Parameter-Efficient Quantum and Hybrid Autoencoders for One-Class Anomaly Detection

Murilo Salem[1], Daniel Pontes[1], João Carrett[1], Luísa Böhm[1]
Henrique dos Reis [1], Anderson Priebe Ferrugem[1]

[1]Centro de Desenvolvimento Tecnológico - Universidade Federal de Pelotas (UFPel)
96010-610 – Pelotas – RS – Brazil

{mcsalem, dhspbarretos, jcacarrett, lcbohm, hdreis, ferrugem}@inf.ufpel.br

Abstract. *This work investigates variational quantum autoencoders (QAE) for one-class anomaly detection under strict parametric constraints, shifting the evaluation focus from absolute performance to performance–capacity trade-offs. We implement a pure QAE baseline and a Hybrid QAE with classical compression/decoding around a quantum latent block, comparing both against classical baselines - AE, VAE, Isolation Forest, and One-Class SVM - on the NSL-KDD and ECG5000 benchmarks. The central contribution is a fairness-oriented evaluation protocol combining standard detection metrics (AUC-ROC, AUC-PR, and F1) with a performance-density measure (AUC-ROC per* $1{,}000$ *trainable parameters) and a fair-budget comparison against a compact classical AE. On NSL-KDD, the QAE baseline achieves* 0.9612 *AUC-ROC with only* 476 *parameters versus* 0.9679 *for a classical AE with* $5{,}580$ *parameters, while attaining* 2.0193 *AUC-ROC/1k parameters against* 0.1735 *for the classical AE - an order-of-magnitude improvement in parametric efficiency. Under a fair-budget setting, the QAE baseline also surpasses the compact classical AE on NSL-KDD (*0.9612 *vs.* 0.9405*). In contrast, ECG5000 favors classical methods, indicating domain dependence rather than universal quantum advantage. Overall, quantum and hybrid autoencoders are not universally superior, but deliver competitive anomaly detection with remarkably high parametric efficiency.*

1. Introduction

Anomaly detection is a fundamental machine learning task aimed at identifying patterns that significantly deviate from the expected behavior of a system [Chandola et al. 2009]. In many real-world applications, however, anomalous events are rare, heterogeneous, or unavailable during training, which makes supervised formulations impractical and motivates the adoption of the *one-class* setting [Pimentel et al. 2014, Chandola et al. 2009]. This paradigm is particularly relevant in critical domains such as cybersecurity and biomedicine, where the ability to detect rare deviations directly affects system reliability and decision quality [Chandola et al. 2009, Pimentel et al. 2014]. In such scenarios, learning the structure of normality from non-anomalous data becomes a principled and operationally viable strategy for identifying abnormal behavior at inference time [Pimentel et al. 2014].

Among the available solutions, reconstruction-based methods, especially Autoencoders (AE), have become prominent due to their ability to learn compact representations of normal data and assign higher anomaly scores to poorly reconstructed samples

[Sakurada and Yairi 2014, Zong et al. 2018]. Nevertheless, strong classical performance is often sustained by architectures with substantial parametric capacity, large memory footprint, and increased computational cost [Zong et al. 2018]. This creates an important asymmetry in model assessment: compact architectures are frequently judged only by absolute predictive metrics, while the architectural cost required to attain such performance remains underexplored. Consequently, model size and parameter efficiency should be treated not as secondary implementation details, but as central scientific variables in the evaluation of anomaly detection systems.

In this context, Quantum Machine Learning (QML) offers an alternative perspective through variational quantum models capable of representing complex correlations with compact parameterizations [Schuld and Killoran 2019, Cerezo et al. 2021]. This perspective is particularly relevant in cybersecurity, where recent work has already explored hybrid quantum machine learning for attack detection [Abreu et al. 2024]. However, the practical relevance of such models should not be framed in terms of universal superiority over classical baselines, especially under the constraints imposed by the NISQ era [Preskill 2018, Cerezo et al. 2021]. Instead, we hypothesize that Variational Quantum Autoencoders (QAE) and Hybrid QAE architectures can remain competitive in one-class anomaly detection while operating under much stricter parametric budgets [Romero et al. 2017, Preskill 2018]. Under this view, the key question is not whether quantum-enhanced models universally achieve the best absolute AUC or F1 values, but whether they can deliver a more favorable performance–capacity trade-off, measured through metrics such as AUC-ROC per 1,000 trainable parameters.

