

Qsco: A Quantum Scoring Module for Open-Set Supervised Anomaly Detection

Yifeng Peng¹, Xinyi Li¹, Zhiding Liang², Ying Wang^{1*}

¹Stevens Institute of Technology, Hoboken, NJ, USA

²Rensselaer Polytechnic Institute, Troy, NY, USA

{ypeng21, xli215, ywang6}@stevens.edu, liangz9@rpi.edu

Abstract

Open set anomaly detection (OSAD) is a crucial task that aims to identify abnormal patterns or behaviors in data sets, especially when the anomalies observed during training do not represent all possible classes of anomalies. The recent advances in quantum computing in handling complex data structures and improving machine learning models herald a paradigm shift in anomaly detection methodologies. This study proposes a Quantum Scoring Module (Qsco), embedding quantum variational circuits into neural networks to enhance the model’s processing capabilities in handling uncertainty and unlabeled data. Extensive experiments conducted across eight real-world anomaly detection datasets demonstrate our model’s superior performance in detecting anomalies across varied settings and reveal that integrating quantum simulators does not result in prohibitive time complexities. At the same time, the experimental results under different noise models also prove that Qsco is a noise-resilient algorithm. Our study validates the feasibility of quantum-enhanced anomaly detection methods in practical applications.

Introduction

Precisely identifying anomalous patterns within extensive and intricate datasets is paramount, especially when data labels are nonexistent, or the underlying data distributions remain undefined. Such conditions delineate the open set anomaly detection (OSAD) field, characterized by its necessity to function robustly across varied and often unforeseen data landscapes. Historically, the challenges inherent in OSAD have been addressed through conventional machine learning strategies (Ding, Pang, and Shen 2022; Li et al. 2021a; Liznerski et al. 2020; Pang et al. 2021; Pang, Shen, and Van Den Hengel 2019). While these methods have proven effective, they increasingly encounter constraints related to data’s escalating scale and complexity.

The advent of quantum computing has foreboded new prospects for augmenting machine learning architectures (Pan et al. 2023; Caro et al. 2022). Quantum computing offers substantial enhancements in computational efficiency and capacity, particularly adept at managing the




Figure 1: Visualization of the anomalies and normal data predicted by DRA (Ding, Pang, and Shen 2022) and **Qsco (Ours)** in the MVTEC AD (Bergmann et al. 2019) dataset (carpet subset). The parameter ℓ controls the depth of the variational circuit in Qsco. With the correct ℓ level, Qsco enhances DRA’s ability to distinguish boundaries between anomalies and normal data. However, if ℓ is too high ($\ell = 3$), it results in over-fitting, as shown in the last graph, while if ℓ is too low ($\ell = 1$), it leads to under-fitting.

high-dimensional data spaces prevalent in anomaly detection tasks. The synergy of quantum computing with machine learning, commonly referred to as quantum machine learning (QML) (Farhi and Neven 2018; Alchieri et al. 2021; Peng et al. 2024), introduces innovative approaches that may surmount some of the fundamental limitations of classical methodologies, including scalability, complexity, and uncertainty. With advancements in quantum hardware, QML is poised for revolutionary applications in AI domains like quantum natural language processing (QNLP) (Abbaszade et al. 2021), quantum computer vision (QCV) (Wei et al. 2023), and quantum anomaly detection (QAD) (Liu and Rebertrost 2018).

This research introduces the Quantum Scoring Module (Qsco), a state-of-the-art quantum model of quantum variational circuits with different rotation and CNOT gates. This model aims to exploit the quantum-enhanced computational power to refine the detection of obscure patterns and anomalies in open-set contexts. With the help of entanglement characteristics in quantum computing, we can discover the distribution of unknown anomalies and solve distribution problems that classical computing cannot simulate. We provide initial findings that indicate the Qsco model’s enhanced

*Corresponding Author.

efficacy in processing high-dimensional data relative to traditional models.

In this paper, we investigate the application of QML in OSAD and endeavor to serve as a foundational text for subsequent inquiries into the fusion of quantum and classical computing paradigms. Through this scholarly endeavor, we aspire to forge the development of novel anomaly detection frameworks that are more competent, precise, and equipped to handle the multifaceted demands of contemporary data landscapes.

Related Work

Open Set Anomaly Detection The recent advancements in open set recognition have led to various innovative approaches to detect and classify unknown samples effectively. (Liu et al. 2019) introduced the Margin Learning Embedded Prediction (MLEP) framework, emphasizing the importance of margin learning for enhancing open set recognition. Building on this, (Acintoae et al. 2022) proposed a general open-set method, GOAD, which uses random affine transformations to extend the applicability of transformation-based methods to non-image data, thus relaxing current generalization assumptions. (Sun et al. 2020) tackled the limitations of variational auto-encoders (VAEs) in providing discriminative representations for known classifications by proposing Conditional Gaussian Distribution Learning (CGDL), which enhances open set recognition. (Kong and Ramanan 2021) demonstrated that a well-chosen GAN discriminator on real outlier data can achieve state-of-the-art results and further improved performance by augmenting real open training examples with adversarially synthesized "fake" data, culminating in the OpenGAN framework. (Yoshihashi et al. 2019) focused on joint classification and reconstruction of input data through their Classification-Reconstruction learning for Open-Set Recognition (CROSR), which utilizes latent representations to enhance the separation of unknowns from knowns without compromising known-class classification accuracy. Lastly, (Zhang et al. 2020) proposed the OpenHybrid framework, which combines an encoder, a classifier, and a flow-based density estimator to detect unknown samples while classifying inliers accurately and effectively.

