Quorum: Zero-Training Unsupervised Anomaly Detection using Quantum Autoencoders

Jason Zev Ludmir
Rice University
Houston, USA

Sophia Rebello
Rice University
Houston, USA

Jacob Ruiz
Stanford University
Stanford, USA

Tirthak Patel
Rice University
Houston, USA

***Abstract*—Detecting mission-critical anomalous events and data is a crucial challenge across various industries, including finance, healthcare, and energy. Quantum computing has recently emerged as a powerful tool for tackling several machine learning tasks, but training quantum machine learning models remains challenging, particularly due to the difficulty of gradient calculation. The challenge is even greater for anomaly detection, where unsupervised learning methods are essential to ensure practical applicability. To address these issues, we propose Quorum, the first quantum anomaly detection framework designed for unsupervised learning that operates without requiring any training.**

I. Introduction

Anomaly detection plays an essential role in various industries, from identifying fraudulent transactions in finance to detecting irregularities in power grids [26], [8], [25], [36] As datasets grow in complexity, traditional machine learning (ML) methods struggle with scalability and accuracy. Quantum computing offers a promising new approach, with its potential to accelerate computations and detect subtle patterns and correlations in data. Leveraging quantum algorithms for ML tasks, particularly for anomaly detection, could transform how we tackle these challenges [28], [27], [33].

The Challenge. Applying quantum computing to anomaly detection presents significant hurdles. Quantum machine learning (QML) models typically rely on parameterized circuits that require training, which is challenging due to the complexity of computing gradients in quantum systems [16], [31]. Quantum systems require gradient calculations from first principles using the parameter shift rule, and these gradients are prone to exponential vanishing in "barren plateau" regions [13]. Moreover, anomaly detection, by nature, is unsupervised, adding another layer of difficulty since no labeled data is available to guide the training [29], [15]. These two factors – quantum training complexity and unsupervised learning requirements – create a considerable challenge for developing efficient quantum anomaly detection methods [29], [14].

The Gap. *Current quantum-based approaches for anomaly detection fall short because they still require training and often rely on supervised or semi-supervised learning.* These methods involve optimizing quantum circuits with labeled data, which limits their applicability in real-world scenarios where such data is scarce. This dependency on training creates both computational overheads and reduces the flexibility needed for fully unsupervised tasks [16], [13], [21], [12], [29], [14], [1].

Our Solution. To address this, we propose Quorum, the first quantum anomaly detection framework that requires zero training and is designed for unsupervised learning. Quorum leverages quantum principles such as amplitude encoding, random quantum transformations, and SWAP tests to identify anomalies without needing any parameter optimization [16], [13]. By utilizing quantum correlations and random projections, Quorum detects anomalies based on deviations from the statistical norms of quantum transformations [17], [21], [13].

First, Quorum's open-source technique carefully distributes the data into buckets based on the likelihood of anomalous events in the dataset before embedding it into quantum states using amplitude encoding [12], [2]. Then, Quorum applies random quantum transformations to this data and uses a SWAP test to measure similarly between quantum states [13], [16]. Quorum constructs an "embarrassingly parallelizable" ensemble of such random transformations and leverages statistical measures to identify anomalies. This approach is scalable and flexible, allowing for efficient anomaly detection without the computing cost of gradient calculation and training [29], [21].

Quorum's Evaluation. We evaluate Quorum through extensive experiments on various datasets, including medical, industrial, and lexical data. We compare its performance to a state-of-the-art method that uses a quantum neural network (QNN) which relies on training and supervised labels. Our results show that Quorum has a 23% higher average F1 score over the QNN across evaluated datasets. We also provide an ablation study of how Quorum performs with different-sized subsamples (buckets). *Quorum consistently identifies subtle anomalies that the state-of-the-art method may overlook, proving to be an effective zero-training quantum solution for unsupervised anomaly detection [29], [21].*

II. Background

Before we present the design of Quorum, we first provide some brief but necessary background.

A. Quantum Computing

The fundamental unit of quantum computation is the qubit, which exists in a superposition of binary states, represented as $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, where α and β are complex numbers that correspond to the amplitudes of the basis states $|0\rangle$ and $|1\rangle$. The probabilities of measuring the qubit in the $|0\rangle$ or $|1\rangle$ state are $|\alpha|^2$ and $|\beta|^2$, respectively, and these probabilities must

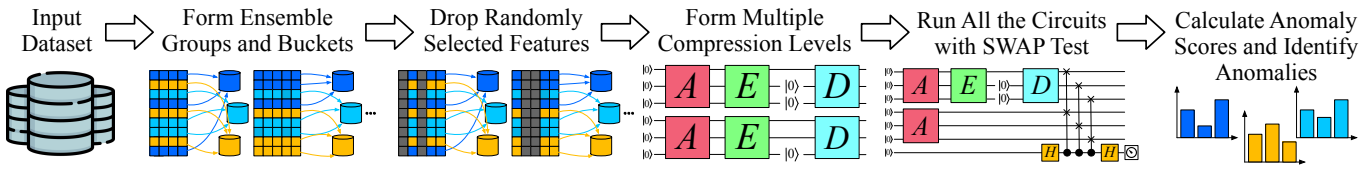

Fig. 1: Visual representation of the steps taken by Quorum to detect the anomalies in a given dataset using quantum computing.

sum to 1: $\|\alpha\|^2 + \|\beta\|^2 = 1$. Qubits can also be entangled, meaning their quantum states are correlated in such a way that they cannot be described independently of one another, which is a key property leveraged in quantum computing applications and algorithms to achieve a quantum advantage [4].