1.1. Main Contributions

This work makes four main contributions. First, we implement and evaluate two quantum-enhanced architectures for one-class anomaly detection, namely a pure Variational Quantum Autoencoder baseline and a Hybrid QAE that combines classical compression and decoding modules with a quantum latent block, both trained on a statevector simulator [Romero et al. 2017, Cerezo et al. 2021]. Second, we compare these models against strong classical baselines - a classical AE, a Variational Autoencoder, Isolation Forest, and One-Class SVM - on two heterogeneous benchmarks, NSL-KDD and ECG5000 [Kingma and Welling 2014, Liu et al. 2008, Schölkopf et al. 2001, Tavallaee et al. 2009, Chen et al. 2015]. Third, we introduce an evaluation framework centered on parametric efficiency, combining standard anomaly detection metrics with performance-density measures and fair-budget comparisons against compact classical autoencoders. Finally, we provide empirical evidence that the usefulness of quantum and hybrid models is domain-dependent: they remain highly competitive and markedly efficient in network intrusion detection, while classical methods still retain an advantage in more challenging biomedical time-series data.

2. Related Work / Background

This section positions the present study at the intersection of classical anomaly detection, variational quantum representation learning, and resource-constrained quantum computing. In one-class anomaly detection, classical methods have long been dominated by two major families: reconstruction-based approaches and boundary- or isolation-based models [Chandola et al. 2009, Pimentel et al. 2014]. Among the former, Autoencoders (AE) remain a standard solution due to their ability to learn compact latent representations of normal data and detect anomalies through elevated reconstruction error, while Variational Autoencoders (VAE) introduce probabilistic regularization that often improves robustness and generalization under uncertainty [Kingma and Welling 2014, Kingma and Welling 2019]. Complementarily, Isolation Forest and One-Class SVM constitute strong non-neural baselines based on distinct principles of anomaly identification, namely recursive isolation of rare instances and boundary estimation around the normal data manifold [Liu et al. 2008, Schölkopf et al. 2001]. Within the CSBC ecosystem, quantum formulations related to support vector machines have also been explored, reinforcing the relevance of SVM-based perspectives in quantum machine learning research [Pinheiro and Kowada 2024].

Quantum Autoencoders (QAE) extend the logic of latent compression to the quantum domain by employing parameterized quantum circuits (PQCs) to encode relevant information into a reduced set of qubits while preserving reconstructive fidelity [Romero et al. 2017]. Conceptually, QAEs offer a compact representation-learning mechanism in Hilbert spaces and have therefore attracted interest as potentially expressive models for anomaly detection. However, their practical utility is inherently conditioned by the limitations of the NISQ era, including restricted qubit counts, shallow feasible circuit depth, limited connectivity, and noise accumulation [Preskill 2018, Cerezo et al. 2021]. These issues are also reflected in recent discussions on hardware-aware NISQ architectures and compilation strategies [Cambiucci et al. 2024].

In response to these constraints, hybrid quantum-classical architectures have emerged as a pragmatic alternative [Benedetti et al. 2019, Cerezo et al. 2021]. Their central rationale is to delegate high-dimensional compression and final reconstruction to classical neural modules, while reserving the quantum block for latent-space transformation, where compact non-linear modeling may be most beneficial. Accordingly, hybrid models occupy an important intermediate position between fully classical baselines and purely quantum autoencoders, making them especially suitable for investigating whether a quantum latent block contributes meaningful value under severe capacity restrictions.

Despite the growing literature on quantum and hybrid anomaly detection, including recent work on semi-supervised quantum generative models for attack detection [Abreu et al. 2025], a central methodological gap remains: comparisons are often conducted under substantially unequal representational budgets. Classical baselines are frequently allowed to operate with much larger numbers of trainable parameters and more mature optimization pipelines, whereas quantum-enhanced models are judged primarily by raw predictive metrics. Such comparisons obscure whether compact quantum and hybrid architectures offer any true representational advantage relative to their cost. For this reason, a more appropriate perspective is to evaluate anomaly detection models under a *fair parametric budget*, in which predictive performance is interpreted jointly with architectural capacity.

3. Methodology

This section describes the formulation of the one-class anomaly detection problem, the datasets and preprocessing pipeline adopted in the experiments, the evaluated architectures, and the protocol used to ensure fair comparisons between quantum, hybrid, and classical models.