Supervised Anomaly Detection Recent anomaly detection advancements have introduced various innovative methods and frameworks to improve how anomalies are identified and classified across various domains. (Görnitz et al. 2013) established a foundation with their supervised approach, paving the way for further research such as (Ruff et al. 2019), who enhanced accuracy through deep semi-supervised methods leveraging both labeled and unlabeled data. (Li et al. 2021b) introduced Cutpaste, a self-supervised technique that emphasized the growing relevance of self-supervision in anomaly detection.

Further innovations include (Akcay, Atapour-Abarghouei, and Breckon 2019), who developed Ganomaly using GANs for semi-supervised anomaly detection, and (Han et al. 2022), who created Adbench, a comprehensive benchmark for method evaluation. In video anomaly detection, (Georgescu et al. 2021) and (Tian et al.

Model	Algorithm	Supervised	VQC
(Liu and Rebertrost 2018)	Quantum PCA	○	✗
(Guo et al. 2022)	AE	✗	✗
(Kottmann et al. 2021)	VQC based	✗	✓
(Guo et al. 2023)	Quantum LOF	✗	✗
(Wang et al. 2023)	QHDNN	✓	✓
(Moro and Prati 2023)	AQAs	✓	●
Ours (Qsco)	VQC based	✓	✓

○ is semi-supervised and ● is both of VQC and quantum mechanics.

Table 1: Overview of Quantum Anomaly Detection Algorithms.

2021) demonstrated the effectiveness of self-supervised and weakly-supervised learning for handling complex data scenarios and robust temporal feature extraction, respectively.

Additionally, (Tang et al. 2024) and (Yao et al. 2023) contributed by benchmarking graph-based methods with Gadbench and enhancing supervised detection with boundary-guided contrastive learning. In the realm of weak supervision, (Cho et al. 2023) and (Bozorgtabar and Mahapatra 2023) explored relational and attention-conditioned learning approaches, while (Chen et al. 2023) combined magnitude-contrastive learning with glance-and-focus mechanisms for video anomaly detection. (Sultani, Chen, and Shah 2018) focused on real-time anomaly detection, highlighting the necessity for efficient, real-time solutions.

Quantum Anomaly Detection Quantum anomaly detection utilizes principles of quantum computing to significantly enhance the identification of data outliers, which is crucial for fraud detection, network security, and medical diagnostics applications. This emerging field has seen promising developments, as shown in Table 1: (Liu and Rebertrost 2018) introduced quantum algorithms for kernel principal component analysis (PCA) and one-class support vector machine (SVM), which can operate with resources logarithmic in the dimensionality of quantum states. (Guo et al. 2022) proposed an efficient quantum algorithm for the computationally demanding ADDE algorithm, using Amplitude Estimation (AE) to handle large datasets effectively. Similarly, (Kottmann et al. 2021) developed a variational quantum anomaly detection algorithm capable of extracting phase diagrams from quantum simulations without prior physical knowledge, fully operable on the quantum device used for the simulations. Extending further, (Guo et al. 2023) transformed the Local Outlier Factor (LOF) algorithm into a quantum version, achieving exponential speedup on data dimensions and polynomial speedup on data quantity. Additionally, (Wang et al. 2023) engineered a quantum-classical hybrid DNN (QHDNN) that learns from normal images to isolate anomalies, exploring various quantum layer architectures and implementing a VQC-based solution. Lastly, (Moro and Prati 2023) explored the potential

