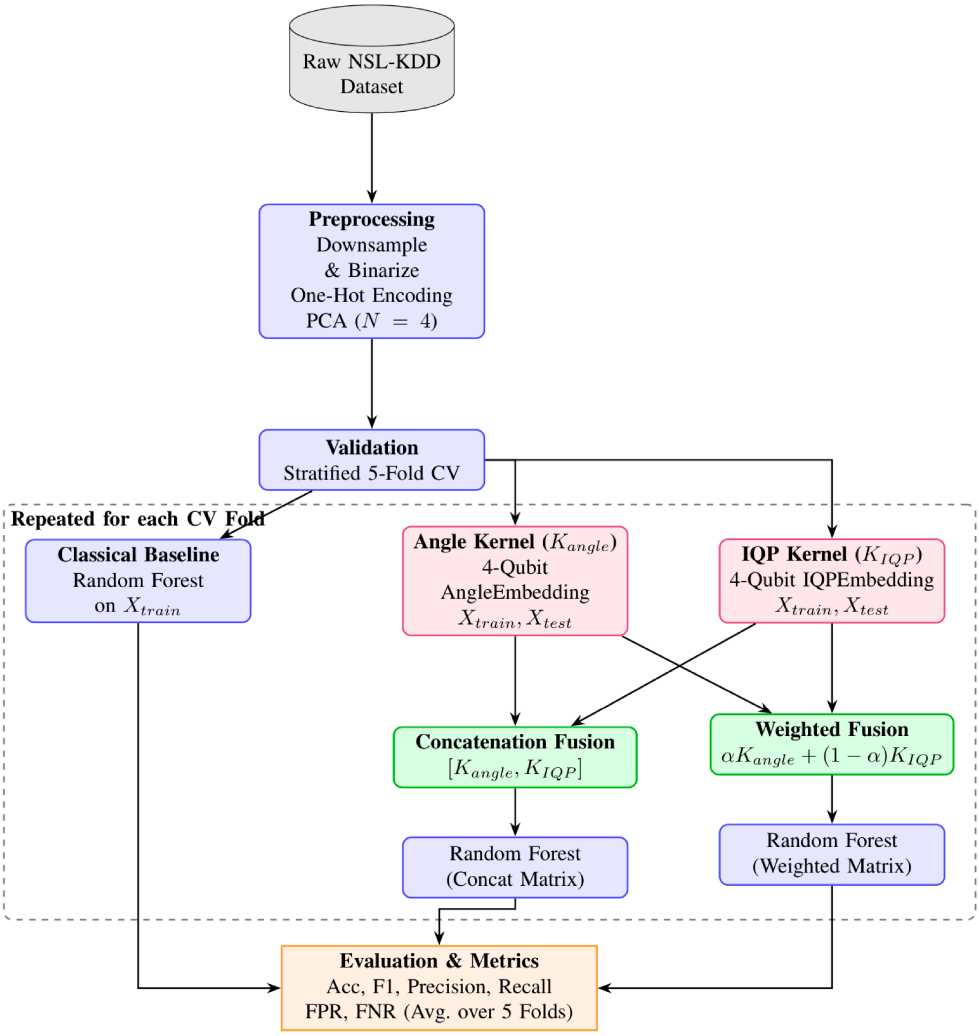


Figure 1. Architecture of the proposed Adaptive Multi-Embedding Quantum Ensemble (AME-QE) framework integrating classical preprocessing, parallel quantum feature mapping, and fusion strategies.

The adaptive aspect of the proposed framework is associated with the controllable fusion mechanism, where the contribution of different quantum embeddings can be adjusted through the parameter α. By varying α, the framework adaptively balances the influence of

Angle and IQP embeddings according to the observed classification behavior and feature representation characteristics.

3.1. Dataset and Preprocessing

The experiments were conducted using the NSL-KDD dataset, a widely used benchmark for intrusion detection research [12]. The dataset contains both normal and malicious network traffic records described by numerical and categorical features.