3.1. Problem Formulation

The anomaly detection task is formulated under the *one-class* setting. Let $x \in R^d$ denote an input sample and let f_θ be an autoencoding model trained exclusively on normal data. Given an input x, the model produces a reconstruction $\hat{x} = f_\theta(x)$, and the anomaly score is defined as the mean squared reconstruction error:

$$s(x) = \frac{1}{d} x - \hat{x}_2^2 . \tag{1}$$

The model parameters are obtained by minimizing the empirical reconstruction risk over the normal training distribution:

$$\theta^* = \arg\min_\theta \; E_{x \sim P_{\mathrm{norm}}} \left[\frac{1}{d} x - f_\theta(x)_2^2 \right] . \tag{2}$$

A sample is classified as anomalous when $s(x) \geq \tau$, where τ is a threshold calibrated on the validation set. This induces a decision function $\hat{y} : R^d \to \{0, 1\}$ given by:

$$\hat{y}(x) = I[s(x) \geq \tau] , \tag{3}$$

where $I[\cdot]$ denotes the indicator function.

In the main protocol, τ is selected by maximizing the validation F1-score, which can be written as:

$$\tau^* = \arg\max_{\tau \in \mathcal{T}} \frac{2 \cdot \mathrm{Precision}(\tau) \cdot \mathrm{Recall}(\tau)}{\mathrm{Precision}(\tau) + \mathrm{Recall}(\tau)} . \tag{4}$$

This preserves the strict one-class nature of training, since only normal samples are used to optimize the model, while labeled validation and test data are used solely for threshold calibration and final evaluation.

3.2. Datasets and Preprocessing

Experiments were conducted on two benchmarks: *NSL-KDD* and *ECG5000* [Tavallaee et al. 2009, Chen et al. 2015]. In NSL-KDD, the `normal` label is treated as the only normal class, and all attack categories are grouped as anomalies. Because the dataset contains both numerical and categorical attributes, preprocessing includes median imputation and standardization for numerical features, together with mode imputation and one-hot encoding for categorical variables (`protocol_type`, `service`, and `flag`). In ECG5000, class 1 is treated as normal and all remaining classes as anomalous; since all variables are numerical, preprocessing is restricted to median imputation and standard scaling.

The one-class split is constructed by allocating 60% of the normal samples to training, while the remaining normal instances are combined with all anomalous samples and divided equally between validation and test. Thus, training remains strictly one-class, whereas validation and test contain both normal and anomalous data. To prevent leakage, preprocessing is fitted only on the training partition and then applied unchanged to validation and test. In the main configuration, PCA with eight components is applied after preprocessing, and the parameter `train_subset_fraction` is used in the few-data experiments to subsample only the normal training set.

PCA is applied to reduce the input dimensionality to a representation compatible with the available qubit budget, removing redundant features and enabling direct encoding into the quantum circuit.

3.3. Models Evaluated

The study compares a *QAE baseline*, a *Hybrid QAE*, and four classical baselines. The QAE baseline uses a variational quantum block with four qubits and depth two. Each layer applies $R_Y(\theta_i x_i)$ gates for data encoding, followed by a circular CZ entanglement ring connecting qubit i to qubit $(i+1) \bmod 4$, and trainable $R_Y(\phi_i)$ rotation gates; this structure is repeated for each depth layer. Measurements in the Z-basis produce classical expectation values fed into a classical decoder, totaling approximately 476 trainable parameters. The Hybrid QAE introduces a small classical compressor before the same quantum block and a classical decoder after measurement, increasing flexibility while preserving the quantum bottleneck [Romero et al. 2017, Benedetti et al. 2019]; it uses approximately 860 parameters in the main setting and 444 in the fair-budget variant.