Quantum gates, represented as unitary matrices, are the basic operations applied to qubits. Common single-qubit gates include the parameterized rotation gates, such as the RX, RY, and RZ gates, which rotate a qubit around the x, y, and z axes, respectively. Two-qubit gates, such as the controlled-X (CX) gate, are used to create entanglement between qubits. These single- and two-qubit gates are defined as:

$$RX(\theta) = \begin{pmatrix} \cos\frac{\theta}{2} & -i\sin\frac{\theta}{2} \\ -i\sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{pmatrix}, \quad RY(\theta) = \begin{pmatrix} \cos\frac{\theta}{2} & -\sin\frac{\theta}{2} \\ \sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{pmatrix},$$

$$RZ(\theta) = \begin{pmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{pmatrix}, \quad CX = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$$

Quantum circuits are constructed by applying sequences of these gates to qubits, followed by measurements. Due to the probabilistic nature of quantum measurements, a circuit's execution yields a probability distribution of possible outcomes, which often requires many repetitions (or shots) to obtain statistically significant results [11], [6]

B. Quantum Encoding and SWAP Test

In quantum computing, classical data must be encoded into quantum states. One common method is *Amplitude Encoding*, [35] which maps classical data points $x_0, \ldots, x_{n-1}$ to the quantum state: $|\psi\rangle = \sum_{j=0}^{n-1} x_j |j\rangle$. This method allows n classical data points to be encoded into $\log_2 n$ qubits, enabling efficient representation of high-dimensional data.

The *SWAP Test* is a quantum algorithm used to determine how similar two quantum states are [7], [22]. It computes the inner product $\langle\phi|\psi\rangle$ between two states $|\phi\rangle$ and $|\psi\rangle$. If the states are identical, the test returns a high overlap; otherwise, a lower overlap indicates dissimilarity. This is a key tool in quantum anomaly detection, where dissimilarity between two states may indicate an anomalous data point [22].

C. Anomaly Detection

Anomaly detection involves identifying data points that deviate significantly from the norm [34], [9]. In unsupervised anomaly detection, we assume no labeled data is available, and the goal is to detect anomalies based solely on patterns within the data [5], [30]. Common classical techniques include clustering and Isolation Forests. Clustering groups data points based on similarity, with anomalies detected as points far from cluster centroids [12], [1]. Isolation Forests, a tree-based algorithm, isolates anomalies by recursively splitting the data based on random feature values, where fewer splits indicate an anomaly [19], [18].

In high-dimensional spaces, autoencoders – neural networks trained to reconstruct their input – are often used for anomaly detection [23], [24]. Autoencoders learn compressed representations of normal data, and anomalies are detected when reconstruction errors for new data points are high [1], [20]. This concept extends to the quantum realm, where a quantum autoencoder can be used to compress and decompress quantum states, flagging quantum outliers based on reconstruction errors [16], [13].

III. Related Work

Several quantum-based approaches for anomaly detection have emerged recently, leveraging various quantum algorithms and machine learning techniques. Liu at al. [16] provided early evidence of the speedup and resource efficiency experienced with quantum anomaly detection. Building on this effort, Herr et al.[13] introduce quantum autoencoders trained with generative-adversarial networks. Both methods require structured queries or training processes, which prohibit their adaptability in unsupervised settings. Taking a hybrid quantum-classical approach, Sakhnenko et al. [29] propose supervised training-based solutions that require significant quantum-classical communication. Most recently, Hdaib et al.[12] propose quantum-enhanced anomaly detection with classical post-processing, still requiring circuit training.

Other works provide application-specific anomaly detection. For example, Ngairangbam et al. [21] utilize quantum classifiers that rely on supervised training for high-energy physics applications. On the other, Kukliansky et al.[14] develop quantum anomaly detectors for network anomalies. These approaches rely on domain-specific anomaly properties and are thus not generally applicable. *Thus, there is a strong need for a fully unsupervised, training-free generalized quantum anomaly detection method that leverages the unique strengths of quantum systems without the overhead of classical optimization.*

IV. Design

The design of Quorum leverages quantum dynamics to identify correlations among data features and detect anomalies by comparing quantum-transformed data samples. As shown in Fig. 1, Quorum is structured around several key stages: preprocessing, quantum embedding, bucketing, feature selection, and circuit-based statistical analysis. Each of these steps

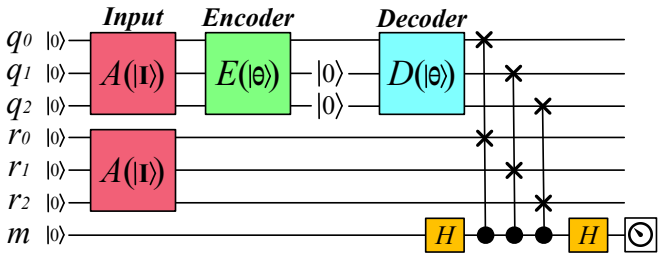

Fig. 2: Quorum's use of SWAP test to determine the similarity between the compressed and original data sample.