For the binary classification setting, all attack categories were grouped into a single attack class, while legitimate traffic was labeled as normal. Due to the computational cost of quantum kernel evaluation, a balanced subset of 400 samples was selected for experimentation.

Categorical attributes were transformed using one-hot encoding, followed by feature standardization. Principal Component Analysis (PCA) was then applied to reduce the dimensionality to four components, matching the number of qubits used in the quantum circuits.

The final dataset distribution consisted of approximately equal numbers of normal and attack samples to avoid class imbalance during training and evaluation.

3.2. Data Processing and Dimensionality Reduction

The NSL-KDD dataset is preprocessed through one-hot encoding of categorical features followed by feature standardization. Principal Component Analysis (PCA) is applied to reduce the dimensionality of the data to match the number of qubits in the quantum circuits. In this study, a reduced subset of 400 samples is used due to computational constraints associated with quantum kernel evaluation.

3.3. Variational Quantum Circuit Architecture

Variational Quantum Circuits (VQCs) are a fundamental building block in quantum machine learning. A VQC consists of three main components: data encoding, a parameterized quantum circuit, and measurement.

In this architecture Figure 2, classical data is first encoded into quantum states using a feature map. A parameterized quantum circuit U_Θ then transforms the state through a sequence of quantum gates. Finally, measurements are performed to extract classical outputs, which can be used for classification or regression tasks.

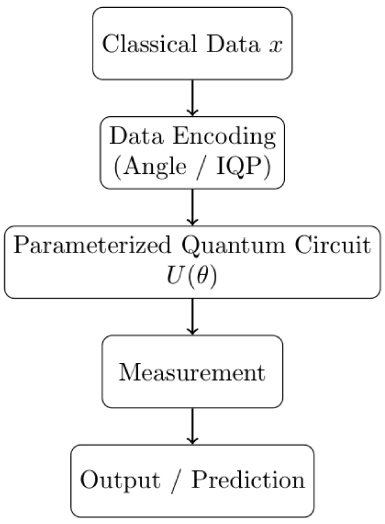


Figure 2. Standard Variational Quantum Circuit (VQC) architecture used in quantum machine learning.

3.4. Quantum Feature Mapping

Two quantum embedding strategies are employed:

Angle Embedding: Encodes classical features as rotation angles in quantum circuits, providing a straightforward mapping.

IQP Embedding: Introduces entanglement between qubits, enabling more complex feature interactions.

For each embedding, a quantum kernel matrix is computed by evaluating the fidelity between quantum states.

3.5. Multi-Embedding Fusion

The proposed fusion strategy combines the two kernel matrices as follows:

$$K_{fusion} = \alpha K_{angle} + (1 - \alpha) K_{iqp}$$

where α in [0,1] controls the contribution of each embedding.

Additionally, a concatenation-based fusion strategy is evaluated for comparison.

The adaptive behavior of the proposed framework arises from the ability to regulate the influence of each embedding representation through the fusion parameter. Different α values are experimentally evaluated to analyze how the balance between embeddings affects classification robustness and detection performance.

3.6. Classification

The fused quantum kernel representations are used as input features for a Random Forest (RF) classifier. In this framework, RF operates as a classical learning model applied to transformed quantum feature representations generated from the fusion of Angle and IQP kernels.

Random Forest was selected for several reasons. First, RF has demonstrated strong robustness and generalization performance on structured cybersecurity datasets such as NSL-KDD. Second, ensemble tree methods are effective at capturing nonlinear decision boundaries while remaining relatively stable on small datasets. Finally, RF provides an efficient mechanism for evaluating whether the fused quantum representations contain discriminative information that can improve intrusion detection performance.

For comparison purposes, a classical RF model is also trained directly on the original preprocessed feature space to serve as a baseline.

3.7. Computational Considerations

Due to the lack of large-scale quantum hardware, all quantum computations are performed using the PennyLane simulator version 0.45.1 [13]. The computational complexity of kernel evaluation scales as $O\left(\mathrm{N}^{(2)}\right)$, which necessitates the use of a reduced dataset size.