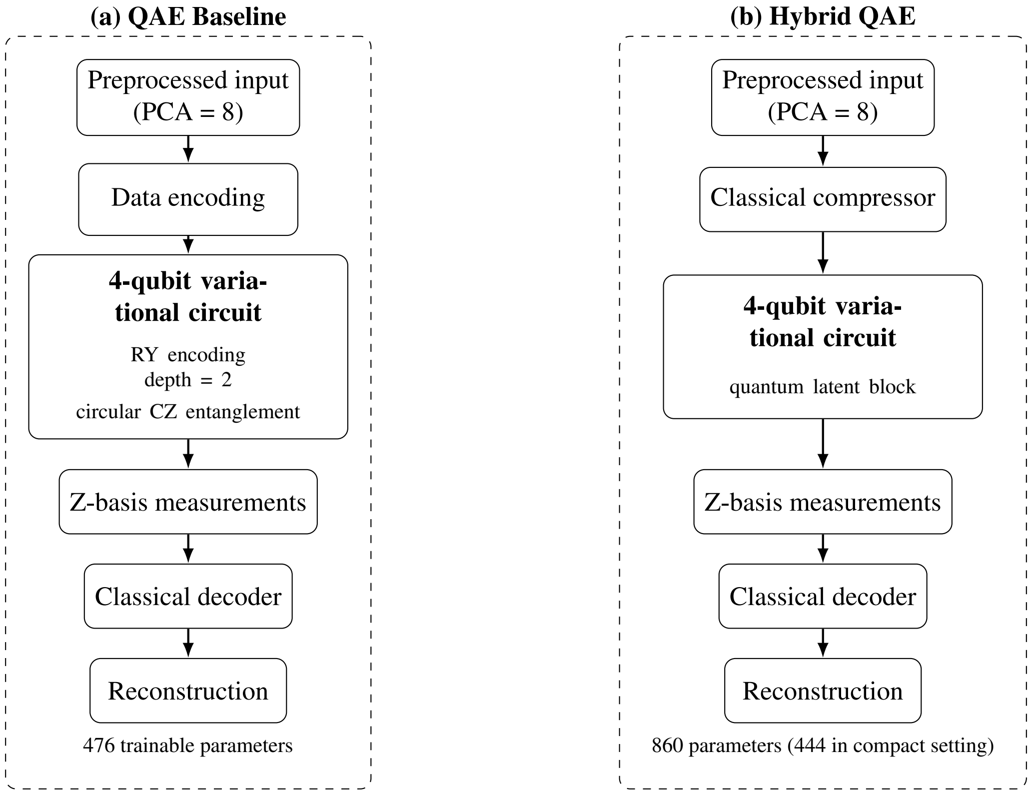

Figure 1. Overview of the proposed architectures. The QAE baseline directly maps the reduced input to a 4-qubit variational circuit, whereas the Hybrid QAE introduces a classical compressor before the same quantum latent block and a classical decoder after measurement.

Classical baselines include a standard Autoencoder (AE), a Variational Autoencoder (VAE), Isolation Forest, and One-Class SVM. The main AE configuration uses hidden layers $[64, 32]$ and latent dimension 4, totaling approximately $5{,}580$ trainable parameters, whereas the VAE contains approximately $5{,}712$ parameters [Kingma and Welling 2014]. For the fair-budget analysis, a compact AE with hidden layers $[8, 4]$ and latent dimension 2 is also evaluated, totaling roughly 242 parameters. Isolation Forest and One-Class SVM are included as strong non-neural references [Liu et al. 2008, Schölkopf et al. 2001].

3.4. Experimental Protocol

All models are evaluated under the same one-class data construction, preprocessing pipeline, and threshold-calibration strategy. Neural models are trained with Adam, learning rate 10^{-3}, batch size 64, and up to 100 epochs, with early stopping based on validation loss [Kingma and Ba 2014]. The primary evaluation metrics are AUC-ROC, AUC-PR, and F1-score. AUC-ROC and AUC-PR assess ranking quality independently of threshold choice, whereas F1-score evaluates the final binary decision after threshold calibration on the validation set [Fawcett 2006, Saito and Rehmsmeier 2015, Powers 2011].

To ensure fair comparisons, the number of trainable parameters is explicitly recorded for each neural model, and results are analyzed from two complementary perspectives: absolute predictive performance and performance under parametric constraints. In addition to the main comparison across all models, a *fair-budget* evaluation contrasts the QAE baseline, the compact Hybrid QAE, and a compact classical AE under a more balanced capacity regime. The protocol also includes a few-data experiment on NSL-KDD, in which the normal training set is reduced to 10%, 25%, 50%, and 100% of its original size while validation and test remain fixed, enabling a controlled assessment of data efficiency.

4. Results and Discussion

This section reports the empirical results obtained on the NSL-KDD and ECG5000 benchmarks and interprets them in light of the central hypothesis of this work. Taken together, the results indicate that quantum and hybrid autoencoders do not exhibit universal superiority over strong classical baselines, yet they remain highly competitive under strict parametric constraints and reveal a markedly favorable performance–capacity trade-off.

