Secure Systems Design and Technology Development. Veliko Tarnovo, Bulgaria
Proceedings of the International Scientific Conference, 2026. Volume I, 191-197

Lightweight Comparative Evaluation of Quantum and Hybrid Quantum-Classical Machine Learning Models for Intrusion Detection

Genadiy Gospodinov
Computer Sciences
University of Library Studies and Information Technologies
Sofia, Bulgaria
g.gospodinov@unibit.bg
ORCID: 0009-0007-0792-2162

Tamara Ristovska
Computer Sciences
University of Library Studies and Information Technologies
Sofia, Bulgaria
t.ristovska@unibit.bg
ORCID: 0009-0005-9870-1327

Tito Titov
Computer Sciences
University of Library Studies and Information Technologies
Sofia, Bulgaria
t.titov@unibit.bg

Kiril Nikolov
Computer Sciences
University of Library Studies and Information Technologies
Sofia, Bulgaria
k.nikolov@unibit.bg

Andriyan Stoilov
Computer Sciences
University of Library Studies and Information Technologies
Sofia, Bulgaria
a.stoilov@unibit.bg

Kliment Toshkov
Computer Sciences
University of Library Studies and Information Technologies
Sofia, Bulgaria
k.toshkov@unibit.bg
ORCID: 0009-0000-7449-9455

***Abstract*— Quantum machine learning (QML) has recently emerged as a promising research direction at the intersection of artificial intelligence, quantum computing, and cybersecurity. Among its potential applications, intrusion detection remains particularly relevant due to the growing complexity of cyber threats and the increasing need for adaptive and efficient detection mechanisms. This paper explores the applicability of lightweight quantum and hybrid quantum-classical machine learning approaches for intrusion detection in resource-constrained experimental settings. The study is designed as a comparative analysis of selected models, including variational and kernel-based quantum techniques, alongside a conventional classical baseline, in order to evaluate their practical potential for cybersecurity tasks.**
The experimental framework relies on a reduced and preprocessed cybersecurity dataset and is implemented in a simulated quantum environment using accessible open-source tools. The comparison considers standard performance indicators such as classification effectiveness and computational efficiency, while also reflecting on broader dimensions of model suitability for security-oriented applications. In addition to the technical evaluation, the paper addresses several ethical considerations associated with the use of QML in cybersecurity, including limited transparency, risks related to biased or imbalanced training data, and the operational implications of incorrect classifications in intrusion detection scenarios. Rather than claiming immediate quantum advantage, the study aims to provide a realistic and methodologically grounded perspective on the current role of QML in cybersecurity research. The expected contribution is twofold - first, to outline a feasible experimental approach for evaluating QML methods in intrusion detection and second, to highlight the importance of combining performance-oriented assessment with ethical and interpretability-related considerations when examining emerging intelligent security solutions.

Keywords – Cybersecurity, hybrid quantum-classical models, intrusion detection, quantum machine learning

I. INTRODUCTION

The increasing complexity of contemporary cyber threats has intensified the need for adaptive, data-driven, and computationally efficient intrusion detection mechanisms. Intrusion Detection Systems (IDS) remain a central component of cybersecurity architectures because they support the identification of suspicious, anomalous or malicious activities within networked environments. Traditional rule-based IDS approaches are often limited by their dependence on predefined signatures and their reduced ability to detect previously unseen attacks. For this reason, machine learning has become a widely investigated direction in intrusion detection research, offering the

Online ISSN 3044-7771
https://doi.org/10.68302/std2026.vol1.207

possibility to learn patterns from traffic data and to distinguish between normal and malicious behavior under different operational conditions. [1] At the same time, the use of machine learning in network intrusion detection is not straightforward, since real network environments are dynamic, imbalanced, noisy, and difficult to represent through stable training data. [2]