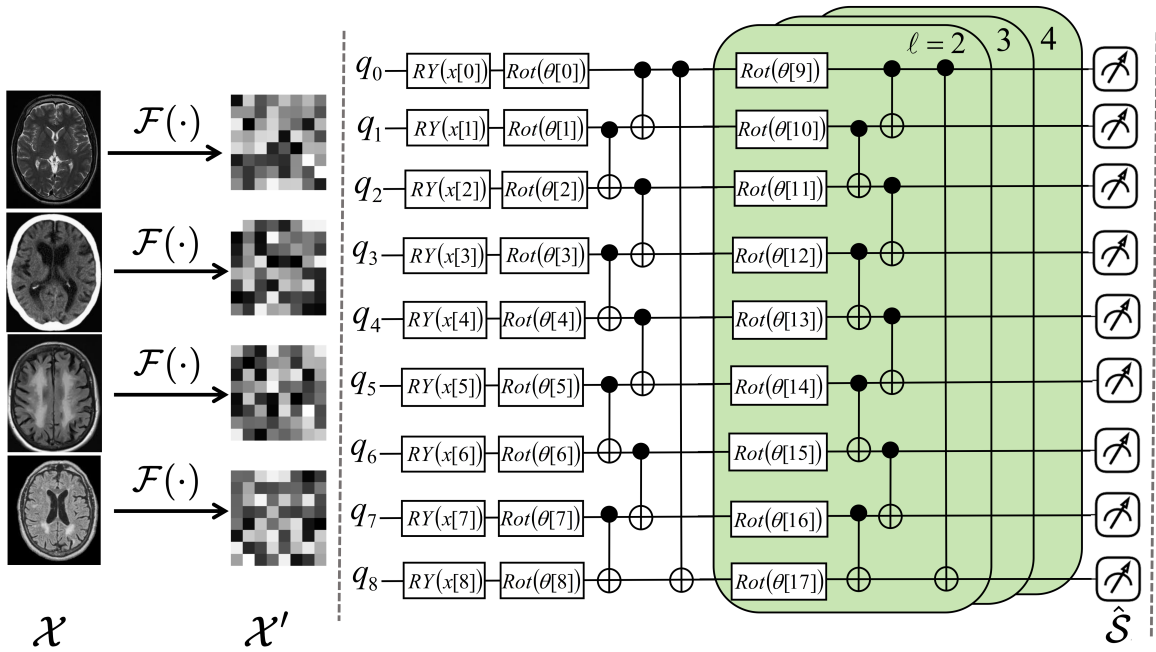


Figure 2: **Overview of our proposed Qsco module.** The left-hand section displays dataset samples from open set data, denoted as the input \mathcal{X} , which are fed into the feature extractor $\mathcal{F}(\cdot)$ to obtain the feature map \mathcal{X}' . This feature map \mathcal{X}' is the input to the quantum variational circuit. In the circuit, the x values are used in the $R_Y(\cdot)$ gates, while the θ parameters are optimized during the training process in the $Rot(\cdot)$ operations. The green-shaded area represents the optimized layer in the Qsco module. The Qsco module output score is $\hat{\mathcal{S}}$.

of adiabatic quantum annealers (AQAs) for quantum speed-up in anomaly detection.

Proposed Approach

Problem Statement The studied problem, open-set supervised AD, can be formally stated as follows. Given a set of training samples $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^{N+M}$, in which $\mathcal{X}_n = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ is the normal sample set and $\mathcal{X}_a = \{\mathbf{x}_{N+1}, \mathbf{x}_{N+2}, \dots, \mathbf{x}_{N+M}\}$ ($M \ll N$) is a very small set of annotated anomalies that provide some knowledge about true anomalies, and the M anomalies belong to the seen anomaly classes $\zeta \subset \varpi$, where $\varpi = \{c_i\}_{i=1}^{|\varpi|}$ denotes the set of all possible anomaly classes, and then the goal is detect both seen and unseen anomaly classes by learning an anomaly scoring function $g : \mathcal{X} \rightarrow \mathbb{R}$ that assigns larger anomaly scores to both seen and unseen anomalies than normal samples.

Quantum Scoring Module In open-set anomaly detection, the model needs to identify new categories or anomalies not seen during the training phase. This type of anomaly detection is critical for many applications, such as security monitoring, medical diagnostics, and industrial quality control, where new and unknown anomalies often arise.

For an image in the dataset $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^{N+M}$, $\mathbf{x}_i \in \mathbb{R}^{H \times W \times C}$. Usually, this image can be represented as a three-dimensional tensor, which for color images usually contains

three color channels (red, green, blue). Let the height of the image be H , the width is W , and the color channel is C .

The feature extractor $\mathcal{F}(\cdot)$ is a function that converts the input image into a feature vector or matrix. It may contain multiple convolution layers, pooling layers, and activation layers and can be expressed as:

$$\mathcal{X}' = \mathcal{F}(\mathcal{X}) : \mathbb{R}^{H \times W \times C} \longrightarrow \mathbb{R}^{H' \times W' \times D}, \quad (1)$$

where D represents feature depth and ($D \gg C$).

The hyperbolic tangent function, or \tanh , is used to normalize input data to a fixed interval. For quantum computing's rotational gates (like R_Y), input typically needs to be in terms of angles, making \tanh ideal for scaling data to the $[-1, 1]$ range, which helps control the rotation angles of these gates. The next step is to apply the hyperbolic tangent activation function \tanh to the feature vector \mathcal{X}' . The hyperbolic tangent function performs a nonlinear transformation on each element \mathcal{X}' , and each element of the output is $[-1, 1]$, mathematically expressed as :

$$\hat{\mathcal{X}} = \tanh(\mathcal{X}') = [\tanh(\mathbf{x}'_1), \tanh(\mathbf{x}'_2), \dots, \tanh(\mathbf{x}'_{N+M})]. \quad (2)$$

We applied a set of K qubits, denoted as $\tau \in \mathbb{C}^K$, where each τ_m represents the state of the m th qubit, in the Qsco module. A quantum state could be represented as:

$$|\psi_m\rangle = \alpha_m |0\rangle + \beta_m |1\rangle, \quad (3)$$

where $|\psi_m\rangle$ is the state of qubit, α_m and β_m are the complex probability amplitudes, which satisfy:

$$|\alpha_m|^2 + |\beta_m|^2 = 1. \quad (4)$$

According to Eq. 4, Eq. 3 could be represented as:

$$|\psi_m\rangle = e^{i\gamma_m} \left(\cos \frac{\theta_m}{2} |0\rangle + e^{i\varphi_m} \sin \frac{\theta_m}{2} |1\rangle \right), \quad (5)$$

where θ_m , φ_m , and γ_m are real numbers specific to the m th qubit. The factor $e^{i\gamma_m}$ can be ignored due to its negligible effect (Nielsen and Chuang 2010a). Therefore, we can rewrite Eq. 5 as:

$$|\psi_m\rangle = \cos \frac{\theta_m}{2} |0\rangle + e^{i\varphi_m} \sin \frac{\theta_m}{2} |1\rangle. \quad (6)$$