4.1. Main Comparison

Table 1 summarizes the primary comparison across all evaluated models. In terms of absolute performance, the strongest classical baselines remain dominant overall. On NSL-KDD, the classical AE achieved the best AUC-ROC (0.9679), followed closely by the Hybrid QAE (0.9651) and the QAE baseline (0.9612). Notably, both quantum-enhanced models outperformed the VAE, Isolation Forest, and One-Class SVM in AUC-ROC, while the Hybrid QAE also achieved the highest AUC-PR among all methods (0.9825), slightly exceeding the classical AE (0.9822).

A different pattern emerges on ECG5000. In this benchmark, classical methods retained clear superiority in absolute detection metrics. One-Class SVM achieved the best overall performance, with AUC-ROC of 0.9887, AUC-PR of 0.9889, and F1-score of 0.9764, followed by Isolation Forest and the classical AE. The QAE baseline and Hybrid QAE remained competitive but did not surpass the strongest classical baselines, reaching AUC-ROC values of 0.9446 and 0.9469, respectively. Overall, the contrast between datasets indicates that the effectiveness of the quantum latent block is domain-dependent, being substantially more favorable in NSL-KDD than in ECG5000.

Table 1. Main comparison across all evaluated models on NSL-KDD and ECG5000.

Dataset	Model	AUC-ROC	AUC-PR	F1	Parameters
NSL-KDD	Classical AE	0.9679	0.9822	0.9420	5580
	Hybrid QAE	0.9651	0.9825	0.9330	860
	QAE Baseline	0.9612	0.9801	0.9340	476
	Isolation Forest	0.9535	0.9754	0.9261	–
	VAE	0.9489	0.9739	0.9160	5712
	One-Class SVM	0.9240	0.9703	0.9157	–
ECG5000	One-Class SVM	0.9887	0.9889	0.9764	–
	Isolation Forest	0.9790	0.9772	0.9696	–
	Classical AE	0.9566	0.9502	0.9542	5580
	VAE	0.9544	0.9509	0.9531	5712
	Hybrid QAE	0.9469	0.9293	0.9484	860
	QAE Baseline	0.9446	0.9274	0.9488	476

4.2. Parametric Efficiency

A central contribution of this study lies in shifting the evaluation from absolute predictive performance to *parametric efficiency*. Table 2 reports the density of performance achieved by each trainable model, measured as AUC-ROC and AUC-PR per $1{,}000$ trainable parameters. Under this perspective, the ranking changes substantially. On NSL-KDD, the QAE baseline reaches 2.0193 AUC-ROC per $1{,}000$ parameters, whereas the classical AE achieves only 0.1735. Likewise, on ECG5000, the QAE baseline attains 1.9845, compared to 0.1714 for the classical AE. In both datasets, this corresponds to more than an order-of-magnitude improvement in performance density relative to large classical autoencoders.

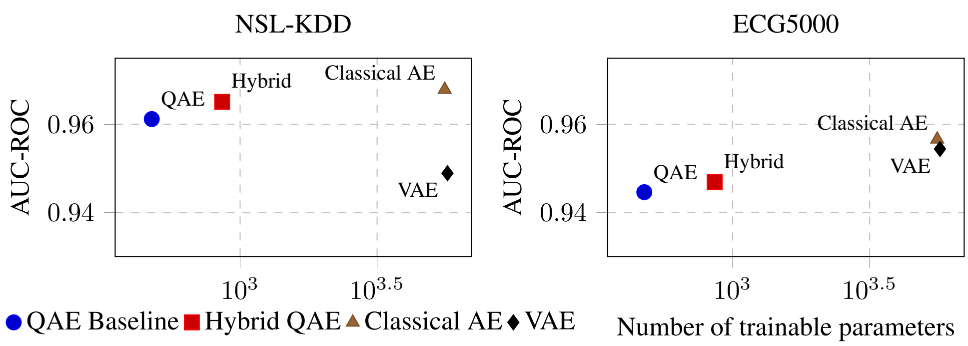

Figure 2. Performance–capacity trade-off of trainable models on NSL-KDD and ECG5000. Quantum-enhanced architectures achieve near-classical detection performance while operating with substantially fewer trainable parameters, occupying a favorable low-capacity region of the performance–capacity space.

The Hybrid QAE also exhibits a favorable trade-off, achieving 1.1222 AUC-ROC per $1{,}000$ parameters on NSL-KDD and 1.1010 on ECG5000. Although less dense than the pure QAE baseline, it improves absolute competitiveness while preserving a low-capacity regime. By contrast, the classical AE and VAE achieve stronger absolute scores among trainable neural models only at the expense of parameter counts above $5{,}500$, which causes their efficiency-per-parameter to drop sharply. Therefore, the relevance of the quantum-enhanced models lies not in universal dominance, but in their markedly superior representational density under strict capacity constraints.