A significant body of research has explored the application of classical machine learning methods such as support vector machines, decision trees, random forests, logistic regression, and neural networks for intrusion detection. These methods have shown promising results on benchmark datasets, especially when combined with appropriate preprocessing, feature selection, and class-balancing strategies. Public datasets such as NSL-KDD intrusion detection dataset, derived as an improved benchmark based on the Knowledge Discovery and Data Mining Cup 1999 (KDD'99) dataset, and the University of New South Wales Network Benchmark 2015 (UNSW-NB15) dataset have played an important role in enabling comparative IDS research, although they also have well-known limitations related to representativeness, redundancy, imbalance, and the gap between benchmark traffic and real operational environments. [3], [4] Therefore, IDS evaluation should not rely exclusively on high aggregate accuracy, but should also consider false positives, false negatives, feature reduction, computational cost, and the practical interpretability of the resulting model.

In parallel with the development of classical machine learning for cybersecurity, quantum computing has introduced new possibilities for computational modelling, optimization, and pattern recognition. Quantum machine learning (QML) has emerged as an interdisciplinary research field that investigates how quantum information processing can be used to enhance or reformulate machine learning tasks. [5] In particular, quantum kernel methods and variational quantum algorithms have attracted attention because they are compatible with the current stage of noisy intermediate-scale quantum technologies and can be implemented in simulated or cloud-based quantum environments. [6], [7] Quantum kernel approaches are based on the idea of embedding classical data into a quantum feature space, while variational quantum classifiers use parameterized quantum circuits trained through a classical optimization loop. [8], [9]

Despite this growing interest, the practical role of QML in cybersecurity remains exploratory. Current quantum hardware is still constrained by noise, limited qubit counts, circuit depth restrictions, and execution costs. As a result, many QML experiments are performed in simulated environments and rely on reduced datasets, a small number of features, and binary classification settings. These restrictions do not invalidate QML research, but they require careful methodological framing. In intrusion detection, where datasets may contain high-dimensional, imbalanced, and heterogeneous network traffic records, lightweight QML evaluation is particularly relevant because it allows researchers to test whether quantum and hybrid quantum-classical models can be meaningfully compared with conventional machine learning baselines under realistic resource constraints.

This paper addresses this research context by presenting a lightweight comparative evaluation of selected quantum and hybrid quantum-classical machine learning models for intrusion detection. The study is not designed to claim immediate quantum advantage over established classical methods. Instead, it aims to provide a realistic and reproducible experimental perspective on how QML models can be implemented, evaluated, and interpreted in a constrained cybersecurity scenario. The proposed experimental design compares a conventional classical baseline with quantum and hybrid approaches, such as quantum kernel-based classification and variational quantum classification. The comparison focuses not only on classification effectiveness, but also on computational efficiency and broader suitability for security-oriented applications.

The motivation for this approach is twofold. First, intrusion detection represents a meaningful cybersecurity use case for testing emerging machine learning paradigms because it requires the identification of complex patterns in network data. Second, QML methods are still at an early stage of practical adoption, which makes small-scale, reproducible, and critically interpreted experiments more valuable than overly ambitious claims about quantum superiority. In this context, a lightweight experimental setting provides a balanced way to examine both the potential and the limitations of QML for intrusion detection. It also allows the research to remain technically feasible using open-source tools such as Python, scikit-learn, and Qiskit within a controlled virtualized environment. [10], [11], [12]

Another important aspect of this study is the inclusion of ethical and interpretability-related considerations. In cybersecurity, incorrect classifications may have operational consequences – false negatives can allow attacks to remain undetected, while false positives may overload analysts and reduce trust in the system. These concerns are especially relevant for models whose internal decision logic is difficult to interpret. Although QML models may offer new representational possibilities, their practical adoption in security contexts requires attention to transparency, reliability, bias in training data, and the responsible use of automated decision-support systems. This perspective is consistent with broader AI risk management principles, which emphasize trustworthy, accountable, and context-aware deployment of AI systems. [13]

The main contribution of this paper is a methodologically grounded and lightweight comparative framework for evaluating quantum and hybrid quantum-classical machine learning models in an intrusion detection scenario. The authors' specific contribution consists in designing and implementing a reproducible experimental workflow that adapts selected QML models to a reduced and preprocessed NSL-KDD dataset, compares them with a classical SVM baseline under the same experimental

conditions, and evaluates the obtained results through classification metrics, execution time, error distribution, and repeated-run stability. From an applied perspective, the obtained results can be used as a reference for researchers and security practitioners who need to assess the practical suitability of QML-based IDS models before extending the analysis to larger datasets, multiclass attack detection, noise-aware simulation, or real quantum backends. By doing so, the study positions QML not as a replacement for classical intrusion detection methods, but as an emerging research direction whose practical value must be assessed through careful, transparent, and reproducible experimentation.

