Towards Practical Quantum Kernels for Network Intrusion Detection

Mary L. Cotrupi[1], **Brian R. Callahan**[2]

[1]Dept. of Computer Science, Rensselaer Polytechnic Institute, 110 8th Street, Troy, NY 12180-3590 USA
[2]Dept. of Computer Science and Software Engineering, Monmouth University, 400 Cedar Avenue, West Long Branch, NJ 07764-1828 USA
cotrum@rpi.edu, bcallaha@monmouth.edu

Abstract

With cyber-attacks becoming increasingly sophisticated, modern network intrusion detection systems (NIDSs) are relying on machine learning (ML) methods for their flexibility in detecting subtle anomalous patterns in huge amounts of network data. However, classical ML methods such as support vector machines (SVMs) often rely on the conversion of low-dimensional data into a high-dimensional space, creating complex linear systems that are time-consuming to evaluate on large data inputs such as network flow logs. We propose addressing this limitation by employing a hybrid quantum-classical ML model to leverage quantum computing's (QC's) superiority in high-dimensional areas. We constructed a quantum kernel with an SVM model and evaluated it on four different network attacks from a modern intrusion detection dataset. Results revealed an average hardware accuracy rate of 85% with noticeably small deviations between runs, suggesting that quantum kernels may be a noise-resistant solution. We evaluated these results alongside classical and noiseless quantum simulator benchmarks.

Introduction

Quantum machine learning (QML) is an emerging field in the realm of quantum computing, subject already to many early-stage research efforts in domains such as healthcare and drug discovery, finance, and cybersecurity (Lamichhane and Rawat 2025; Corli et al. 2025; La Cour 2023). The most realistic near-term applications of quantum will likely be in hybrid usage with classical methods (Lamichhane and Rawat 2025). This preview paper focuses on the extension of a proposed hybrid QML model, the quantum kernel with classical SVM (Marcantonio et al. 2023), for detecting a subset of advanced network intrusion attacks and evaluating its performance on noisy quantum hardware.

Our hope is to improve the capabilities of cyber practitioners with quantum-enhanced tooling to detect a wide variety of threats and anomalies. As quantum machines exist today, it is worth exploring what extant machines are capable of and further developing frameworks as quantum hardware matures.

Materials and Methods

Data Preparation

We selected the BCCC-CSE-CIC-IDS2018 dataset (BCCC-CSE-CIC-IDS2018; Shafi, Lashkari, and Roudsari 2025), an updated version of the University of New Brunswick's well-known CSE-CIC-IDS2018 (Bönninghausen, Uetz, and Henze 2024; Liu et al. 2022; CSE-CIC-IDS2018), with 46 million records and 300+ features for ample sample selection and feature flexibility for 16 different cybersecurity attacks. For this preliminary paper, we have selected a subset of four attacks, shown in Figures 2 and 3.

Our data subset creation and cleaning process involved removing missing columns, rows, highly collinear data, and using the ML gradient boosting framework (GBF) XGBoost (Chen and Guestrin 2016) to select the top eight features. We then randomly sampled 30 points from each corresponding benign and attack dataset to create a 1:1 anomaly:benign ratio in order to conduct smaller tests for our quantum kernel hardware comparisons, with the 1:1 ratio eliminating accuracy metric bias.

Quantum Kernel Construction

We make heavy use of the open-source QuASK (Quantum Advantage Seeker with Kernels) module (Marcantonio et al. 2023) for our quantum kernel construction. Similar to a classical kernel, quantum kernels are defined as the inner products in a Hilbert space. To calculate these products, we must use quantum state measurement (Incudini, Martini, and Pierro 2024). The quantum kernel maps the classical data into the Hilbert space of a quantum system (encodes it), and the pair of encoded samples is tested via the overlap or swap test, simple procedures—whose circuits are shown in Figure 1—that allow us to estimate the inner products of quantum states. These tests have low circuit depth overhead, a crucial aspect for achieving high attack classification accuracies on noisy intermediate-scale quantum (NISQ) devices (Wang et al. 2021; Shaib et al. 2023).

We then employ the Scikit C-Support Vector Classification model (SVC — scikit-learn 1.7.1 documentation) to handle the classical evaluation of our quantum-generated matrices to ultimately create a hybrid quantum-classical pipeline. For now we provide just an overview of kernel construction.