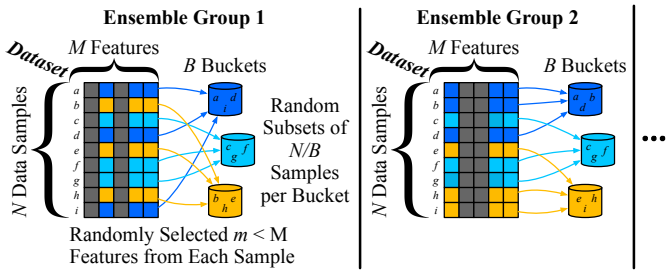

Fig. 3: Quorum's bucketing procedure for distributing the subsampled data samples across different ensemble groups.

contributes to an efficient and scalable framework for anomaly detection using quantum methods.

A. Preprocessing and Normalization

Quorum begins with preprocessing the dataset, which involves a variety of steps depending on the dataset's initial format. This typically includes transforming all non-numeric features into float values (e.g., via hashing), removing any label data that could indicate whether a sample is anomalous, and performing a range-based normalization. The normalization process is essential to ensure that all features contribute equally to the quantum state, which is critical for the subsequent quantum encoding.

Given a dataset with M features, Quorum normalizes each feature so that its maximum possible value is $\frac{1}{M}$. This ensures that the sum of the squares of all features for any sample does not exceed 1. The normalization is performed as follows:

$$\text{normalized feature value} = \frac{\text{raw feature value for sample}}{\text{max feature value}} \times \frac{1}{M}$$

This normalization serves two key purposes. First, it equalizes the contribution of all features to the final quantum state, preventing any feature from dominating due to its scale. Second, it simplifies the process of amplitude embedding for quantum states, as the normalized values now range between $[0, \frac{1}{M}]$ for all of the features.

B. Quantum Embedding

After normalization, Quorum embeds the data into quantum states using amplitude encoding. The normalized feature values are squared to convert them into probabilities, and an "overflow state" probability is added to account for any remaining probability mass, ensuring the total probability sums to 1. This ensures that the total probability mass of the quantum state is preserved.

A quantum circuit is then created to prepare a state vector corresponding to these probability amplitudes using amplitude embedding. This process is repeated for each data point (with 2^n features) using two sets of n qubits within the same circuit, creating two identical encodings: one for the transformation and one as a reference (Fig. 2). This dual encoding allows us to compare the transformed data with the original using a SWAP test (we'll discuss in Sec. IV-D why this is required).

Each quantum circuit consists of $2n + 1$ qubits, where the additional qubit serves as an ancilla for reading the SWAP test result. The SWAP test measures the similarity between the original and transformed quantum states, preserving the relative magnitudes of the features in the quantum state.

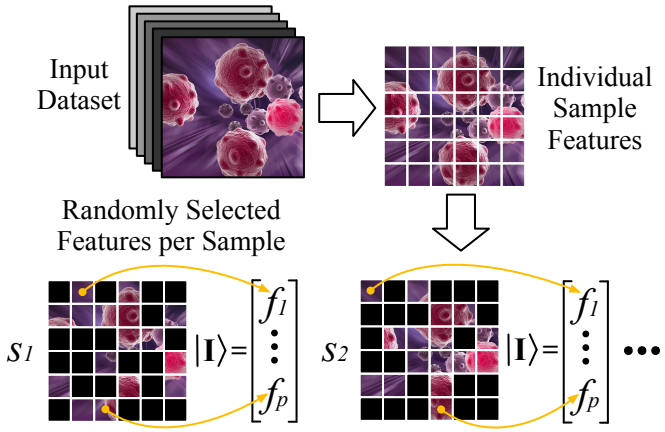

Fig. 4: Quorum subsamples features from the input space.

C. Bucketing and Feature Selection

Following the embedding process, Quorum employs a bucketing strategy and performs feature selection to prepare the data for anomaly detection. As shown in Fig. 3, the dataset is divided into a series of B random subsets, or buckets. Given a dataset of size N, there are N/B buckets. This bucketing strategy enhances the visibility of anomalies by allowing data points to be compared against smaller, more localized subsets of the dataset. The size of the buckets (B) is determined based on the total number of data points and the estimated proportion of anomalies, both of which determine the probability of having at least one anomaly in each bucket. By distributing the data into smaller buckets, Quorum increases the contrast between normal and anomalous points, making it easier to detect outliers (anomalous data points).

Feature selection is performed after bucketing to ensure that the data samples can fit on quantum circuits, with Quorum using a uniform random selection strategy (Fig. 4). Unlike traditional downsampling or dimensionality reduction techniques like Principal Component Analysis (PCA), random selection offers several advantages. *It is computationally faster, avoids bias towards features that might not indicate anomalies, and allows exploration of feature combinations that might*

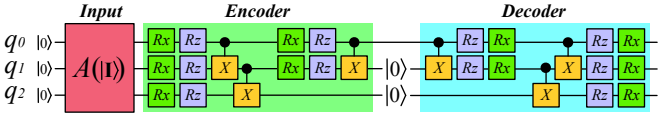

Fig. 5: The ansatz leveraged by Quorum for its autoencoder to establish correlations across input features.

otherwise be overlooked. For each quantum circuit with n qubits, Quorum selects $m = (2^n - 1)$ features from the dataset, leaving room for an overflow state. This ensures the selected features fit within the quantum state space of n qubits.

D. Quantum Circuit Design and Ansatz

The core of Quorum's anomaly detection framework lies in the design of its quantum circuits. Unlike traditional quantum autoencoders, which train parameterized gate angles to optimize encoding-decoding processes, Quorum does not rely on learning optimal parameters. Instead, it utilizes random quantum transformations and applies statistical analysis to detect anomalies without training.