The remainder of the paper is structured as follows. The Materials and Methods section describes the dataset preparation, preprocessing procedure, experimental environment, selected models, and evaluation metrics. The Results and Discussion section presents the comparative results, including classification performance, computational efficiency, and interpretation of the observed model behavior. Finally, the Conclusion summarizes the main findings, outlines the limitations of the study, and identifies possible directions for future research, including larger datasets, multiclass intrusion detection, noise-aware quantum simulation, and evaluation on real quantum backends.

II. Materials and Methods

The methodology follows a simple but structured workflow – dataset selection, preprocessing and feature reduction, model implementation, evaluation, and visual comparison of the obtained results. A classical machine learning model is used as a baseline, while two QML-oriented approaches are included to represent kernel-based and variational quantum learning. This makes it possible to compare not only predictive performance, but also the practical behavior of the models in terms of execution time, stability, and suitability for intrusion detection scenarios. The experimental implementation relies on open-source tools, including Python, scikit-learn, Qiskit, and Qiskit Machine Learning, which supports the reproducibility of the study and keeps the technical setup accessible for further validation. Qiskit provides open-source tools for developing and executing quantum programs, while Qiskit Machine Learning includes building blocks such as quantum kernels and quantum neural networks for classification and regression tasks. [11]

A. Experimental Design and Dataset

The experiment was designed as a binary classification task for intrusion detection. Network traffic records were grouped into two main classes: normal traffic and attack traffic. This binary formulation was selected to keep the experiment computationally feasible while still preserving the central logic of intrusion detection, namely the distinction between legitimate and potentially malicious behaviour. Although multiclass intrusion detection can provide a more detailed view of attack categories, it would also increase the complexity of the experimental setting and require a larger feature space, which is less suitable for a lightweight QML evaluation.

For the purposes of the study, the NSL-KDD dataset was selected as the experimental data source. NSL-KDD is widely used in intrusion detection research and was introduced as an improved version of the original KDD'99 dataset, addressing some of its known redundancy-related issues. At the same time, it should not be treated as a perfect representation of real contemporary network traffic. The Canadian Institute for Cybersecurity explicitly notes that, although NSL-KDD remains useful as a benchmark for comparing intrusion detection methods, it still has limitations and does not fully reflect modern operational networks. This makes it appropriate for a controlled comparative experiment, but not for making broad deployment claims.

In this study, the dataset is used in a reduced form. The reduction is not only a practical decision, but also a methodological requirement. Quantum and hybrid quantum-classical models executed in simulated environments are sensitive to the number of input features and training samples. Therefore, a smaller and balanced subset of records is used to allow fair execution of all selected models under the same experimental conditions. The aim is not to maximize dataset scale, but to establish a reproducible comparison between a classical baseline, a quantum kernel-based model, and a variational quantum classifier.

B. Data Preprocessing and Feature Reduction

Before model training, the dataset undergoes several preprocessing steps. First, categorical attributes are transformed into numerical form, since the selected machine learning models require numerical input. The class labels are also converted into a binary representation, where normal records are assigned to one class and attack records to the other. After that, the data are scaled to avoid disproportionate influence from features with larger numeric ranges. This step is especially important for distance-based and kernel-based models, including support vector machines and quantum kernel classifiers.