4. Results and Discussion

In this section we present the global results and deep discussion with comparison.

4.1. Overall Performance

The results indicate that the classical Random Forest achieves the highest overall performance. However, the proposed fusion approach remains competitive, particularly for α = 0.8, suggesting that combining quantum embeddings can enhance classification robustness.

4.2. Multi-Metric Comparative Analysis

Figure 3 provides a comprehensive visualization of the performance across multiple evaluation metrics. Unlike Table 1, which summarizes key indicators, this figure highlights the consistency and trade-offs between models.

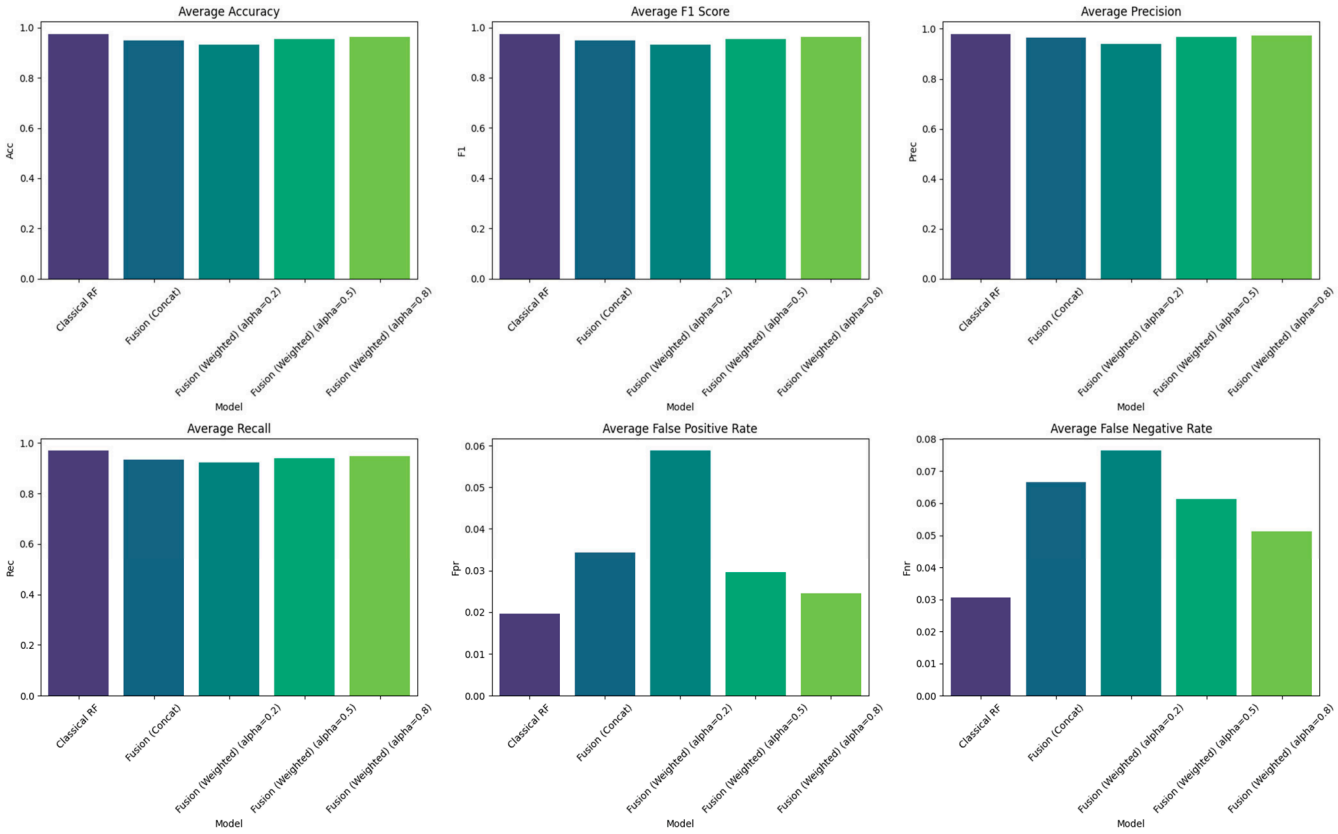


Figure 3. Comparison of average performance metrics across models: Accuracy, F1-score, Precision, Recall, False Positive Rate (FPR), and False Negative Rate (FNR).