Each quantum circuit begins with an amplitude encoding of the data, as described above. The first set of encoded qubits is then passed through an ansatz, which consists of layers of RX and RZ rotations and CNOT gates. The ansatz performs random transformations on the encoded data, ensuring a high degree of variability in the quantum states (Fig. 5).

The ansatz includes three main components: an encoder circuit with randomly initialized parameters, a partial reset operation that simulates an information bottleneck by resetting a subset of qubits, and a decoder circuit that applies the inverse of the encoder. The random angles for the quantum gates in the encoder are initialized from a uniform distribution $U(0, 2\pi)$, and the decoder negates these angles to revert the transformations. A SWAP test is then performed between the transformed and original quantum states to measure their similarity. This similarity score forms the basis for detecting anomalies, as normal and anomalous data points will behave differently under random quantum transformations.

This design enables Quorum to compress data through the autoencoder and then decode it on the other end of the reset, with the anticipation that anomalous data would be more likely to deviate from the original state when the two states are compared using the SWAP test.

E. Ensemble Groups and Statistical Analysis

To ensure robustness, Quorum processes each data point through multiple ensemble groups. Each ensemble group involves randomly initialized quantum circuits (θs), ensuring that the data points are processed differently each time. Additionally, the dataset is divided into new buckets for each ensemble group, providing different perspectives on the data.

Each ensemble group also explores multiple compression levels (Fig. 6). The compression level, determined by the number of qubits reset in the partial reset operation, varies systematically from the highest compression (fewest qubits retained) to the lowest (most qubits retained). Quorum creates a multi-dimensional view of each data point's behavior under various quantum transformations. This approach leverages the principle of random projections, where projecting data into multiple random subspaces can reveal structural information that might not be apparent in any single projection. Each ensemble group, with its unique combination of buckets, feature subsets, circuit parameters, and compression levels, represents a different "quantum projection" of the data.

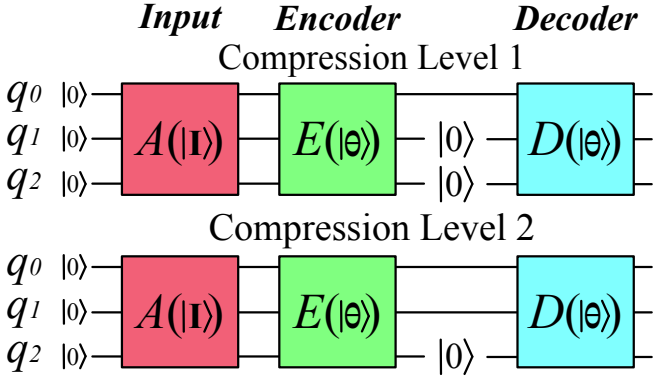

Fig. 6: Quorum leverages multiple compression levels across different ensemble groups to improve anomaly detection.

Statistical analysis is then performed on the SWAP test outputs. *A key innovation in Quorum's approach lies in how these quantum circuit outputs are analyzed to detect anomalies.* For each bucket and each run, the mean and standard deviation of the SWAP test results are calculated. The anomaly score for each data point is derived by calculating the normalized deviation from the bucket mean divided by the standard deviation. These deviations are summed across all runs and buckets, producing an overall anomaly score for each data point (Fig. 7). A higher score indicates a greater likelihood of an anomaly.

The use of random angles in the quantum circuit is a critical aspect of Quorum's design. Rather than learning to reconstruct inputs accurately, this approach creates a random projection of the data in high-dimensional Hilbert space. The statistical analysis then captures how differently each data point behaves under these random quantum transformations compared to the average behavior of points in its various buckets.

This design allows Quorum to detect anomalies without explicitly learning the structure of normal data or performing any optimization, instead relying on the statistical properties of how data points respond to random quantum transformations. *The effectiveness of this approach lies in the fact that anomalous data points tend to behave differently under these transformations compared to normal data points, as the randomizations add more deviation to anomalous data compared to normal data.* By aggregating these behaviors across multiple random initializations, buckets, and compression levels, Quorum aims to build a comprehensive profile of each data point's response to quantum transformations,

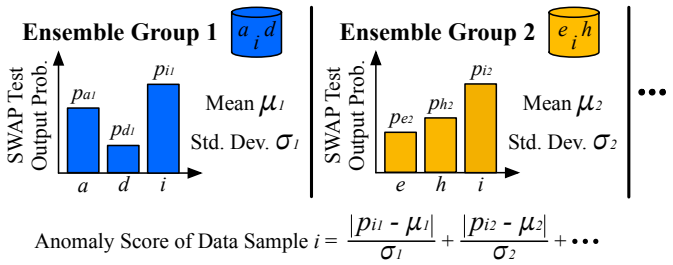

Fig. 7: Quorum's anomaly score calculates the statistical deviation of a data sample across ensemble groups.

potentially unveiling subtle anomalies that more conventional techniques might miss.

F. Scalability and Flexibility

A significant advantage of Quorum's design is its inherent scalability and flexibility. While our initial experiments utilized 3-qubit encodings (resulting in 7-qubit circuits), the approach can be readily scaled to accommodate larger encodings. For instance, moving to 4-qubit or higher encodings would introduce additional "moments" to our results for each compression level, potentially capturing even more nuanced relationships in the data. This scalability allows users to tailor the depth and complexity of the quantum transformations to their specific needs and computational resources.