A train-test split is then applied to separate the data used for model training from the data used for evaluation. The same split is preserved across all compared models in order to ensure that the results are directly comparable. Because the study focuses on lightweight QML experimentation, the dimensionality of the feature space is reduced before the models are trained. In the proposed setup, Principal Component Analysis (PCA) is used to reduce the input space to four numerical components. This number is small enough to be compatible with a simulated quantum circuit and, at the same time, sufficient to preserve a compact representation of the original traffic patterns.

The choice of four features is also linked to the quantum model design. In many QML implementations, the number of input features is closely related to the number of qubits or to the complexity of the feature map. A high-dimensional input space would make the experiment unnecessarily heavy and could obscure the comparative purpose of the study. For this reason, feature reduction is treated not as a loss of methodological depth, but as a

necessary step for aligning the cybersecurity dataset with the practical constraints of quantum simulation.

C. Compared Models and Experimental Environment

Three models are included in the comparative evaluation. The first model is a classical Support Vector Machine, used as a baseline. SVMs are suitable for this role because they are established supervised learning methods for classification and are commonly used in benchmark classification tasks. The scikit-learn documentation describes support vector machines as supervised methods applicable to classification, regression, and outlier detection, with particular usefulness in high-dimensional spaces. [14]

The second model is a Quantum Support Vector Classifier. QSVC represents the quantum kernel-based part of the experiment. Instead of relying only on a classical kernel function, this approach uses a quantum kernel to map the input data into a quantum feature space. In Qiskit Machine Learning, QSVC is implemented as a classifier that uses a quantum kernel while keeping a familiar fit-and-predict workflow similar to classical machine learning models. [15] This makes it suitable for a comparative study, because it can be positioned close to the classical SVM baseline while introducing a quantum feature mapping mechanism.

The third model is a Variational Quantum Classifier. VQC represents the hybrid quantum-classical part of the experiment. It uses a parameterized quantum circuit whose output is interpreted as a classification result, while the parameters are adjusted through a classical optimization loop. According to the Qiskit Machine Learning documentation, VQC is a variational algorithm in which measured bitstrings are interpreted as classifier outputs. [16] In this study, the VQC model is included not because it is expected to outperform the classical baseline, but because it allows the experiment to examine how a variational QML approach behaves in a small-scale intrusion detection setting.

The implementation is carried out in a Kali Linux virtual machine using Python. Classical preprocessing and baseline modelling are performed with scikit-learn, while the quantum and hybrid models are implemented using Qiskit and Qiskit Machine Learning. Quantum execution is performed in a simulated environment rather than on real quantum hardware. This decision keeps the experiment reproducible and avoids additional variability caused by hardware noise, queue times, or backend availability. Qiskit Aer is suitable for this type of setup because it provides high-performance quantum circuit simulators and supports execution with or without noise models. [17]

TABLE 1 SUMMARY OF THE EXPERIMENTAL SETUP

Component	Description
Task	Binary intrusion detection
Dataset	Reduced and preprocessed NSL-KDD dataset
Classes	Normal traffic and attack traffic
Feature reduction	PCA to four components
Classical baseline	Support Vector Machine
Quantum model	Quantum Support Vector Classifier
Hybrid model	Variational Quantum Classifier
Environment	Kali Linux virtual machine
Main tools	Python, scikit-learn, Qiskit, Qiskit Machine Learning
Execution mode	Local quantum simulation

D. Evaluation Metrics

The models are evaluated using standard classification metrics and simple computational efficiency indicators. The selected performance metrics are accuracy, precision, recall, and F1-score. Accuracy provides a general view of the proportion of correctly classified instances, but it is not sufficient on its own for intrusion detection tasks. In IDS scenarios, the cost of different errors is not equal. A false negative may allow malicious activity to remain undetected, while a false positive may generate unnecessary alerts and increase the workload of security analysts. For this reason, precision, recall, and F1-score are included to provide a more balanced interpretation of model behaviour.

In addition to classification performance, training time and prediction time are measured. This is important because a model that achieves acceptable accuracy but requires disproportionate computational resources may be less suitable for practical intrusion detection. Runtime indicators are especially relevant in the context of QML, where simulation overhead can become significant even when the dataset is small. Therefore, computational efficiency is treated as part of the comparative evaluation rather than as a secondary implementation detail.