Table 2. Parametric efficiency of trainable models, expressed as performance per $1{,}000$ trainable parameters. AUC-ROC/1k ratio can exceed 1 when the number of trainable parameters is smaller than $1{,}000$, since AUC-ROC is divided by the parameter count expressed in thousands.

Dataset	Model	AUC-ROC	Params	AUC-ROC/1k	AUC-PR/1k
NSL-KDD	QAE Baseline	0.9612	476	2.0193	2.0590
	Hybrid QAE	0.9651	860	1.1222	1.1424
	Classical AE	0.9679	5580	0.1735	0.1760
	VAE	0.9489	5712	0.1661	0.1705
ECG5000	QAE Baseline	0.9446	476	1.9845	1.9483
	Hybrid QAE	0.9469	860	1.1010	1.0806
	Classical AE	0.9566	5580	0.1714	0.1703
	VAE	0.9544	5712	0.1671	0.1665

4.3. Fair-Budget Comparison

To reduce the criticism that the main comparison favors classical models through substantially larger capacity, we conducted an additional fair-budget analysis using a compact classical AE as reference. This choice follows the broader need for carefully controlled benchmarking in quantum machine learning, where the conclusions can depend strongly on how quantum and classical models are matched in capacity and evaluation protocol [Bowles et al. 2024]. Table 3 reports the results in this restricted regime. On NSL-KDD, the QAE baseline achieved the best absolute AUC-ROC (0.9612), followed by the compact Hybrid QAE (0.9596), both outperforming the compact classical AE (0.9405).

On ECG5000, the compact classical AE remained slightly ahead in AUC-ROC (0.9472), while the Hybrid QAE achieved a nearly identical value (0.9469) and the highest F1-score in this restricted comparison (0.9514). The QAE baseline remained close, with AUC-ROC of 0.9446. Thus, even in the dataset least favorable to the quantum hypothesis, compact quantum and hybrid models remain competitive under tighter capacity constraints. Overall, the fair-budget comparison reinforces the main interpretation of this study: the strongest case for quantum-enhanced autoencoders is not unrestricted absolute superiority, but robust competitiveness in low-capacity regimes.

4.4. Few-Data Regime

In the few-data experiment on NSL-KDD [Tavallaee et al. 2009], the QAE baseline improved consistently as the fraction of normal training data increased, rising from AUC-

Table 3. Comparison under a fairer parametric budget.

Dataset	Model	AUC-ROC	AUC-PR	F1	Params	AUC-ROC/1k
NSL-KDD	QAE Baseline	0.9612	0.9801	0.9340	476	2.0193
	Hybrid QAE	0.9596	0.9795	0.9314	444	2.1612
	Classical AE Compact	0.9405	0.9720	0.9297	242	3.8863
ECG5000	Classical AE Compact	0.9472	0.9297	0.9497	242	3.9139
	Hybrid QAE	0.9469	0.9291	0.9514	444	2.1326
	QAE Baseline	0.9446	0.9274	0.9488	476	1.9845

ROC 0.9358 at 10% to 0.9442 at 100%, whereas the compact classical AE remained nearly stable and slightly decreased from 0.9361 to 0.9341. Although limited to a single benchmark, this suggests that the quantum latent block retains useful representational headroom in low-capacity settings and benefits more consistently from additional data [Benedetti et al. 2019]. Overall, the results reinforce a restrained conclusion: there is no universal quantum advantage in one-class anomaly detection, yet quantum and hybrid autoencoders provide compact and competitive solutions, particularly when performance is evaluated jointly with parameter budget and data regime [Bowles et al. 2024].