Table 1. Average performance across the 5-fold cross-validation.

Model	Acc	F1	FPR	FNR
Classical RF	0.9750	0.9744	0.0196	0.030
Fusion (Concat)	0.9425	0.9403	0.0391	0.0765
Fusion $\alpha = 0.2$	0.9250	0.9215	0.0490	0.1019
Fusion $\alpha = 0.5$	0.9375	0.9341	0.0343	0.0919
Fusion $\alpha = 0.8$	0.9475	0.9450	0.0294	0.0762

From the figure, several important observations can be made:

4.2.1. Consistency of Classical RF

The Classical Random Forest maintains superior performance across all metrics, achieving near-optimal Accuracy, F1-score, Precision, and Recall, while simultaneously minimizing both FPR and FNR. This confirms its robustness and stability in low-dimensional tabular data, which is consistent with prior studies showing that ensemble tree methods remain highly competitive for structured cybersecurity datasets.

4.2.2. Performance Degradation in Fusion Models

All fusion strategies exhibit a noticeable drop across all metrics, particularly in Recall and F1-score. This degradation suggests that the integration of quantum embeddings

introduces additional variance without sufficiently improving feature separability. This aligns with findings in quantum machine learning literature, where hybrid models often struggle to outperform classical baselines on small datasets due to noise sensitivity and limited expressivity [11].

4.2.3. Effect of Weighted Fusion α

The figure clearly shows a monotonic improvement as α increases from 0.2 to 0.8:

- Accuracy and F1-score increase steadily.
- FPR and FNR decrease correspondingly.

This indicates that Angle Embedding contributes more effectively to classification performance than IQP Embedding. Similar observations have been reported in [14], where simpler encoding strategies were found to be more stable for classical data distributions.

4.2.4. False Positive vs. False Negative Behavior

A critical insight from Figure 3 is the imbalance introduced by fusion models:

- Lower α values significantly increase FNR (missed attacks).
- Higher α values reduce FNR but still do not reach classical RF performance.

In intrusion detection systems, minimizing FNR is crucial, as missed attacks have higher security risks than false alarms. This suggests that fusion strategies must be carefully tuned to prioritize recall over precision, as also emphasized in cybersecurity literature.

4.2.5. Precision Stability Across Models

Precision remains relatively stable across all models, indicating that when attacks are detected, they are generally classified correctly. However, the drop in Recall implies that fusion models fail to detect a significant portion of attacks, reducing overall system reliability.

4.2.6. Insight into Quantum Feature Representations

The results suggest that quantum feature maps do not necessarily enhance separability in low-dimensional tabular datasets. This supports recent theoretical work showing that quantum kernels may not provide advantage unless specific data structures or high-dimensional Hilbert spaces are effectively exploited.

Overall, Figure 3 reinforces that while hybrid quantum–classical fusion introduces diversity, it does not yet translate into superior predictive performance under current constraints (small dataset, simulation environment, and limited qubit capacity).

4.3. Confusion Matrix Analysis

Figures 4–6 present the confusion matrices for the last fold.

The Classical RF achieves perfect classification in the last fold (41 true negatives and 39 true positives), with zero false positives and false negatives. This indicates excellent separability in the reduced feature space.