Furthermore, the design of Quorum lends itself to extensive parallelization. Each ensemble group is entirely independent of the others, making the technique embarrassingly parallel. The scalability of Quorum extends beyond just increasing the number of qubits or ensemble groups. The flexibility in choosing the number of compression levels, the size of buckets, and the number of features selected allows users to fine-tune the balance between computational cost and the granularity of anomaly detection.

In the following section, we delve into the methodology used to obtain our evaluated results.

V. Experimental Methodology

Experimental Setup. We evaluate Quorum using Qiskit Aer's quantum circuit simulator. We use Python 3.10.12 and Qiskit 1.2, IBM's quantum computing language [3], which is used to run simulations of quantum circuits locally. Each quantum circuit is generated via a Qiskit QuantumCircuit object, and the quantum circuits are run through noiseless simulations. We run simulations with Quorum and competitors on a local research cluster with Ubuntu 22.04.2 LTS on a 32-core 2.0 GHz AMD EPYC 7551P processor with 32 GB RAM.

We also perform noisy simulations, where we model our hardware after IBM's Brisbane quantum computer to provide realistic error rates. The noise parameters were obtained directly from IBM's Brisbane backend specifications; median properties are as follows: coherence times ($T_1 = 230.42$ μs, $T_2 = 143.41$ μs), gate errors (single-qubit SX gate error = 2.274×10^{-4}, two-qubit gate error = 2.903×10^{-3}), and readout error (1.38×10^{-2}). Note: due to the over 100,000 runs required for Quorum's evaluation, it was cost-prohibitive to execute on real hardware. Thus, we use faithful simulations.

TABLE I: Datasets used for Quorum's evaluation. The rightmost column refers to the likelihood of at least one anomaly being placed in each bucket.

Dataset	Samples	Anomalies	Features	Pr [Anomaly ∈ Bucket]
Breast Cancer	367	10	30	0.75
Pen-Global	809	90	16	0.6
Letter	533	33	32	0.95
Power Plant	1,000	30	5	0.75

Datasets. We evaluate Quorum on four distinct datasets representing different anomaly detection scenarios (Table I). Three of the four datasets are directly derived from a related and widely-cited survey of classical unsupervised anomaly detection techniques by Goldstein and Uchida [10]. The fourth dataset, which contains measurements taken from a combined cycle power plant, was taken from UCI's machine learning repository [32]. For the power plant dataset, we inserted "plausible" anomalies into the dataset based on ranges of values that are possible for each feature. All datasets have labels stripped for all operations until the evaluation is performed to facilitate unsupervised anomaly detection.

Experimental Framework. Each experiment consists of multiple ensemble groups, where an ensemble group represents a complete run of Quorum with different random initializations. We use 3-qubit encodings for our primary experiments, resulting in 7-qubit circuits (including the ancilla qubit). We chose this circuit size mainly due to limitations on our computational resources as we ran simulations on our local systems. For each dataset, we execute with 1,000 ensemble members, with each member using different random bucket assignments and feature selections. We executed 4,096 shots per circuit for measurements. Increasing both shot count and ensemble members has significant impacts on performance, with benefits diminishing as they increase past a certain point. For noisy simulations, we similarly executed 4,096 shots per circuit. We use different target probabilities for bucket size determination (see Table I) and provide an ablation study on the effect of bucket sizes for each of the different datasets.

Evaluation Metrics. We evaluate Quorum's performance using several metrics: (1) **Detection Rate/Accuracy** at various percentile thresholds, measuring the fraction of true anomalies captured in the top k% of anomaly scores. We chose to include this metric to showcase how well our technique separates anomalies from the rest of the dataset. (2) **Precision**, calculated as the ratio of correctly identified anomalies to the total number of samples flagged as anomalous. (3) **Recall**, measured as the ratio of correctly identified anomalies to the total number of true anomalies in the dataset. And (4) **F1 Score**, which is the harmonic mean of precision and recall.

Competitive Techniques. We compare Quorum against a

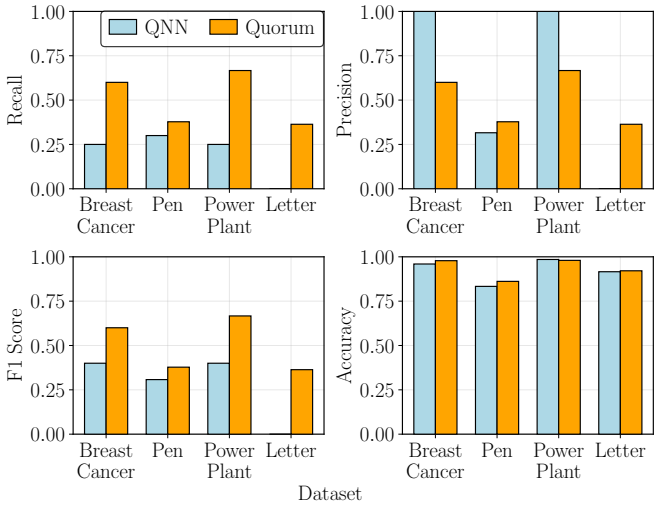

Fig. 8: Comparison of performance metrics between QNN and Quorum across four datasets. Note that the QNN did not detect any anomalies for the letter dataset, and thus, precision, recall, and F1 scores are all zero.

state-of-the-art quantum anomaly detection technique that utilizes quantum neural networks to find anomalies in labeled datasets [14] (we refer to this technique as "QNN"). The technique is the most suitable among competitors as it shows improved performance over them and is designed for noisy quantum hardware, making it a practical benchmark. We adapted QNN for generic use since it was initially developed for network anomaly detection. Further, comparing against QNN's supervised, training-based approach helps demonstrate the advantages of our zero-training unsupervised method, particularly in scenarios where labeled data is unavailable. Note: no unsupervised, zero-training work exists to compare against as Quorum is the first of its kind in this domain.