The input of the quantum circuits is the $\hat{\mathcal{X}} = \{\hat{\mathbf{x}}_1, \hat{\mathbf{x}}_2, \dots, \hat{\mathbf{x}}_{N+M}\}$ as shown in Fig. 2. We take $\hat{\mathbf{x}}_m$ and utilize the Rotation Y gate (Ry gate) as below:

$$Ry(\hat{\mathbf{x}}_m) = \begin{pmatrix} \cos(\frac{\hat{\mathbf{x}}_m}{2}) & -\sin(\frac{\hat{\mathbf{x}}_m}{2}) \\ \sin(\frac{\hat{\mathbf{x}}_m}{2}) & \cos(\frac{\hat{\mathbf{x}}_m}{2}) \end{pmatrix}. \quad (7)$$

For a qubit in state in Eq. 3, we can calculate the result as below:

$$Ry(\hat{\mathbf{x}}_m)|\psi_m\rangle = \begin{pmatrix} \alpha_m \cos(\frac{\hat{\mathbf{x}}_m}{2}) - \beta_m \sin(\frac{\hat{\mathbf{x}}_m}{2}) \\ \alpha_m \sin(\frac{\hat{\mathbf{x}}_m}{2}) + \beta_m \cos(\frac{\hat{\mathbf{x}}_m}{2}) \end{pmatrix}. \quad (8)$$

Then we initialize the parameters for the rotation gates $Rot(\cdot)$, $\theta \in \mathbb{R}^{m \times n \times k}$ defined as below:

$$Rot(\cdot) = R_z(\lambda)R_y(\phi)R_x(\theta). \quad (9)$$

Rx gate (Rotation around the X-axis gate) in quantum computing is used to rotate qubits around the X-axis of the Bloch ball, and the Ry gate is shown in Eq. 7. The Rotation Z gate (Rz gate) is used in quantum computing to rotate the qubit around the Z-axis as shown in Fig. 3.

$$Rx(\theta) = \begin{pmatrix} \cos(\frac{\theta}{2}) & -i \sin(\frac{\theta}{2}) \\ -i \sin(\frac{\theta}{2}) & \cos(\frac{\theta}{2}) \end{pmatrix}, R_z(\lambda) = \begin{pmatrix} e^{-i\frac{\lambda}{2}} & 0 \\ 0 & e^{i\frac{\lambda}{2}} \end{pmatrix}. \quad (10)$$

So the quantum state after the $Rot(\cdot)$ could be expressed in summary as:

$$|\psi''_m\rangle = R_z(\lambda)R_y(\phi)R_x(\theta)|\psi'_m\rangle \quad (11)$$

The function of the CNOT gate is as follows: if the control bit is in the $|1\rangle$ state, it will perform a Pauli-X gate operation on the target bit (flip the state of the target bit). The target bit is unchanged if the control bit is in the $|0\rangle$ state.

We utilize the CNOT gate to implement conditional logic between qubits:

$$|\tilde{\psi}_m\rangle = |\psi''_m\rangle \otimes |\psi''_{m+1}\rangle, m \in [0, K-1]. \quad (12)$$

We implement a Pauli-Z gate (Z gate for short) on a superposition state that will affect the phase of the qubit without changing its probability amplitude as:

$$\sigma_z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \quad (13)$$




Figure 3: Visualization of the quantum state $|\rho\rangle$ on the Bloch sphere during the quantum downsampling process. The left, middle, and right figures depict the effects of the single-qubit rotation gates RX , RY , and RZ .

When we apply the Pauli-Z gate in Eq. 13 to a qubit in Eq. 3, we can get:

$$|\psi\rangle = \alpha |0\rangle - \beta |1\rangle. \quad (14)$$

Physically, the Z-gate acts like a 180 degree rotation of the qubit's phase. This phase rotation does not change the measured probability distribution of the qubit, but it affects the global phase of the quantum state:

$$\tilde{s}_m = \sigma_z \cdot |\tilde{\psi}_m\rangle. \quad (15)$$

The output from Qsco as the input from the feature extractor $\hat{\mathcal{X}}_n$ to the output \hat{S}_n is shown as below:

$$\hat{\mathcal{X}} \longrightarrow \hat{S}, S = \{s_1, s_2, \dots, s_{N+M}\} \quad (16)$$

\hat{S} is the output of the Qsco module, which can be added to the scoring part of any model.

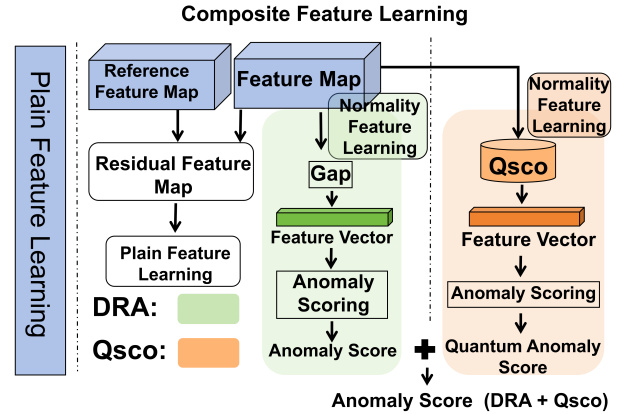


Figure 4: Overview of our proposed Qsco module and classical DRA (Ding, Pang, and Shen 2022) latent residual abnormality learning in a composite feature space.