In contrast, the fusion models introduce misclassifications. For the concatenation-based fusion, 3 false positives and 3 false negatives are observed, indicating balanced but degraded performance. The weighted fusion $\alpha = 0.5$ reduces false positives (2 cases) but slightly increases false negatives (4 cases), suggesting a shift toward conservative detection.

https://doi.org/10.3390/engproc2026154006

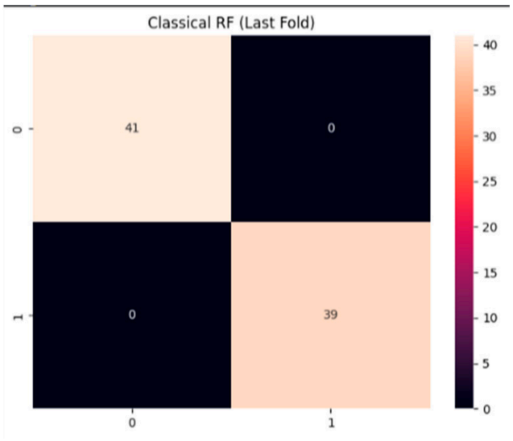


Figure 4. Confusion Matrix—Classical RF.

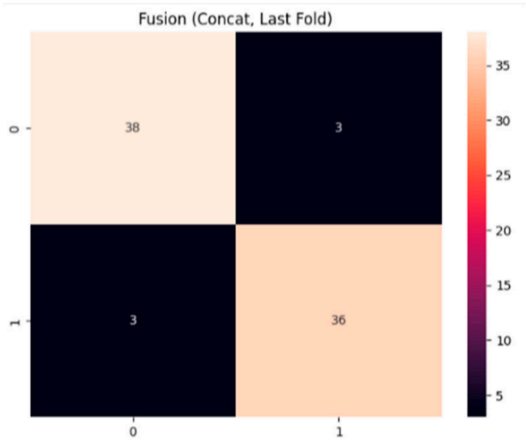


Figure 5. Confusion Matrix—Fusion (Concat).

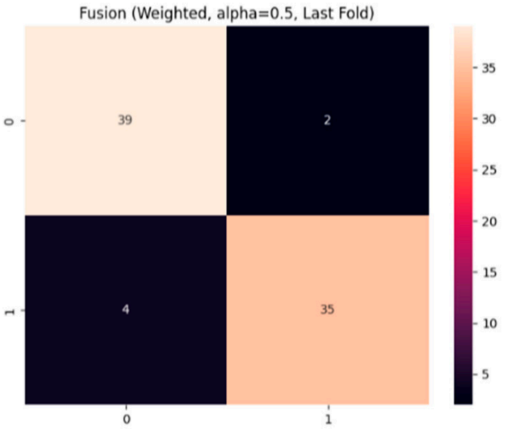


Figure 6. Confusion Matrix—Fusion $\alpha = 0.5$.

4.4. False Positive and False Negative Trade-Off

A detailed analysis of FPR and FNR reveals important trade-offs:

- Classical RF maintains the lowest FPR (0.0196) and FNR (0.0306), ensuring both high detection accuracy and low false alarm rate.
- Fusion methods exhibit increased FNR, particularly for lower α values, indicating a higher rate of missed attacks.
- Increasing α improves recall and reduces FNR, confirming that Angle Embedding contributes more effectively to attack detection.

https://doi.org/10.3390/engproc2026154006

This trade-off highlights the sensitivity of quantum models to feature representation and fusion strategy.

Impact of Fusion Parameter α

The performance varies significantly with α, specifically:

- α = 0.2 results in the worst performance, indicating that IQP Embedding alone is insufficient.
- α = 0.5 provides moderate performance but does not outperform concatenation.
- α = 0.8 achieves the best fusion performance, suggesting that Angle Embedding dominates the feature representation.

This demonstrates that the fusion weight parameter plays a critical role in balancing the contribution of the two quantum embeddings.

4.5. Computational Complexity

A major limitation of the proposed approach lies in computational cost. While classical RF training requires less than 1 s, quantum kernel computation takes approximately 550 s per fold.

This overhead arises from the quadratic complexity of kernel evaluation:

$$O\left(N^2\right)$$

To address this limitation, experiments were conducted on a reduced dataset of 400 samples. This highlights the current scalability challenges of QML methods when applied to real-world intrusion detection scenarios.

4.6. Discussion

The experimental results, supported by both quantitative metrics Table 1 and visual analysis Figure 3, provide several key insights into the effectiveness of quantum–classical fusion for intrusion detection.