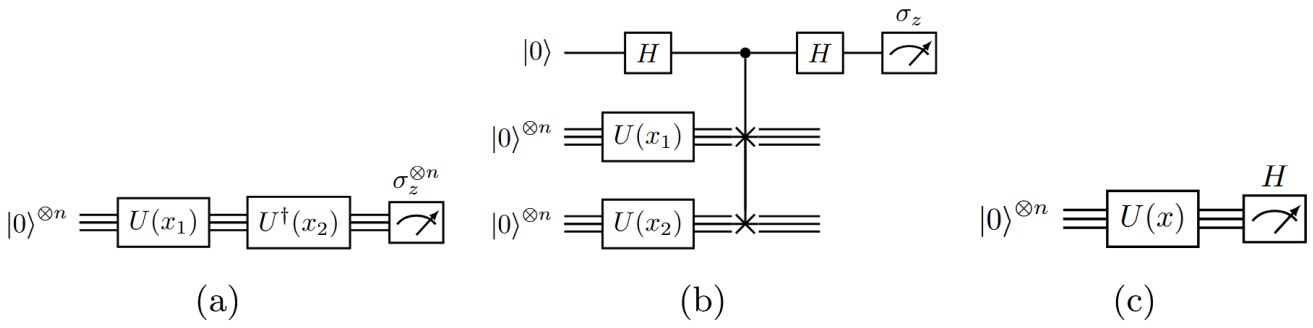

Figure 1: (1a)-(1b) Fidelity and SWAP test for quantum kernel estimation, where U is the feature map associated with the quantum kernel. (1c) Quantum circuit for the feature map associated with the projected kernel, the Hermitian observable H can be arbitrary (Marcantonio et al. 2023).