VI. Evaluation

Our experimental results demonstrate the efficacy of Quorum's anomaly detection approach across diverse datasets.

Flagship Results. Quorum demonstrates balanced anomaly detection performance across all evaluated datasets – Fig. 8. The recall measurements particularly highlight Quorum's strengths, where it consistently outperforms the QNN method. While the QNN achieves perfect precision scores on both the breast cancer and power plant datasets, this comes at the significant cost of being overly conservative in anomaly detection, leading to its notably poor recall performance. Quorum thus achieves superior F1 performance across every dataset tested – 23% higher on average. Quorum's more nuanced approach leads to better overall detection capabilities while maintaining comparable accuracy levels across all datasets. *These results demonstrate that Quorum provides an effective approach to anomaly detection, successfully balancing precision and recall without sacrificing overall classification performance.*

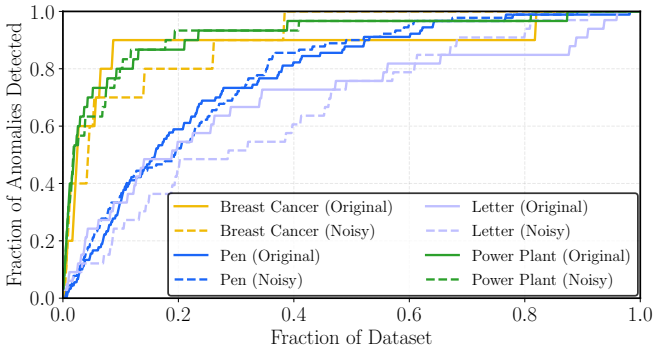

Fig. 9: Quorum groups anomalies consistently amongst data points with the highest absolute average deviations.

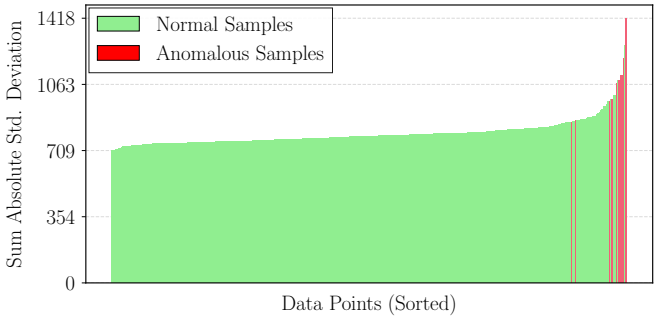

Fig. 10: An example of how Quorum separates anomalies from normal samples on the breast cancer dataset (16K shots).

Detection Rates. In Fig. 9, Quorum's detection rate curves exhibit notably steep initial gradients, particularly for the breast cancer and power plant datasets, which achieve approximately 80% detection rate within the first 10% of the highest deviation samples in noiseless simulations (see Fig. 10 for a detailed look at how Quorum grouped the breast cancer data samples). The performance hierarchy among datasets likely reflects their inherent separability characteristics, with the breast cancer dataset showing the most distinctive anomaly patterns (achieving near-perfect detection at the 10th percentile), followed by power plant data. The letter and pen datasets, while requiring a larger percentile of samples for complete detection, still maintain clear separation from random performance, achieving a roughly 60% detection rate within the top 20% of sample deviations for noiseless runs. When subjected to realistic quantum noise, Quorum demonstrates a high degree of resilience. Noisy simulations closely track their noiseless counterparts across all datasets, with only minimal degradation in performance. *Such inherent noise resilience is a significant advantage of Quorum for near-term applications, as Quorum can maintain its effectiveness even on noisy hardware without requiring high-overhead error mitigation or correction methods.*

Bucket Size Ablation. Analysis of F1 scores across different bucket size configurations in Table II reveals that, as expected, very small bucket sizes generally lead to degraded perfor-

TABLE II: F1 Scores for different bucket sizes (p is the probability of at least one anomaly in a bucket).

Dataset	$p = 0.5$	$p = 0.6$	$p = 0.75$	$p = 0.95$	$p = 0.98$
Breast Cancer	0.500	0.500	0.600	0.500	0.600
Pen Digits	0.333	0.389	0.367	0.389	0.389
Letter	0.152	0.182	0.242	0.273	0.273
Power Plant	0.600	0.600	0.633	0.533	0.600

mance. However, we observe that moderately sized buckets oftentimes outperform larger ones – for instance, the letter dataset achieves its peak F1 score of 0.273 at $p = 0.95$, while the breast cancer and power plant datasets show improved performance at $p = 0.75$. This behavior can be explained by the trade-off between statistical significance and local sensitivity - while larger buckets provide more robust statistical estimates, smaller buckets may better capture localized anomaly patterns by reducing the "averaging out" effect of mixing anomalous and normal samples from the datasets.

VII. Conclusion

In this paper, we introduced Quorum, a novel quantum anomaly detection framework that operates without requiring any training. The framework's design, which incorporates strategic data bucketing, feature selection, and ensemble analysis, proves particularly effective at identifying anomalies while remaining computationally efficient through its inherent parallelizability. Our evaluation shows that Quorum can achieve up to 80% detection rate within the first 10% of highest-deviation samples, demonstrating strong anomaly-detection power, and high resilience to noise on near term quantum systems. These results suggest that Quorum represents a promising and wholly quantum step forward in anomaly detection.