First, the Classical Random Forest consistently outperforms all hybrid approaches. This result is not unexpected, as ensemble learning methods are well-known for their strong performance on tabular datasets with limited samples. The ability of RF to capture non-linear feature interactions without requiring complex feature transformations gives it a clear advantage.

Second, while fusion models introduce richer feature representations through quantum embeddings, they fail to translate this increased expressivity into improved classification performance. This can be attributed to several factors:

1. The limited size of the dataset;
2. The noise and approximation errors introduced by quantum simulation;
3. The mismatch between quantum feature maps and classical data distributions. Similar limitations have been highlighted in recent QML studies [14].

Third, the analysis of the fusion parameter α reveals that Angle Embedding plays a dominant role in performance. As α increases, all metrics improve, indicating that simpler encoding strategies are more effective for this task. This observation is consistent with prior work suggesting that highly entangled quantum circuits do not necessarily provide advantages for classical machine learning problems [7].

Fourth, the trade-off between FPR and FNR is particularly critical in the context of intrusion detection. Fusion models tend to increase FNR, which is undesirable since missed attacks pose significant security risks. This highlights the need for cost-sensitive optimization strategies in future work.

https://doi.org/10.3390/engproc2026154006

Although the proposed AME-QE framework demonstrates the feasibility of integrating multiple quantum embeddings for intrusion detection, the current implementation should primarily be considered an exploratory proof-of-concept rather than a production-ready cybersecurity solution. Real-world intrusion detection systems typically operate on large-scale streaming traffic with strict latency and scalability requirements, whereas the present study relies on simulation-based quantum kernel evaluation on a reduced subset of the NSL-KDD dataset.

Nevertheless, the framework provides valuable insights into the behavior of hybrid quantum–classical feature fusion mechanisms and highlights important research directions for future scalable quantum intrusion detection architectures.

Finally, the computational complexity remains a major limitation. The quadratic scaling of kernel-based quantum methods severely restricts their applicability to large-scale datasets. This scalability issue is widely recognized as one of the main challenges in practical quantum machine learning.

In summary, while the proposed fusion framework demonstrates the feasibility of integrating multiple quantum embeddings, it does not yet outperform classical approaches. However, it provides valuable insights into the behavior of hybrid models and establishes a foundation for future research, particularly in the context of real quantum hardware and larger datasets.

5. Limitations

Despite the promising results obtained with the proposed Adaptive Multi-Embedding Quantum Ensemble (AME-QE), several limitations must be acknowledged.

5.1. Scalability of Quantum Kernels

A primary limitation of the proposed approach is the computational complexity associated with quantum kernel evaluation. The construction of kernel matrices scales quadratically with the number of samples, i.e., $O(N^2)$, making it impractical for large-scale datasets. In this study, experiments were restricted to 400 samples to ensure feasibility. This limitation is consistent with prior work highlighting scalability challenges in quantum kernel methods [3].

5.2. Dependence on Classical Simulation

All experiments were conducted using the PennyLane simulator rather than real quantum hardware. While simulators enable reproducibility and flexibility, they do not capture noise, decoherence, and hardware constraints present in real quantum devices. As noted in [15], current Noisy Intermediate-Scale Quantum (NISQ) devices remain limited in qubit count and stability, restricting practical deployment.

5.3. Limited Dataset Size

Due to the high computational cost, the dataset was significantly reduced from its original size. While this enables experimentation, it may limit the generalization capability of the model. Previous studies emphasize the importance of large-scale datasets for reliable IDS evaluation [4].

5.4. Performance Gap with Classical Models

The experimental results show that classical Random Forest consistently outperforms the proposed quantum fusion methods. This performance gap suggests that, for structured tabular data such as NSL-KDD, classical models remain highly competitive. Similar observations have been reported in recent comparative studies between classical and quantum machine learning approaches.

https://doi.org/10.3390/engproc2026154006

5.5. Sensitivity to Embedding Strategy

The results demonstrate that the performance of the proposed framework is highly sensitive to the choice of embedding and fusion parameter α. In particular, Angle Embedding contributes more effectively than IQP Embedding, indicating that more complex quantum feature maps do not necessarily yield better performance. This aligns with findings in [14], where the effectiveness of feature maps depends strongly on data characteristics.