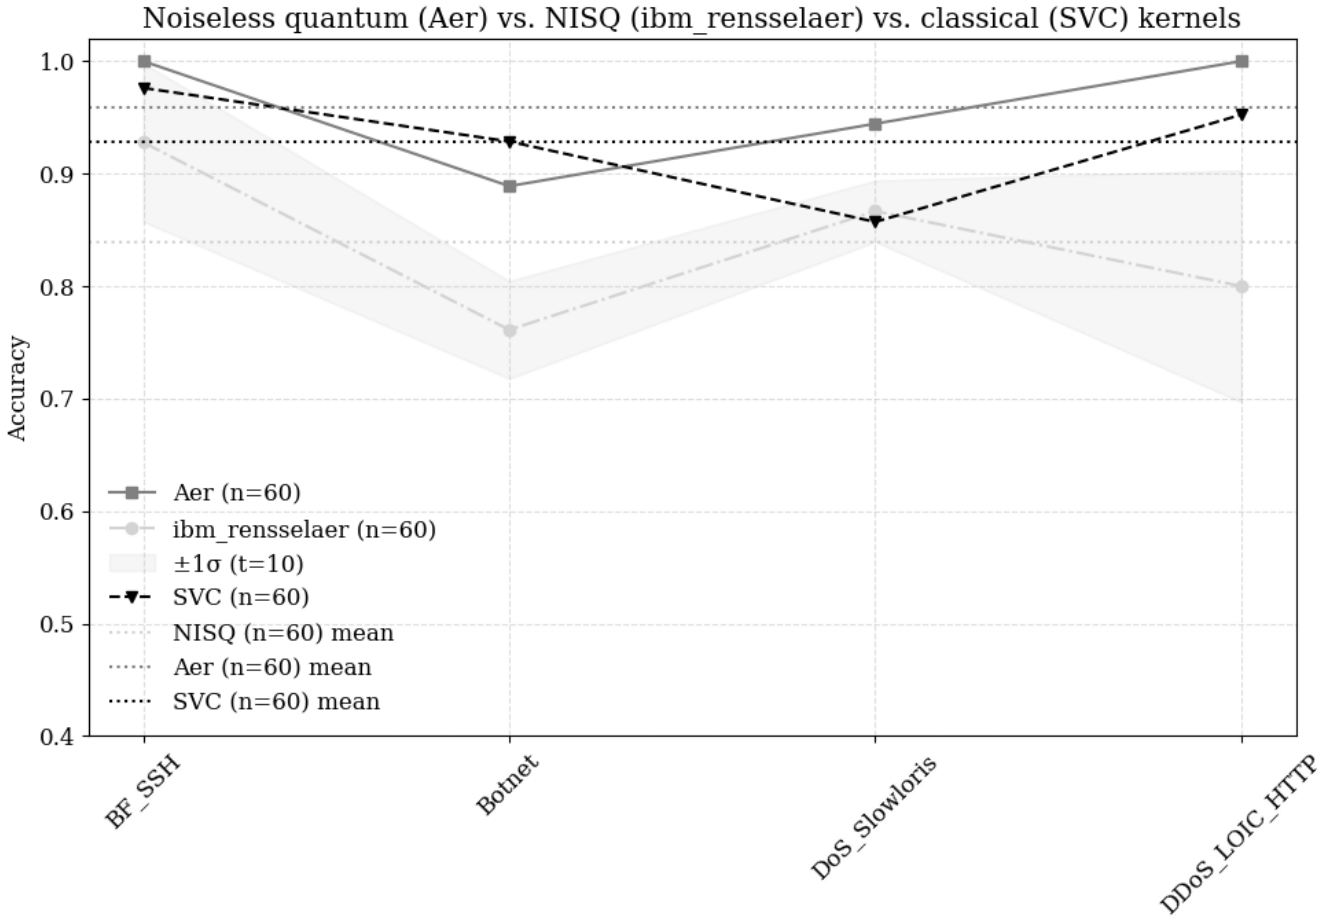

Figure 2: Quantum hardware vs. simulator vs. classical kernel predictions for the test dataset

Preliminary Results

The results of this study align with prior work (Payares and Martinez-Santos 2021; Kalinin and Krundyshev 2023; Wang et al. 2021) confirming that quantum machine learning can achieve high threat detection accuracy. Our preliminary findings (Figures 2 and 3) demonstrate that this accuracy persists on noisy quantum hardware, supporting the potential for quantum speedup over classical methods as quantum computers advance.