Quorum's code is open-sourced at: `https://github.com/positivetechnologylab/Quorum`.

Acknowledgment

This work was supported by Rice University, the Rice University George R. Brown School of Engineering and Computing, and the Rice University Department of Computer Science. This work was supported by the DOE Quantum Testbed Finder Award DE-SC0024301. This work was also supported by the Ken Kennedy Institute and Rice Quantum Initiative, which is part of the Smalley-Curl Institute.

References

[1] Mohiuddin Ahmed, Abdun Naser Mahmood, and Jiankun Hu. A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60:19–31, 2016.

[2] Javier Alcazar, Vicente Leyton-Ortega, and Alejandro Perdomo-Ortiz. Classical versus quantum models in machine learning: insights from a finance application. *Machine Learning: Science and Technology*, 1(3):035003, 2020.

[3] Gadi Aleksandrowicz, Thomas Alexander, Panagiotis Barkoutsos, Luciano Bello, Yael Ben-Haim, David Bucher, Francisco Jose Cabrera-Hernández, Jorge Carballo-Franquis, Adrian Chen, Chun-Fu Chen, Jerry M. Chow, Antonio D. Córcoles-Gonzales, Abigail J. Cross, Andrew Cross, Juan Cruz-Benito, Chris Culver, Salvador De La Puente González, Enrique De La Torre, Delton Ding, Eugene Dumitrescu, Ivan Duran, Pieter Eendebak, Mark Everitt, Ismael Faro Sertage, Albert Frisch, Andreas Fuhrer, Jay Gambetta, Borja Godoy Gago, Juan Gomez-Mosquera, Donny Greenberg, Ikko Hamamura, Vojtech Havlicek, Joe Hellmers, Łukasz Herok, Hiroshi Horii, Shaohan Hu, Takashi Imamichi, Toshinari Itoko, Ali Javadi-Abhari, Naoki Kanazawa, Anton Karazeev, Kevin Krsulich, Peng Liu, Yang Luh, Yunho Maeng, Manoel Marques, Francisco Jose Martín-Fernández, Douglas T. McClure, David McKay, Srujan Meesala, Antonio Mezzacapo, Nikolaj Moll, Diego Moreda Rodríguez, Giacomo Nannicini, Paul Nation, Pauline Ollitrault, Lee James O'Riordan, Hanhee Paik, Jesús Pérez, Anna Phan, Marco Pistoia, Viktor Prutyanov, Max Reuter, Julia Rice, Abdón Rodríguez Davila, Raymond Harry Putra Rudy, Mingi Ryu, Ninad Sathaye, Chris Schnabel, Eddie Schoute, Kanav Setia, Yunong Shi, Adenilton Silva, Yukio Siraichi, Seyon Sivarajah, John A. Smolin, Mathias Soeken, Hitomi Takahashi, Ivano Tavernelli, Charles Taylor, Pete Taylour, Kenso Trabing, Matthew Treinish, Wes Turner, Desiree Vogt-Lee, Christophe Vuillot, Jonathan A. Wildstrom, Jessica Wilson, Erick Winston, Christopher Wood, Stephen Wood, Stefan Wörner, Ismail Yunus Akhalwaya, and Christa Zoufal. Qiskit: An Open-source Framework for Quantum Computing, January 2019.

[4] MD Arafath and A Niranjil Kumar. Quantum computing based neural networks for anomaly classification in real-time surveillance videos. *Computer Systems Science & Engineering*, 46(2), 2023.

[5] Sounak Bhowmik and Himanshu Thaplival. Quantum machine learning for anomaly detection in consumer electronics. In *2024 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, pages 544–550. IEEE, 2024.

[6] Giuseppe Bisicchia, Jose García-Alonso, Juan M. Murillo, and Antonio Brogi. Distributing quantum computations, by shots. In Flavia Monti, Stefanie Rinderle-Ma, Antonio Ruiz Cortés, Zibin Zheng, and Massimo Mecella, editors, *Service-Oriented Computing*, pages 363–377, Cham, 2023. Springer Nature Switzerland.

[7] Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Phys. Rev. Lett.*, 87:167902, Sep 2001.

[8] Wisam Elmasry and Mohammed Wadi. Enhanced anomaly-based fault detection system in electrical power grids. *International Transactions on Electrical Energy Systems*, 2022(1):1870136, 2022.

[9] Joseph A Gallego-Mejia, Oscar A Bustos-Brinez, and Fabio A González. Inqmad: incremental quantum measurement anomaly detection. In *2022 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 787–796. IEEE, 2022.

[10] Markus Goldstein and Seiichi Uchida. A comparative evaluation of unsupervised anomaly detection algorithms for multivariate data. *PLOS ONE*, 11(4):1–31, 04 2016.

[11] Mayank Gupta and Manisha J. Nene. Quantum computing: A measurement and analysis review. *Concurrency and Computation: Practice and Experience*, 33(20):e6344, 2021.

[12] Moe Hdaib, Sutharshan Rajasegarar, and Lei Pan. Quantum autoencoder frameworks for network anomaly detection. In *International Conference on Neural Information Processing*, pages 69–82. Springer, 2023.

[13] Daniel Herr, Benjamin Obert, and Matthias Rosenkranz. Anomaly detection with variational quantum generative adversarial networks. *Quantum Science and Technology*, 6(4):045004, 2021.