5.6. Lack of Statistical Variability Analysis

Due to the high computational cost associated with quantum kernel evaluation and simulation-based execution, the current study reports average performance metrics obtained from stratified 5-fold cross-validation without additional repeated training cycles for statistical variability estimation. Consequently, standard deviations, confidence intervals, and formal statistical significance testing were not included in the present work.

Future research will focus on conducting repeated experimental runs and statistical validation using larger computational resources and optimized quantum implementations to provide a more comprehensive assessment of model stability and reproducibility.

5.7. Future Directions

Future work will focus on addressing these limitations by:

- Exploring kernel approximation techniques to reduce computational complexity;
- Evaluating the framework on real quantum hardware;
- Item Extending experiments to larger datasets and multiclass scenarios;
- Investigating adaptive or learnable fusion mechanisms.

6. Conclusions

In this paper, we introduced an Adaptive Multi-Embedding Quantum Ensemble (AME-QE) framework for intrusion detection, combining multiple quantum feature maps through concatenation and weighted fusion strategies. The proposed approach was evaluated on the NSL-KDD dataset under a binary classification setting using a stratified 5-fold cross-validation protocol.

The experimental results demonstrate that while classical Random Forest achieves superior overall performance, the proposed quantum fusion models provide competitive results and offer valuable insights into hybrid quantum–classical learning. In particular, the analysis reveals that Angle Embedding contributes more effectively to classification performance than IQP Embedding, and that appropriately tuning the fusion parameter α can significantly improve the balance between precision and recall.

Beyond raw performance, this work highlights an important observation: increasing the complexity of quantum feature maps does not necessarily translate into better results for structured tabular data. Instead, simpler and more stable embeddings may offer better generalization in practical cybersecurity scenarios. Additionally, the confusion matrix and error rate analysis show that quantum fusion models tend to maintain low false positive rates, which is a desirable property in intrusion detection systems where false alarms can be costly.

However, the study also underscores a critical limitation of current QML approaches, namely the high computational cost associated with quantum kernel evaluation. The need to restrict experiments to a small subset of the dataset illustrates the scalability challenges that must be addressed before real-world deployment becomes feasible.

Overall, this work contributes to the growing body of research on quantum machine learning for cybersecurity by proposing a novel multi-embedding fusion strategy and providing a detailed empirical analysis of its behavior. While classical methods remain

dominant at this stage, the insights gained from this study pave the way for future research on scalable, adaptive, and hardware-efficient quantum learning models.

Future work will focus on extending the framework to multiclass intrusion detection, incorporating additional embedding strategies, and evaluating performance on larger datasets and real quantum hardware. Furthermore, integrating statistical validation and adaptive fusion mechanisms represents a promising direction for improving robustness and generalization.

Author Contributions: Conceptualization, R.A.K. and H.B.; methodology, R.A.K., H.B. and S.S.; software, R.A.K. and H.B.; validation, S.S., H.B. and A.E.; formal analysis, H.B. and S.S.; investigation, R.A.K. and H.B. and S.S.; resources, R.A.K.; data curation, R.A.K.; writing—original draft preparation, R.A.K. and S.S.; writing—review and editing, H.B. and S.S.; visualization, R.A.K.; supervision, H.B. and A.E.; project administration, H.B.; funding acquisition, S.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by project 2026-FEEA-01, "Systematic Study of Methodological and Architectural Approaches to Modeling Digital Transformation Based on Artificial Intelligence," funded by the Research Fund of University of Ruse"Angel Kanchev".