We tested our quantum kernel on the 127-qubit IBM

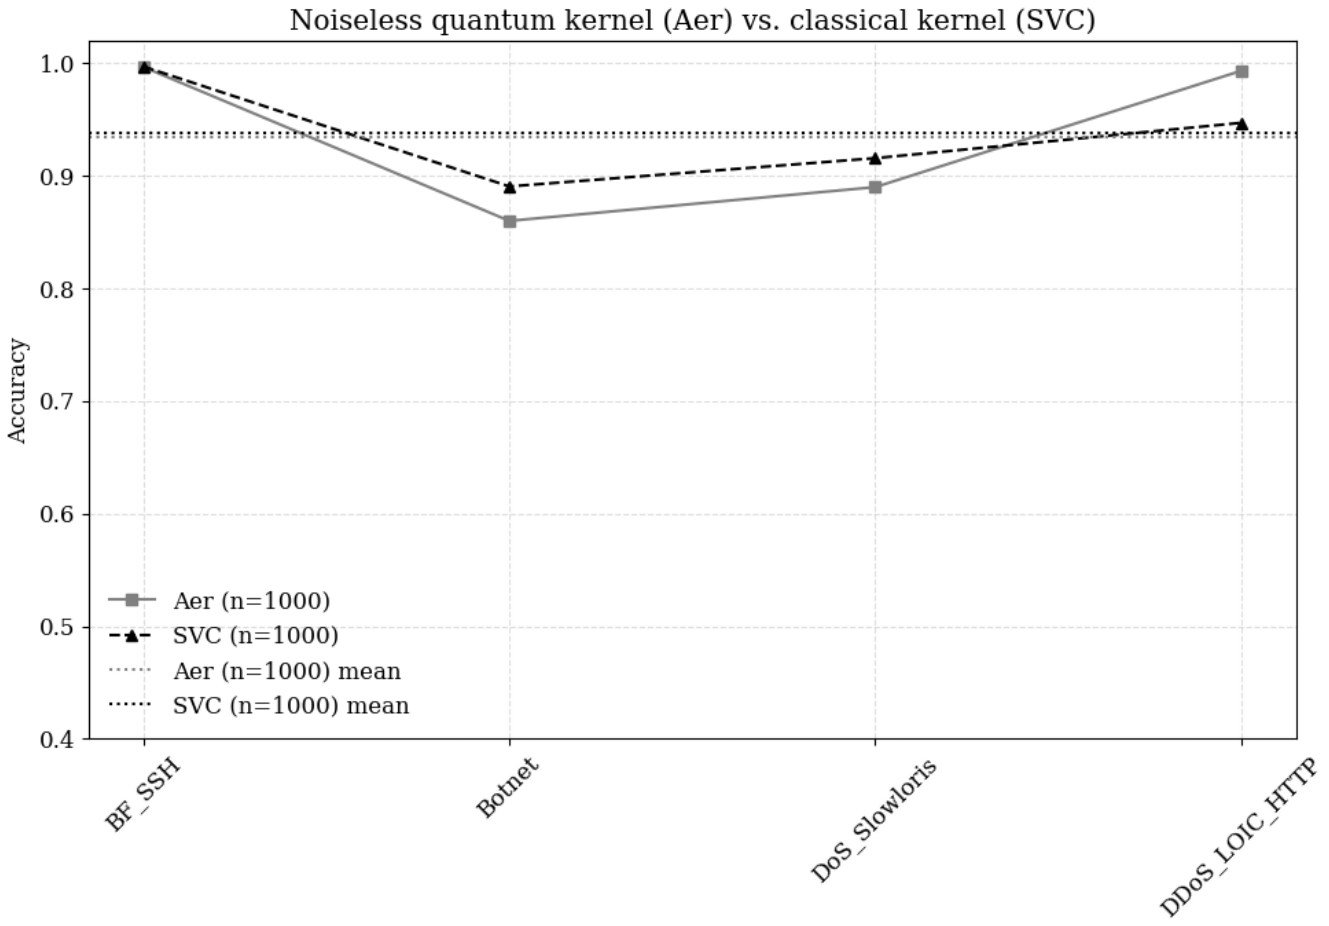

Figure 3: Quantum simulator vs. classical predictions for the test dataset

Quantum System One at Rensselaer Polytechnic Institute. Figure 2 compares these results against a classical benchmark kernel (Scikit's radial-based function [RBF] default) and a noiseless quantum simulator (Qiskit Aer) for n=60 points (30 anomalous and 30 benign distributed in an 8:2 training:testing ratio) for four attacks selected from BCCC-CSE-CIC-IDS2018: brute-force SSH attacks (BF_SSH), botnet attacks (Botnet), the Slowloris denial of service attack (DoS_Slowloris), and the Low Orbit Ion Cannon distributed denial of service attack (DDoS_LOIC_HTTP). We performed ten hardware iterations (t=10), with the shaded region depicting divergence due to noise, decoherence, and quantum error. Accuracy is the percentage of correct predictions for the test data.

Figure 3 follows a similar trend comparing only the quantum simulator and classical kernels for an increased sample size of n=1000 points, showing very close (93.5% Aer vs. 93.7% RBF) average accuracies between the quantum simulator and classical kernels and revealing the potential of quantum as a viable alternative to traditional methods in intrusion detection. We did not perform testing with n=1000 points on the Rensselaer quantum computer due to current limited qubit-power of the hardware and impractical runtime requirements.

Discussion & Future Work

The complexity of modern intrusions requires managing large datasets, and classical ML-based classification on such data often significantly degrades IDS performance during training and testing (Kalinin and Krundyshev 2023). Theoretical research has proven the potential for quantum speedup, from current polynomial-time SVM methods to logarithmic time of the vector lengths when using inner-product quantum evaluation (Ding, Bao, and Huang 2022).

Furthermore, many of the aforementioned studies lack hardware testing due to limited access to real quantum hardware, and this becomes a significant limitation in the studies (Gouveia and Correia 2020; Kalinin and Krundyshev 2023), forcing them to rely on noisy simulators that use simplified approximations that cannot be compared to experimentally obtained error rates from hardware (IonQ — noise model documentation).

However, our preliminary results in Figures 2 and 3 reveal that it is possible to achieve high accuracies on NISQ-era quantum hardware. Thus, the theoretical speedup studied is

entirely possible as quantum machines become more powerful.

Further Research

In addition to performing more hardware testing, we intend to increase the number and types of attacks tested in greater sample sizes to determine if our QML model is more suitable for a particular attack type, and to establish a firmer conclusion on whether the quantum kernel is a viable QML application in the current noisy era.

Conclusion

Our results demonstrate that it is feasible to use QML on currently existing quantum hardware to analyze network traffic data and classify that data as anomalous or benign. As quantum computing matures, we believe it is all the more likely that cybersecurity practitioners will face threats posed by quantum computing as well as use quantum computing to enhance their defenses (Hossain Faruk et al. 2022; Abellan and Pruneri 2018). As such, this preliminary study demonstrates the efficacy of one such avenue for enhancing cyber defenses using quantum computing and paves the way for further exploration at the intersections of quantum computing, artificial intelligence, and cybersecurity.

Acknowledgments

The authors gratefully acknowledge the generous support of an anonymous donor to the RPI School of Science whose contribution made this research possible.

References

Abellan, C.; and Pruneri, V. 2018. The future of cybersecurity is quantum. *IEEE Spectrum*, 55(7): 30–35.