The results are planned to be presented through a summary table and several visual elements in the following section. These include a comparison of classification metrics, a comparison of execution times, and confusion matrices for the selected models. This combination of numerical and visual reporting allows the analysis to move beyond a single performance value and to discuss the models in terms of practical suitability, stability, and limitations. In line with the exploratory nature of the study, the results will be interpreted cautiously and will not be used to claim general quantum advantage. Instead, the experiment is intended to show how lightweight QML models can be tested, compared, and critically evaluated in the context of intrusion detection.

III. RESULTS AND DISCUSSION

The results reported in this section should be read as a compact experimental profile of the evaluated models rather than as a simple ranking by accuracy. The analysis therefore considers how each approach performs, how costly it is to execute, what types of errors it produces and how consistently it behaves across repeated runs.

A. Classification Performance of the Compared Models

The first part of the experimental evaluation focuses on the classification performance of the three selected models: the classical Support Vector Machine, the Quantum Support Vector Classifier, and the Variational Quantum Classifier. All models were trained and tested on the same reduced and pre-processed NSL-KDD subset, using a binary classification formulation with two classes: normal traffic and attack traffic. The dataset was balanced before training, and the feature space was reduced to four principal components in order to keep the comparison compatible with the lightweight quantum simulation setting.

The obtained results are presented in Table 2 and Figure 1. The classical SVM achieved the highest overall performance, with an accuracy, precision, recall, and F1-score of 0.9667. The identical values for these four metrics are the result of a balanced test set and a symmetric distribution of classification errors across the two classes. In other words, the model performed almost equally well in identifying both normal and attack instances.

TABLE 2 CLASSIFICATION PERFORMANCE OF THE EVALUATED MODELS

Model	Accuracy	Precision	Recall	F1-score
SVM	0.9667	0.9667	0.9667	0.9667
QSVC	0.9417	0.9649	0.9167	0.9402
VQC	0.8000	0.8103	0.7833	0.7966

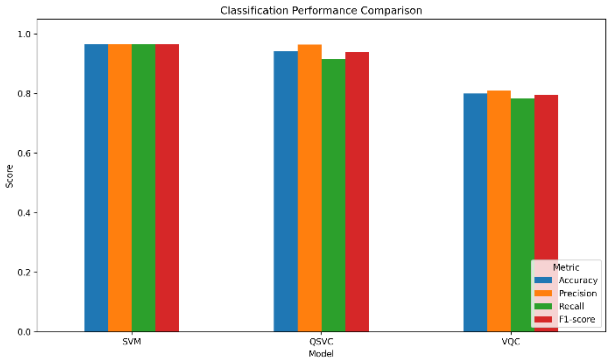

Figure 1 Classification performance comparison of SVM, QSVC and VQC

The QSVC model also produced competitive results, reaching an accuracy of 0.9417 and an F1-score of 0.9402. Its precision was slightly higher than its recall, which suggests that when the model predicted an attack, this prediction was usually reliable, but it missed a slightly larger number of attack instances than the SVM. This behaviour is important in intrusion detection, where recall has direct practical relevance because undetected attacks may remain active within a system.

The VQC model achieved lower performance across all classification metrics, with an accuracy of 0.8000 and an F1-score of 0.7966. Although this result is weaker than those of SVM and QSVC, it does not make the model irrelevant for the study. Rather, it shows that a hybrid variational quantum-classical approach can be implemented in a lightweight intrusion detection setting, but its performance remains more limited and more sensitive to the constraints of the experimental configuration.

Overall, the classification results indicate that the classical SVM remains the strongest model in this experimental setting. QSVC appears promising from a predictive perspective, since its results are relatively close to the SVM baseline. However, the VQC results suggest that variational quantum models may require more careful tuning, larger repeated-run validation, or alternative circuit configurations before they can reach comparable performance in this type of task.