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available upon request from the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

AME-QE	Adaptive Multi-Embedding Quantum Ensemble
IDS	Intrusion Detection System
QML	Quantum Machine Learning
RF	Random Forest
SVM	Support Vector Machine
QSVM	Quantum Support Vector Machine
QNN	Quantum Neural Network
VQC	Variational Quantum Circuit
IQP	Instantaneous Quantum Polynomial
PCA	Principal Component Analysis
FPR	False Positive Rate
FNR	False Negative Rate
NSL-KDD	Network Security Laboratory Knowledge Discovery Dataset
NISQ	Noisy Intermediate-Scale Quantum
CV	Cross-Validation
Acc	Accuracy
F1	F1-Score

References

1. Nielsen, M.A.; Chuang, I.L. *Quantum Computation and Quantum Information*; Cambridge University Press: Cambridge, UK, 2010.
2. Havlíček, V.; Córcoles, A.D.; Temme, K.; Harrow, A.W.; Kandala, A.; Chow, J.M.; Gambetta, J.M. Supervised learning with quantum-enhanced feature spaces. *Nature* **2019**, *567*, 209–212. [CrossRef]

3. Kerkatou, R.A.; Belhadef, H.; Eutamene, A.; Stefanova, S.P. Hybrid Quantum Machine Learning for Intrusion Detection: A Comparative Study of QNN and QSVM Models. *IEEE Access* **2026**, *14*, 77540–77556. [CrossRef]
4. Liao, H.-J.; Lin, C.-H.R.; Lin, Y.-C.; Tung, K.-Y. Intrusion detection system: A comprehensive review. *J. Netw. Comput. Appl.* **2013**, *36*, 16–24. [CrossRef]
5. Kim, G.; Lee, S.; Kim, S. LSTM recurrent neural network classifier for intrusion detection. In Proceedings of the International Conference on Platform Technology and Service (PlatCon), Jeju, Republic of Korea, 15–17 February 2016.
6. Zeng, T.; Wen, Q.; Huang, Y.; Hu, Y.; Wang, C. Adaptive multi-embedding fusion via denoising bidirectional Mamba for sequential recommendation. *J. King Saud Univ. Comput. Inf. Sci.* **2026**, *38*, 231. [CrossRef]
7. Chen, X.; Xu, S.; Hu, S.; Ma, X. MambaFN: An adaptive fusion network based on an improved Mamba to fuse multi-modal images. *Displays* **2026**, *91*, 103251. [CrossRef]
8. Matra, A. MaTra4LS: Mamba-transformer dual-architecture for long and short-term interests fusion in sequential recommendation. *Inf. Sci.* **2026**, *746*, 123445. [CrossRef]
9. Liu, H.; Li, X.; Huang, Z.; Ye, T.; Liu, Y.; Tan, H. CAWM-Mamba: A unified model for infrared-visible image fusion and compound adverse weather restoration. *Expert Syst. Appl.* **2026**, *314*, 131665. [CrossRef]
10. Xie, X.; Cui, Y.; Tan, T.; Zheng, X.; Yu, Z. FusionMamba: Dynamic feature enhancement for multimodal image fusion with Mamba. *arXiv* **2024**, arXiv:2404.09498.
11. Schuld, M. Supervised quantum machine learning models are kernel methods. *Quantum* **2021**, *5*, 512. [CrossRef]
12. Tavallaee, M.; Bagheri, E.; Lu, W.; Ghorbani, A.A. A detailed analysis of the KDD CUP 99 dataset. In Proceedings of the IEEE Symposium on Computational Intelligence for Security and Defense Applications, Ottawa, ON, Canada, 8–10 July 2009.
13. Bergholm, V.; Izaac, J.; Schuld, M.; Gogolin, C.; Ahmed, S.; Ajith, V.; Alam, M.S.; Alonso-Linaje, G.; AkashNarayanan, B.; Asadi, A.; et al. PennyLane: Automatic differentiation of hybrid quantum-classical computations. *arXiv* **2018**, arXiv:1811.04968.
14. Schuld, M.; Killoran, N. Quantum machine learning in feature Hilbert spaces. *Phys. Rev. Lett.* **2019**, *122*, 040504. [CrossRef]
15. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2018**, *2*, 79. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.