We chose DRA to be our backbone model because it is the SOTA method in the OSAD task, and we insert Qsco to DRA as an additional part to participate in the scoring process as shown in Fig. 4.

Experimental Results

Hyper parameters We employ Adam as the optimizer, initialized with a learning rate of 0.0002 and a weight decay

Dataset (ϖ)	DRA	DRA + Qsco ($\ell = 2$)	DRA + Qsco ($\ell = 1$)	DRA + Qsco ($\ell = 3$)
Ten Training Anomaly Examples (Random) ♡				
AITEX (12)	0.886 \pm 0.021	0.875 \pm 0.029	0.832 \pm 0.021	0.880 \pm 0.020
SDD (1)	0.972 \pm 0.010	0.980 \pm 0.007	0.979 \pm 0.006	0.974 \pm 0.010
ELPV (2)	0.821 \pm 0.028	0.819 \pm 0.020	0.820 \pm 0.017	0.819 \pm 0.018
Optical (1)	0.967 \pm 0.007	0.969 \pm 0.008	0.950 \pm 0.014	0.960 \pm 0.009
BrainMRI (1)	0.959 \pm 0.011	0.972 \pm 0.017	0.978 \pm 0.017	0.961 \pm 0.019
HeadCT (1)	0.998 \pm 0.003	0.999 \pm 0.003	0.989 \pm 0.008	0.980 \pm 0.012
Hyper-Kvasir (4)	0.840 \pm 0.029	0.813 \pm 0.019	0.811 \pm 0.025	0.812 \pm 0.018
MVTec AD (-)	0.960 \pm 0.011	0.968 \pm 0.010	0.966 \pm 0.013	0.961 \pm 0.024
One Training Anomaly Example (Random) ♡				
AITEX (12)	0.703 \pm 0.038	0.708 \pm 0.027	0.685 \pm 0.044	0.728 \pm 0.018
SDD (1)	0.853 \pm 0.056	0.876 \pm 0.032	0.803 \pm 0.064	0.866 \pm 0.067
ELPV (2)	0.626 \pm 0.074	0.749 \pm 0.048	0.588 \pm 0.022	0.595 \pm 0.038
Optical (1)	0.894 \pm 0.026	0.910 \pm 0.011	0.859 \pm 0.020	0.879 \pm 0.017
BrainMRI (1)	0.638 \pm 0.045	0.692 \pm 0.051	0.718 \pm 0.043	0.708 \pm 0.039
HeadCT (1)	0.804 \pm 0.010	0.781 \pm 0.007	0.818 \pm 0.006	0.846 \pm 0.003
Hyper-Kvasir (4)	0.712 \pm 0.010	0.768 \pm 0.015	0.771 \pm 0.008	0.770 \pm 0.014
MVTec AD (-)	0.904 \pm 0.033	0.914 \pm 0.029	0.905 \pm 0.031	0.908 \pm 0.035

Table 2: AUC results (mean \pm std) on 8 real-world AD datasets under **general** setting (noted by ♡). The best and second-best results and the third-best are respectively highlighted in **red** and **blue** and **bold**. ϖ is the number of anomaly classes.

factor of $1e - 5$. The StepLR scheduler is utilized with the learning rate decaying by 0.1 every 10 epoch. The batch size is 48, and steps for each training epoch are 20, and there are 30 epochs for training. Many studies use synthetic anomaly detection datasets derived from well-known image classification benchmarks, like MNIST (LeCun et al. 2010) and CIFAR-10 (Krizhevsky 2009), utilizing either one-vs-all or one-vs-one approaches. These methods create anomalies that are significantly different from standard samples. However, the differences between anomalies and standard samples are often subtle and minor in real-world scenarios, such as detecting defects in industrial settings or identifying lesions in medical imaging. Consequently, following previous research (Ding, Pang, and Shen 2022), we prioritize natural anomaly datasets, including five industrial defect inspection datasets: MVTEC AD (Bergmann et al. 2019), AITEX (Silvestre-Blanes et al. 2019), SDD (Tabernik et al. 2019), ELPV (Deitsch et al. 2019) and Optical (Wieler, Hahn, and Hamprecht 2023), and three more real-world datasets: BrainMRI (Salehi et al. 2021) and HeadCT (Salehi et al. 2021), Hyperkvasir (Borgli et al. 2020). Experiments are performed on an NVIDIA 3060 GPU, 48 GB RAM, and an NVIDIA RTX A6000, 48 GB RAM. We built the Qsco quantum circuit on PennyLane (Bergholm et al. 2022) for quantum simulation.

One and Ten Anomaly Examples It means the number

of anomaly data in the training set is one and ten, respectively. We use the following two experiment protocols:

General setting It mimics a typical open-set anomaly detection situation, where a few anomaly examples are randomly selected from all potential anomaly classes in the test set for each dataset. These selected anomalies are then excluded from the test data. This approach aims to reflect real-world conditions where it is uncertain which anomaly classes are known and the extent of the anomaly classes covered by the given examples. Consequently, the datasets may include both seen and unseen anomaly classes or solely the seen anomaly classes, depending on the complexity of the applications, such as the total number of possible anomaly classes.