B. Computational Efficiency

Classification performance alone is not sufficient for evaluating intrusion detection models, especially when the task is associated with operational environments where timely detection is important. For this reason, the second part of the evaluation examines the computational efficiency of the three models. The comparison includes training time, prediction time, and total execution time, measured during the same experimental run used for the classification analysis.

The results shown on figure 2 show a clear difference between the classical and QML-based approaches. The classical SVM required only 0.0021 seconds for training and 0.0005 seconds for prediction, resulting in a total execution time of approximately 0.0026 seconds. This confirms its practical advantage as a lightweight baseline model. In contrast, the QSVC model required 224.59 seconds for training and 165.11 seconds for prediction, reaching a total execution time of 389.70 seconds. The VQC model required 30.73 seconds for training and 0.3050 seconds for prediction, with a total execution time of 31.03 seconds.

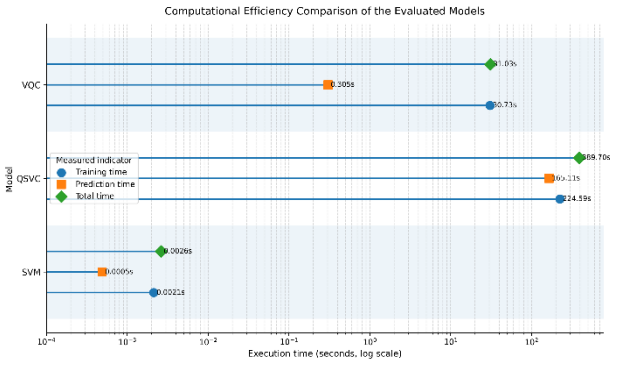

Figure 2 Computational efficiency comparison based on training, prediction and total execution time

These results reveal an important contrast between predictive quality and computational cost. QSVC achieved classification results relatively close to the SVM baseline, but this came with a substantial simulation overhead. Its total execution time was several orders of magnitude higher than that of the classical model. This difference is especially relevant for intrusion detection, where a model that performs well in terms of classification metrics may

still be difficult to justify if its execution time is too high for practical use.

The VQC model showed a different profile. It was clearly slower than the SVM, but considerably faster than the QSVC in this experiment. At the same time, its classification performance was weaker. This makes VQC an interesting but less competitive option in the present configuration. Its behaviour suggests that the hybrid quantum-classical approach may be computationally manageable in a small-scale setting, but further optimization would be necessary to improve its predictive performance.

Taken together, the computational results support the cautious interpretation adopted in this study. The quantum and hybrid models can be executed and compared in a lightweight simulated environment, but their practical value cannot be assessed only through accuracy or F1-score. In the present experiment, the SVM offers the best balance between performance and execution time, while QSVC demonstrates stronger predictive potential than VQC but at a much higher computational cost. This confirms that, under the current experimental constraints, the evaluated QML approaches are better understood as exploratory research models rather than immediately deployable intrusion detection solutions.

C. Error Distribution and Stability Analysis

To better understand the behaviour of the evaluated models, the analysis also considered the distribution of classification errors. This is particularly important in intrusion detection, where false negatives are more critical than ordinary classification mistakes, since they indicate attack instances classified as normal traffic. The confusion matrices for the three models are presented in Figure 3.

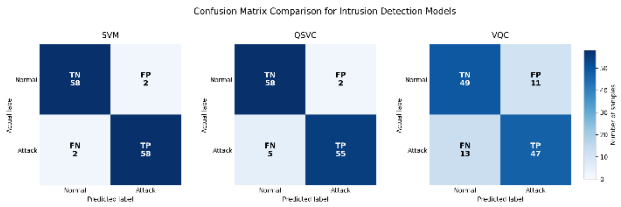

Figure 3 Confusion matrices for SVM, QSVC and VQC

The SVM model showed the most balanced error distribution, with 2 false positives and 2 false negatives. QSVC produced the same number of false positives, but a higher number of false negatives, misclassifying 5 attack samples as normal traffic. This explains its lower recall compared with SVM. The VQC model produced the weakest error profile, with 11 false positives and 13 false negatives, which confirms that its lower overall performance is not caused by a single isolated metric, but by a broader decline in classification reliability.