BCCC-CSE-CIC-IDS2018. 2018. Available online at: https://www.yorku.ca/research/bccc/ucs-technical/cybersecurity-datasets-cds/.

Bönninghausen, P.; Uetz, R.; and Henze, M. 2024. Introducing a Comprehensive, Continuous, and Collaborative Survey of Intrusion Detection Datasets. In *Proceedings of the 17th Cyber Security Experimentation and Test Workshop*, CSET '24, 34–40. New York, NY, USA: Association for Computing Machinery. ISBN 979-8-4007-0957-9.

Chen, T.; and Guestrin, C. 2016. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, 785–794. New York, NY, USA: Association for Computing Machinery. ISBN 9781450342322.

Corli, S.; Moro, L.; Dragoni, D.; Dispenza, M.; and Prati, E. 2025. Quantum machine learning algorithms for anomaly detection: A review. (arXiv:2408.11047). ArXiv:2408.11047 [quant-ph].

CSE-CIC-IDS2018. 2018. Available online at: https://registry.opendata.aws/cse-cic-ids2018/.

Ding, C.; Bao, T.-Y.; and Huang, H.-L. 2022. Quantum-Inspired Support Vector Machine. *IEEE transactions on neural networks and learning systems*, 33(12): 7210–7222.

Gouveia, A.; and Correia, M. 2020. Towards Quantum-Enhanced Machine Learning for Network Intrusion Detection. In *2020 IEEE 19th International Symposium on Network Computing and Applications (NCA)*, 1–8. ISSN: 2643-7929.

Hossain Faruk, M. J.; Tahora, S.; Tasnim, M.; Shahriar, H.; and Sakib, N. 2022. A Review of Quantum Cybersecurity: Threats, Risks and Opportunities. In *2022 1st International Conference on AI in Cybersecurity (ICAIC)*, 1–8.

Incudini, M.; Martini, F.; and Pierro, A. D. 2024. Toward Useful Quantum Kernels. *Advanced Quantum Technologies*, 2024: 2300298.

IonQ — noise model documentation. N.d. Simulation with Noise Models. Available online at: https://docs.ionq.com/guides/simulation-with-noise-models.

Kalinin, M.; and Krundyshev, V. 2023. Security intrusion detection using quantum machine learning techniques. *Journal of Computer Virology and Hacking Techniques*, 19(1): 125–136.

La Cour, B. 2023. Advances in Quantum Computing. *Entropy*, 25(12): 1633.

Lamichhane, P.; and Rawat, D. B. 2025. Quantum Machine Learning: Recent Advances, Challenges, and Perspectives. *IEEE Access*, 13: 94057–94105.

Liu, L.; Engelen, G.; Lynar, T.; Essam, D.; and Joosen, W. 2022. Error Prevalence in NIDS datasets: A Case Study on CIC-IDS-2017 and CSE-CIC-IDS-2018. In *2022 IEEE Conference on Communications and Network Security (CNS)*, 254–262.

Marcantonio, F. D.; Incudini, M.; Tezza, D.; and Grossi, M. 2023. Quantum Advantage Seeker with Kernels (QuASK): a software framework to speed up the research in quantum machine learning. *Quantum Machine Intelligence*, 5(1): 20. ArXiv:2206.15284 [quant-ph].

Payares, E. D.; and Martinez-Santos, J. C. 2021. Quantum machine learning for intrusion detection of distributed denial of service attacks: a comparative overview. In *Quantum Computing, Communication, and Simulation*, volume 11699, 35–43. SPIE.

Shafi, M.; Lashkari, A. H.; and Roudsari, A. H. 2025. Toward generating a large scale intrusion detection dataset and intruders behavioral profiling using network and Transportation Layers Traffic Flow Analyzer (ntlflowlyzer). *Journal of Network and Systems Management*, 33(2): 44.

Shaib, A.; Naim, M. H.; Fouda, M. E.; Kanj, R.; and Kurdahi, F. 2023. Efficient noise mitigation technique for quantum computing. *Scientific Reports*, 13(1): 3912. Publisher: Nature Publishing Group.

SVC — scikit-learn 1.7.1 documentation. N.d. Available online at: https://scikit-learn.org/stable/modules/generated/sklearn.svm.SVC.html.

Wang, X.; Du, Y.; Luo, Y.; and Tao, D. 2021. Towards understanding the power of quantum kernels in the NISQ era. *Quantum*, 5: 531. ArXiv:2103.16774 [quant-ph].