Hard setting It aims to specifically assess the models' ability to detect unseen anomaly classes, which is a critical challenge in open-set AD. To achieve this, anomaly example sampling is restricted to a single anomaly class, and all samples from this class are removed from the test set to ensure it includes only unseen anomaly classes. It is important to note that this setting applies only to datasets with at least two anomaly classes.

In Table 2, we conducted an empirical study on how changing the depth of the Qsco model (the number of layers ℓ) affects its performance across eight different datasets. We categorized the tests into scenarios involving ten training

anomalies and one training anomaly. Our results imply that increasing the parameter ℓ , which controls the depth of layers in Qsco, leads to better anomaly detection AUC scores, especially in scenarios with only one training anomaly. As ℓ increases, Qsco’s variational circuit becomes deeper, enhancing its ability to detect anomalies in more challenging tasks (performance on one training anomaly is better on ten training anomalies). However, increasing the depth can reduce performance in simpler scenarios with ten anomalies in the training set, possibly due to over-fitting. This highlights the importance of selecting an optimal value for ℓ , as the model’s complexity needs to be adjusted based on the specific characteristics of the task. For MVTec AD, more class details (noted as (-)) can be found in **Appendix** Table. 1 at <https://arxiv.org/abs/2405.16368>.

Meanwhile, we measured the simulation time of quantum variational circuits using PennyLane 0.35.1, and the computing resource used was an NVIDIA GeForce RTX 3060 and 64GB RAM. While Qsco performance has been improved as indicated in Table 2, the quantum approach also has a relative shortcoming: it will cause a certain amount of time loss. However, from the table. 3, we observe that the additional time consumption of approximately 20% (6 seconds) is acceptable compared to the performance enhancement, attributed to the fewer quantum bits $\tau = 9$. The reason behind the time loss is the deepening of the quantum variational circuit, which ultimately leads to an increase in the revolving gate operations $Rot(\cdot)$. Nevertheless, a too-deep circuit (bigger ℓ) will lead to over-fitting. On the contrary, a too-shallow circuit will be unable to explore the abnormal data space, leading to under-fitting problems. Therefore, we need to consider fewer quantum gate operations while ensuring performance because, in the NISQ era (Preskill 2018), a controllable number of quantum gates has higher implement ability and operability.

In Table. 4, after applying the best-performing parameter $\ell = 2$, we compared the performance (AUC score) with more baseline methods, namely SAOE (combining data augmentation-based Synthetic Anomalies (Li et al. 2021a; Liznerski et al. 2020; Tack et al. 2020), FLOS (Lin et al. 2017), MLEP (Liu et al. 2019), and DevNet (Pang et al. 2021; Pang, Shen, and Van Den Hengel 2019)). This table indicates that Qsco has better anomaly detection performance than most baseline methods among all datasets.

Table. 5 presents performance quantification following the established general configuration under ten and one anomaly examples. This table computes the average values for each dataset, with additional details available (subset for each dataset) in the **Appendix** Table. 2 for ten anomaly examples and **Appendix** Table. 3 for one anomaly example. Table. 5 demonstrates that for the dataset under "One Training Anomaly Examples (Random)," the introduction of Qsco significantly improves the performance metrics of the DRA. However, for "Ten Training Anomaly Example (Random)," it is important to recognize that integrating DRA and Qsco does not consistently achieve superior outcomes across the datasets. The proposed Qsco improves the performance of DRA under hard settings for ten training anomaly examples. However, DRA + Qsco doesn’t always take the

best performance. This discrepancy arises because the foundational DRA often demonstrates diminished effectiveness compared to alternative methods. Therefore, the combined performance of DRA and Qsco falls short. This shortfall is primarily attributable to the inherent limitations of DRA, despite our efforts to enhance its performance.

We present a comprehensive visualization of the performance metrics of Qsco ($\ell = 2$) across eight datasets, contrasting the results for scenarios involving one anomaly and ten anomaly examples in Fig. 5. Details for more ℓ could be found in **Appendix** Fig.1 and Fig.2. The results reveal that Qsco significantly enhances the effectiveness of the DRA across a broad spectrum of evaluation metrics in situations under one anomaly example. This enhancement is consistent and marked by a notably more substantial improvement than scenarios with ten training anomaly examples. Furthermore, despite the inherently lower task complexity associated with ten training anomaly examples, Qsco demonstrates distinct advantages. The performance improvements, although less dramatic than in single anomaly scenarios, suggest that Qsco’s methodology adapts effectively to varying degrees of complexity within anomaly detection tasks. This adaptability indicates robustness, critical for practical applications where anomaly conditions can vary widely.

Noise Analysis

In the NISQ era, anticipating the deployment of Qsco on actual quantum computing platforms necessitates addressing several critical issues, among which quantum noise is paramount. Accordingly, this section is dedicated to the application of various noise models to Qsco, enhancing our understanding of quantum noise in practical scenarios.

Depolarizing noise. The Depolarizing channel is denoted as Δ_γ with probability γ , replaces the state ρ (density matrix) with the maximally mixed state $\frac{I}{d}$, where I is the identity matrix and d is the dimension of the Hilbert space.

The action of the depolarizing channel on a state ρ is given by:

$$\Delta_\gamma(\rho) = (1 - \gamma)\rho + \frac{\gamma}{d}I. \quad (17)$$

Bit Flip Noise. The bit flip noise flips the state from $|0\rangle$ to $|1\rangle$ or vice versa with probability γ as

$$\mathcal{E}(\rho) = (1 - \gamma)\rho + \gamma X\rho X, \quad (18)$$

where X is the Pauli-X gate.