A repeated-run stability analysis was also performed using three different random seeds. The purpose of this step was to check whether the results remain consistent when the data split and sampling conditions change. The stability comparison based on F1-score is shown in Figure 4.

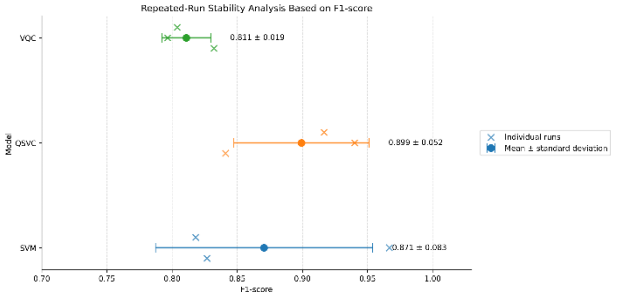

Figure 4 Repeated-run stability analysis based on F1-score

The repeated-run results showed that QSVC achieved the highest mean F1-score across the three executions, with a mean value of 0.8993 and a standard deviation of 0.0518. SVM followed with a mean F1-score of 0.8706, but with a higher standard deviation of 0.0833, indicating stronger variability across the repeated runs. VQC achieved the lowest mean F1-score, 0.8109, although its standard deviation was comparatively small at 0.0187. These results suggest that the single-run advantage of SVM should be interpreted together with the stability results, while QSVC appears more consistent in this reduced experimental setting despite its high computational cost.

D. Discussion of Practical Suitability

The results indicate that no model should be evaluated only through a single metric. The SVM model remains the most practical solution in the present setup because it combines high classification performance with extremely low execution time. Its main advantage is not only its accuracy, but the balance between predictive quality, speed, and relatively stable behaviour in the initial run.

QSVC is the most interesting quantum-oriented result. It achieved classification performance close to the classical baseline and showed the highest mean F1-score in the repeated-run analysis. However, this advantage is limited by the very high execution time observed during simulation. In practical intrusion detection, such computational overhead would be difficult to justify unless future implementations reduce the cost of quantum kernel evaluation.

The VQC model demonstrated that a hybrid quantum-classical classifier can be applied to the selected IDS task, but its results were less competitive. Its lower accuracy, higher number of false positives and false negatives, and weaker F1-score indicate that this configuration requires further tuning before it can be considered practically useful.

Overall, the findings support the main assumption of the study – lightweight QML models can be implemented and compared in an intrusion detection scenario, but the results do not demonstrate immediate quantum advantage. Instead, they show that QML methods are currently more suitable as exploratory research tools, while classical models such as SVM remain stronger candidates for practical deployment under constrained experimental conditions.

IV. Conclusion

This paper presented a lightweight comparative evaluation of classical, quantum, and hybrid quantum-classical machine learning models for intrusion detection. Using a reduced and pre-processed NSL-KDD dataset, the study compared SVM, QSVC, and VQC in terms of classification performance, computational efficiency, error distribution, and repeated-run stability. The applied value of the obtained results lies in providing a compact and reproducible evaluation procedure that helps identify the trade-offs between classification performance, computational cost, error behaviour, and stability before QML-based IDS approaches are considered for more realistic cybersecurity environments.

The results show that the classical SVM achieved the best balance between predictive performance and execution time. QSVC produced competitive classification results and demonstrated strong potential in the repeated-run analysis, but its high computational cost in the simulated environment limits its immediate practical applicability. VQC was successfully implemented as a hybrid quantum-classical model, but its performance remained weaker in the selected experimental configuration.

In conclusion, the findings do not indicate a clear quantum advantage in the examined intrusion detection scenario. Instead, they suggest that QML models should currently be understood as promising exploratory tools rather than direct replacements for established classical methods. Future work may extend this evaluation by using larger and more recent intrusion detection datasets, testing multiclass attack classification, exploring alternative quantum circuits, and comparing simulation-based results with executions on real quantum backends.