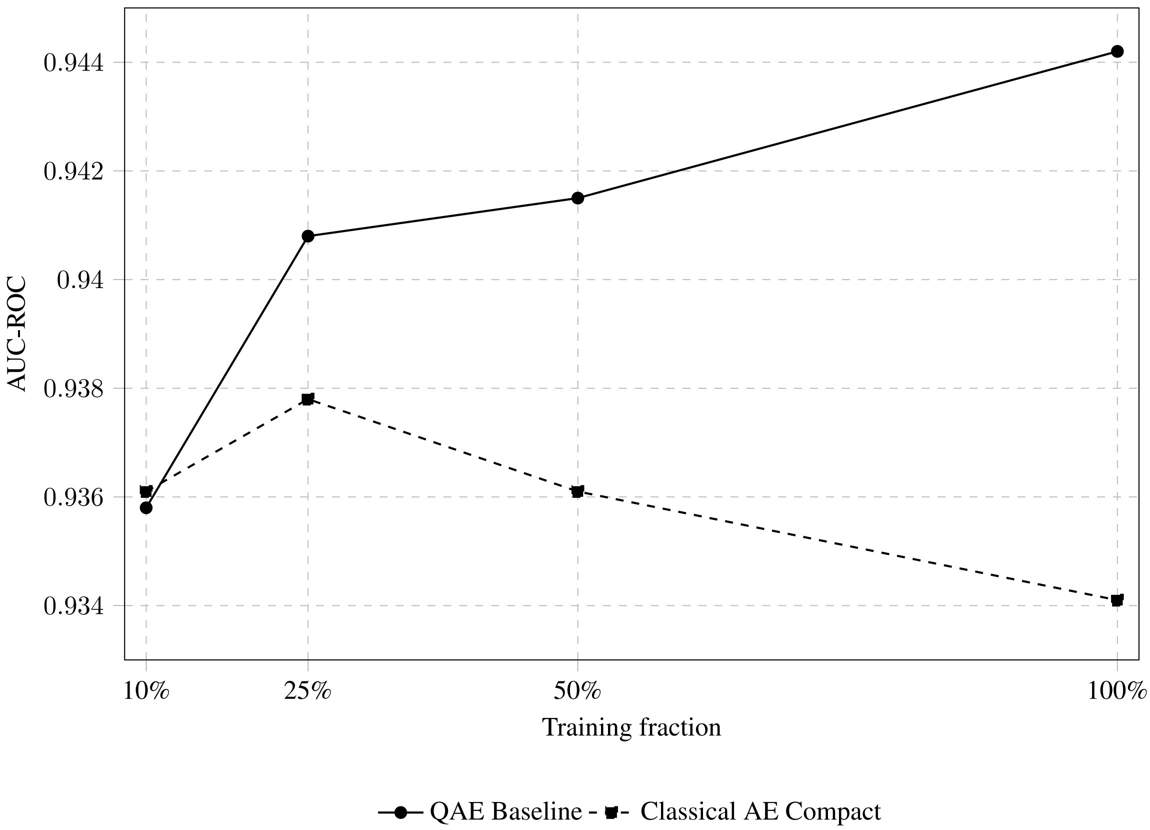

Figure 3. Data-efficiency behavior on NSL-KDD under limited normal training data. The QAE baseline improves more consistently with additional data than the compact classical AE.

The non-monotonic behavior of the compact classical AE, whose AUC-ROC slightly decreases from 10% to 100% of training data, may reflect capacity saturation: with only 242 parameters and latent dimension 2, the model likely reaches its representational limit early, and additional normal samples broaden the learned reconstruction boundary without improving anomaly discrimination.

5. Conclusion

This work investigated variational quantum and hybrid autoencoders for one-class anomaly detection under strict parametric constraints, framing the evaluation in terms of performance–capacity trade-offs rather than universal quantum superiority. The results show that classical methods remain dominant in absolute performance on ECG5000, where One-Class SVM achieved AUC-ROC of 0.9887, while on NSL-KDD the proposed models remain highly competitive, with the QAE baseline reaching 0.9612 AUC-ROC using only 476 parameters and the Hybrid QAE achieving 0.9651 with 860 parameters, compared to 0.9679 for a classical AE with $5{,}580$ parameters. Under a fair-budget setting, the QAE baseline surpasses the compact classical AE on NSL-KDD (0.9612 vs. 0.9405 AUC-ROC), while remaining competitive on ECG5000. Most importantly, both quantum-enhanced architectures exhibit substantially higher parametric efficiency, with the QAE baseline achieving up to 2.0193 AUC-ROC per $1{,}000$ parameters versus 0.1735 for the classical AE, indicating an order-of-magnitude improvement in performance density. In the few-data regime, the QAE baseline also shows more consistent gains as training data increases, suggesting preserved representational capacity under low-parameter settings. Overall, there is no universal quantum advantage, yet quantum and hybrid autoencoders offer a compelling alternative in low-capacity regimes. Future work should include multi-seed statistical analysis, noise-aware experiments on realistic hardware, and broader dataset evaluations to better delineate the practical scope of quantum-enhanced anomaly detection.