[14] Alon Kukliansky, Marko Orescanin, Chad Bollmann, and Theodore Huffmire. Network anomaly detection using quantum neural networks on noisy quantum computers. *IEEE Transactions on Quantum Engineering*, 2024.

[15] W Bernard Lee and Anthony G Constantinides. Computational results for a quantum computing application in real-life finance. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, volume 1, pages 414–423. IEEE, 2023.

[16] Nana Liu and Patrick Rebentrost. Quantum machine learning for quantum anomaly detection. *Physical Review A*, 97(4):042315, 2018.

[17] Javier Mancilla and Christophe Pere. A preprocessing perspective for quantum machine learning classification advantage in finance using nisq algorithms. *Entropy*, 24(11):1656, 2022.

[18] Elisa Marcelli, Tommaso Barbariol, Davide Sartor, and Gian Antonio Susto. Active learning-based isolation forest (alif): Enhancing anomaly detection with expert feedback. *Information Sciences*, page 121012, 2024.

[19] Antonella Mensi, Alessio Franzoni, David MJ Tax, and Manuele Bicego. An alternative exploitation of isolation forests for outlier detection. In *Structural, Syntactic, and Statistical Pattern Recognition: Joint IAPR International Workshops, S+ SSPR 2020, Padua, Italy, January 21–22, 2021, Proceedings*, pages 34–44. Springer, 2021.

[20] Naji Najari, Samuel Berlemont, Grégoire Lefebvre, Stefan Duffner, and Christophe Garcia. Robust variational autoencoders and normalizing flows for unsupervised network anomaly detection. In *International Conference on Advanced Information Networking and Applications*, pages 281–292. Springer, 2022.

[21] Vishal S Ngairangbam, Michael Spannowsky, and Michihisa Takeuchi. Anomaly detection in high-energy physics using a quantum autoencoder. *Physical Review D*, 105(9):095004, 2022.

[22] Vishal S. Ngairangbam, Michael Spannowsky, and Michihisa Takeuchi. Anomaly detection in high-energy physics using a quantum autoencoder. *Phys. Rev. D*, 105:095004, May 2022.

[23] Oyebade K Oyedotun and Djamila Aouada. A closer look at autoencoders for unsupervised anomaly detection. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3793–3797. IEEE, 2022.

[24] Krishna Patra, Rabi Narayan Sethi, and Dhiren Kkumar Behera. Anomaly detection in rotating machinery using autoencoders based on bidirectional lstm and gru neural networks. *Turkish Journal of Electrical Engineering and Computer Sciences*, 30(4):1637–1653, 2022.

[25] Marco Pistoia, Syed Farhan Ahmad, Akshay Ajagekar, Alexander Buts, Shouvanik Chakrabarti, Dylan Herman, Shaohan Hu, Andrew Jena, Pierre Minssen, Pradeep Niroula, et al. Quantum machine learning for finance iccad special session paper. In *2021 IEEE/ACM international conference on computer aided design (ICCAD)*, pages 1–9. IEEE, 2021.

[26] Tahereh Pourhabibi, Kok-Leong Ong, Booi H. Kam, and Yee Ling Boo. Fraud detection: A systematic literature review of graph-based anomaly detection approaches. *Decision Support Systems*, 133:113303, 2020.

[27] John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.

[28] John Preskill. Quantum computing 40 years later. In *Feynman Lectures on Computation*, pages 193–244. CRC Press, 2023.

[29] Alona Sakhnenko, Corey O'Meara, Kumar JB Ghosh, Christian B Mendl, Giorgio Cortiana, and Juan Bernabé-Moreno. Hybrid classical-quantum autoencoder for anomaly detection. *Quantum Machine Intelligence*, 4(2):27, 2022.

[30] Julian Schuhmacher, Laura Boggia, Vasilis Belis, Ema Puljak, Michele Grossi, Maurizio Pierini, Sofia Vallecorsa, Francesco Tacchino, Panagiotis Barkoutsos, and Ivano Tavernelli. Unravelling physics beyond the standard model with classical and quantum anomaly detection. *Machine Learning: Science and Technology*, 4(4):045031, 2023.

[31] Yehui Tang, Junchi Yan, Guoqiang Hu, Baohua Zhang, and Jinzan Zhou. Recent progress and perspectives on quantum computing for finance. *Service Oriented Computing and Applications*, 16(4):227–229, 2022.

[32] Pnar Tfekci and Heysem Kaya. Combined Cycle Power Plant. UCI Machine Learning Repository, 2014. DOI: https://doi.org/10.24432/C5002N.

[33] Sohum Thakkar, Skander Kazdaghli, Natansh Mathur, Iordanis Kerenidis, André J Ferreira-Martins, and Samurai Brito. Improved financial forecasting via quantum machine learning. *Quantum Machine Intelligence*, 6(1):27, 2024.

[34] Maida Wang, Anqi Huang, Yong Liu, Xuming Yi, Junjie Wu, and Siqi Wang. A quantum-classical hybrid solution for deep anomaly detection. *Entropy*, 25(3):427, 2023.

[35] M. Weigold, J. Barzen, F. Leymann, and M. Salm. Encoding patterns for quantum algorithms. *IET Quantum Communication*, 2:141–152, 2021.

[36] Cuicui Zhang, Jiali Sun, Ruixuan Lu, and Peng Wang. Anomaly detection model of power grid data based on stl decomposition. In *2021 IEEE 5th Information Technology,Networking,Electronic and Automation Control Conference (ITNEC)*, volume 5, pages 1262–1265, 2021.