Acknowledgements

This article is the result of research funded by the National Science Fund of Bulgaria under project No. BG-175467353-2024-10-0013, contract No. КП-06-М87/5, dated December 6, 2024. The project is titled Artificial Intelligence and Cybersecurity: Navigating the Ethical Challenges.

References

[1] A. L. Buczak and E. Guven, “A Survey of Data Mining and Machine Learning Methods for Cyber Security Intrusion Detection,” *IEEE Communications Surveys and Tutorials*, vol. 18, no. 2, pp. 1153–1176, Apr. 2016, doi: 10.1109/COMST.2015.2494502;PAGE:STRING:ARTICLE/CHAPTER.

[2] R. Sommer and V. Paxson, “Outside the Closed World: On Using Machine Learning For Network Intrusion Detection”.

[3] “NSL-KDD.” Accessed: May 10, 2026. [Online]. Available: https://www.kaggle.com/datasets/hassan06/nslkdd

[4] “The UNSW-NB15 Dataset | UNSW Research.” Accessed: May 10, 2026. [Online]. Available: https://research.unsw.edu.au/projects/unsw-nb15-dataset

[5] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, “Quantum machine learning,” *Nature 2017 549:7671*, vol. 549, no. 7671, pp. 195–202, Sep. 2017, doi: 10.1038/nature23474.

[6] J. Preskill, “Quantum Computing in the NISQ era and beyond,” *Quantum*, vol. 2, p. 79, Aug. 2018, doi: 10.22331/q-2018-08-06-79.

[7] M. Cerezo *et al.*, “Variational quantum algorithms,” *Nature Reviews Physics 2021 3:9*, vol. 3, no. 9, pp. 625–644, Aug. 2021, doi: 10.1038/s42254-021-00348-9.

[8] M. Schuld and N. Killoran, “Quantum Machine Learning in Feature Hilbert Spaces,” *Phys. Rev. Lett.*, vol. 122, no. 4, p. 040504, Feb. 2019, doi: 10.1103/PhysRevLett.122.040504.

[9] V. Havlíček *et al.*, “Supervised learning with quantum-enhanced feature spaces,” *Nature 2019 567:7747*, vol. 567, no. 7747, pp. 209–212, Mar. 2019, doi: 10.1038/s41586-019-0980-2.

[10] F. Pedregosa FABIANPEDREGOSA *et al.*, “Scikit-learn: Machine Learning in Python,” *Journal of Machine Learning Research*, vol. 12, no. 85, pp. 2825–2830, 2011, Accessed: May 10, 2026. [Online]. Available: http://jmlr.org/papers/v12/pedregosa11a.html

[11] G. Aleksandrowicz *et al.*, “Qiskit: An Open-source Framework for Quantum Computing”, doi: 10.5281/ZENODO.2562111.

[12] “GitHub - qiskit-community/qiskit-machine-learning: An open-source library built on Qiskit for quantum machine learning tasks at scale on quantum hardware and classical simulators · GitHub.” Accessed: May 10, 2026. [Online]. Available: https://github.com/qiskit-community/qiskit-machine-learning

[13] N. Institute of Standards, “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” 2020, doi: 10.6028/NIST.AI.100-1.

[14] A. J. Smola and B. Schölkopf, “A tutorial on support vector regression,” *Stat. Comput.*, vol. 14, no. 3, pp. 199–222, Aug. 2004, doi: 10.1023/B:STCO.0000035301.49549.88.

[15] “QSVC - Qiskit Machine Learning 0.9.0.” Accessed: May 10, 2026. [Online]. Available: https://qiskit-community.github.io/qiskit-machine-learning/stubs/qiskit_machine_learning.algorithms.QSVC.html

[16] “VQC - Qiskit Machine Learning 0.9.0.” Accessed: May 10, 2026. [Online]. Available: https://qiskit-community.github.io/qiskit-machine-learning/stubs/qiskit_machine_learning.algorithms.VQC.html

[17] “Qiskit Aer 0.17.1.” Accessed: May 10, 2026. [Online]. Available: https://qiskit.github.io/qiskit-aer/