References

Abreu, D., Moura, D., Rothenberg, C., and Abelém, A. (2025). Rede generativa adversarial quântica semi-supervisionada (sqgan) para detecção de ataques. In *Anais do XLIII Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos*, pages 252–265. Sociedade Brasileira de Computação.

Abreu, D., Rothenberg, C. R. E., and Abelém, A. (2024). qids: Sistema de detecção de ataques baseado em aprendizado de máquina quântico híbrido. In *Anais do XLII Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos*, pages 295–308. Sociedade Brasileira de Computação.

Benedetti, M., Lloyd, E., Sack, S., and Fiorentini, M. (2019). Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4):043001.

Bowles, J., Ahmed, S., and Schuld, M. (2024). Better than classical? the subtle art of benchmarking quantum machine learning models. *arXiv preprint arXiv:2403.07059*.

Cambiucci, W., Silveira, R. M., and Ruggiero, W. V. (2024). Arquitetura para sistema de computação quântica multicore nisq com compilação sensível ao hardware. In *Anais do I Workshop de Redes Quânticas*, pages 7–12. Sociedade Brasileira de Computação.

Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., Cincio, L., and Coles, P. J. (2021). Variational quantum algorithms. *Nature Reviews Physics*, 3(9):625–644.

Chandola, V., Banerjee, A., and Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3):15:1–15:58.

Chen, Y., Keogh, E., Hu, B., Begum, N., Bagnall, A., Mueen, A., and Batista, G. (2015). The UCR time series classification archive. University of California, Riverside. Available at: `http://www.cs.ucr.edu/~eamonn/time_series_data/`.

Fawcett, T. (2006). An introduction to roc analysis. *Pattern Recognition Letters*, 27(8):861–874.

Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kingma, D. P. and Welling, M. (2014). Auto-encoding variational Bayes. In *International Conference on Learning Representations (ICLR)*.

Kingma, D. P. and Welling, M. (2019). An introduction to variational autoencoders. *Foundations and Trends in Machine Learning*, 12(4):307–392.

Liu, F. T., Ting, K. M., and Zhou, Z.-H. (2008). Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining*, pages 413–422. IEEE.

Pimentel, M. A. F., Clifton, D. A., Clifton, L., and Tarassenko, L. (2014). A review of novelty detection. *Signal Processing*, 99:215–249.

Pinheiro, G. and Kowada, L. A. B. (2024). Automation of the quantum algorithm hhl for implementing two-dimensional svms. In *Anais do I Workshop de Redes Quânticas*, pages 13–18. Sociedade Brasileira de Computação.

Powers, D. M. W. (2011). Evaluation: From precision, recall and f-measure to roc, informedness, markedness and correlation. *Journal of Machine Learning Technologies*, 2(1):37–63.

Preskill, J. (2018). Quantum computing in the NISQ era and beyond. *Quantum*, 2:79.

Romero, J., Olson, J. P., and Aspuru-Guzik, A. (2017). Quantum autoencoders for efficient compression of quantum data. *Quantum Science and Technology*, 2(4):045001.

Saito, T. and Rehmsmeier, M. (2015). The precision-recall plot is more informative than the roc plot when evaluating binary classifiers on imbalanced datasets. *PLOS ONE*, 10(3):e0118432.

Sakurada, M. and Yairi, T. (2014). Anomaly detection using autoencoders with nonlinear dimensionality reduction. In *Proceedings of the 2nd Workshop on Machine Learning for Sensory Data Analysis*, MLSDA '14, pages 4–11, New York, NY, USA. Association for Computing Machinery.

Schölkopf, B., Platt, J. C., Shawe-Taylor, J., Smola, A. J., and Williamson, R. C. (2001). Estimating the support of a high-dimensional distribution. *Neural computation*, 13(7):1443–1471.

Schuld, M. and Killoran, N. (2019). Quantum machine learning in feature hilbert spaces. *Physical Review Letters*, 122(4):040504.

Tavallaee, M., Bagheri, E., Lu, W., and Ghorbani, A. A. (2009). A detailed analysis of the NSL-KDD data set. In *2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications*, pages 1–6. IEEE.

Zong, B., Song, Q., Min, M. R., Cheng, W., Lumezanu, C., Zheng, D., and Chen, H. (2018). Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In *International Conference on Learning Representations (ICLR)*.