Phase Flip Noise. Similar to the bit flipping noise. The probability γ , this noise applies a Pauli-Z gate to the quantum state, adding a negative sign to the $|1\rangle$ as:

$$\mathcal{E}(\rho) = (1 - \gamma)\rho + \gamma Z\rho Z. \quad (19)$$

Amplitude Damping. The amplitude damping channel simulates the effect of energy dissipation from the qubit to the environment, decaying from the excited state $|1\rangle$ to the ground state $|0\rangle$. The effect can be represented using Kraus operators:

$$\mathcal{E}(\rho) = E_0\rho E_0^\dagger + E_1\rho E_1^\dagger, \quad (20)$$

where $E_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{pmatrix}$ and $E_1 = \begin{pmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{pmatrix}$ are the Kraus operators (Nielsen and Chuang 2010b).

Dataset (ϖ)	DevNet	DRA	DRA + Qsco ($\ell = 2$)	DRA + Qsco ($\ell = 1$)	DRA + Qsco ($\ell = 3$)
Ten Training Anomaly Examples (Random) ♡					
AITEX (12)	20.93 \pm 0.76	36.75 \pm 0.50	42.22 \pm 2.01 (14.9% \uparrow)	41.38 \pm 1.02 (12.6% \uparrow)	43.53 \pm 2.66 (18.4% \uparrow)
SDD (1)	21.28 \pm 0.98	36.59 \pm 3.51	42.28 \pm 2.37 (15.6% \uparrow)	40.00 \pm 2.72 (9.3% \uparrow)	45.43 \pm 3.15 (24.2% \uparrow)
ELPV (2)	21.41 \pm 0.95	34.30 \pm 2.30	41.79 \pm 2.29 (21.8% \uparrow)	40.80 \pm 3.12 (19.0% \uparrow)	44.31 \pm 2.93 (29.2% \uparrow)
Optical (1)	21.43 \pm 1.04	35.40 \pm 1.55	42.81 \pm 3.06 (20.9% \uparrow)	40.44 \pm 2.26 (14.2% \uparrow)	43.97 \pm 3.43 (24.2% \uparrow)
BrainMRI (1)	13.81 \pm 0.72	19.12 \pm 0.17	24.25 \pm 0.44 (26.8% \uparrow)	22.85 \pm 0.50 (19.5% \uparrow)	26.45 \pm 0.52 (38.3% \uparrow)
HeadCT (1)	14.03 \pm 0.55	19.18 \pm 0.25	24.64 \pm 0.37 (28.5% \uparrow)	22.67 \pm 0.27 (18.2% \uparrow)	26.26 \pm 0.13 (37.0% \uparrow)
Hyper-Kvasir (4)	21.47 \pm 1.09	34.06 \pm 2.77	41.70 \pm 3.21 (22.4% \uparrow)	41.37 \pm 5.55 (21.5% \uparrow)	53.65 \pm 3.43 (57.5% \uparrow)
One Training Anomaly Example (Random) ♡					
AITEX (12)	20.89 \pm 1.07	32.66 \pm 2.09	44.12 \pm 1.86 (35.1% \uparrow)	43.68 \pm 4.56 (33.7% \uparrow)	44.74 \pm 1.5 (37.0% \uparrow)
SDD (1)	21.90 \pm 1.31	35.14 \pm 2.09	40.54 \pm 3.20 (15.4% \uparrow)	38.64 \pm 3.03 (10.0% \uparrow)	42.06 \pm 3.40 (19.7% \uparrow)
ELPV (2)	21.39 \pm 1.80	33.01 \pm 2.15	46.60 \pm 4.03 (41.2% \uparrow)	43.50 \pm 2.85 (31.8% \uparrow)	47.33 \pm 5.22 (43.4% \uparrow)
Optical (1)	20.94 \pm 1.30	34.01 \pm 1.70	42.92 \pm 3.50 (26.2% \uparrow)	40.84 \pm 2.12 (20.1% \uparrow)	44.18 \pm 0.35 (30.0% \uparrow)
BrainMRI (1)	14.19 \pm 0.87	19.33 \pm 0.82	24.46 \pm 1.97 (26.5% \uparrow)	23.95 \pm 2.11 (24.0% \uparrow)	30.82 \pm 4.07 (59.4% \uparrow)
HeadCT (1)	14.73 \pm 0.71	19.74 \pm 0.79	30.20 \pm 1.67 (53.0% \uparrow)	22.73 \pm 1.65 (15.1% \uparrow)	31.09 \pm 5.10 (57.5% \uparrow)
Hyper-Kvasir (4)	23.08 \pm 1.22	34.55 \pm 2.08	41.75 \pm 3.29 (20.8% \uparrow)	41.15 \pm 4.28 (19.1% \uparrow)	47.70 \pm 2.67 (38.1% \uparrow)

Table 3: **Time complexity (s) per epoch.** ϖ is the number of anomaly classes. The testing computing resource is an NVIDIA GeForce RTX 3060 and 64GB RAM. Simulation is processed by PennyLane 0.35.1. ♡ means general settings. **Red** means the biggest time consumption.

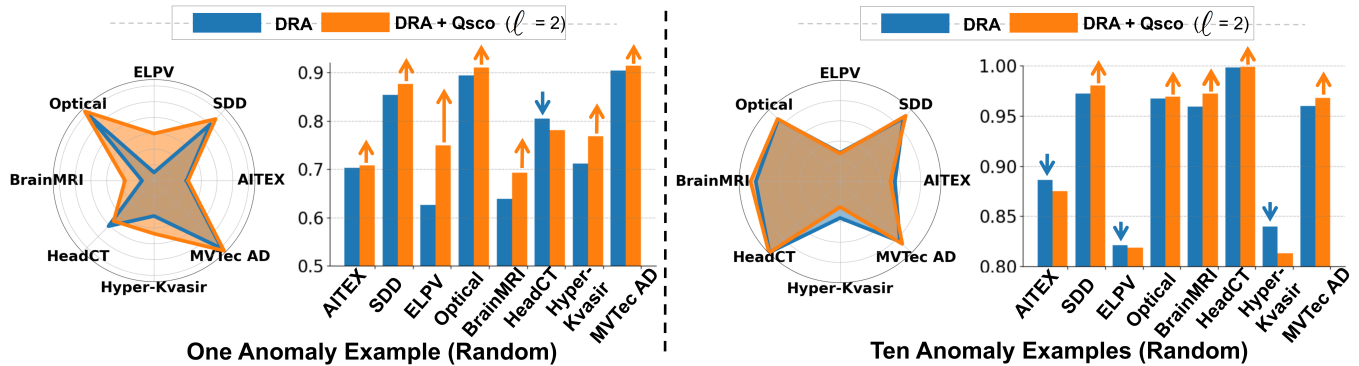


Figure 5: **Performance Comparison of DRA and Qsco ($\ell = 2$) with different numbers of anomaly examples.** This comparison evaluates the performance (AUC score) of DRA and DRA + Qsco $\ell = 2$ under the same **general** settings across eight different datasets. The left side displays results for a single anomaly example, while the right shows results for ten. In the column chart, the \uparrow indicates an improvement from the DRA model, and the \downarrow signifies a decrease from the DRA.

Phase Damping. Unlike amplitude damping, phase damping preserves the amplitude of the quantum state but loses phase information, describing a dephasing process of quantum information without energy loss as:

$$\mathcal{E}(\rho) = E_0 \rho E_0^\dagger + E_1 \rho E_1^\dagger. \quad (21)$$

As illustrated in Table. 6, it is evident that various noise models adversely affect the AITEX, SDD, BrainMRI, HeadCT, and Hyper-Kvasir datasets in the context of Ten

Training Anomaly Examples, with the overall maximum fluctuation around 10%. Conversely, performance on the ELPV and Optical datasets has improved. We hypothesize that introducing noise acts as a form of regularization for these two datasets, potentially preventing overfitting. This may allow the models to improve generalization and robustness in noisy real-world environments. For One and Ten Training Anomaly Examples in the Mvtectad dataset could

Dataset (ϖ)	SAOE	FLOS	MLEP	DevNet	DRA	DRA + Qsco ($\ell = 2$)
Ten Training Anomaly Examples (Random) ♡						
AITEX (12)	0.874 \pm 0.024	0.841 \pm 0.049	0.867 \pm 0.037	0.881\pm0.039	0.886\pm0.021	0.875\pm0.029
SDD (1)	0.955 \pm 0.020	0.967 \pm 0.018	0.783 \pm 0.013	0.977\pm0.010	0.972\pm0.010	0.980\pm0.001
ELPV (2)	0.793 \pm 0.047	0.818 \pm 0.032	0.794 \pm 0.047	0.856\pm0.013	0.821\pm0.028	0.819\pm0.020
Optical (1)	0.941\pm0.013	0.720 \pm 0.055	0.740 \pm 0.039	0.777 \pm 0.022	0.967\pm0.007	0.969\pm0.008
BrainMRI (1)	0.900 \pm 0.041	0.955 \pm 0.011	0.959\pm0.011	0.959 \pm 0.012	0.959\pm0.011	0.972\pm0.017
HeadCT (1)	0.935 \pm 0.021	0.971 \pm 0.004	0.972 \pm 0.014	0.993\pm0.004	0.998\pm0.003	0.999\pm0.003
Hyper-Kvasir (4)	0.666 \pm 0.050	0.773 \pm 0.029	0.600 \pm 0.069	0.876\pm0.012	0.840\pm0.029	0.813\pm0.019
MVTec AD (-)	0.926 \pm 0.010	0.939\pm0.007	0.907 \pm 0.005	0.934 \pm 0.019	0.960\pm0.011	0.968\pm0.010
One Training Anomaly Example (Random) ♡						
AITEX (12)	0.675 \pm 0.094	0.538 \pm 0.073	0.564 \pm 0.055	0.627\pm0.031	0.703\pm0.038	0.708\pm0.027
SDD (1)	0.781 \pm 0.009	0.840 \pm 0.043	0.811 \pm 0.045	0.873\pm0.055	0.853\pm0.056	0.876\pm0.032
ELPV (2)	0.635\pm0.092	0.457 \pm 0.056	0.578 \pm 0.062	0.691\pm0.069	0.626 \pm 0.074	0.749\pm0.048
Optical (1)	0.815\pm0.014	0.518 \pm 0.003	0.516 \pm 0.009	0.537 \pm 0.016	0.894\pm0.026	0.910\pm0.011
BrainMRI (1)	0.531 \pm 0.060	0.693 \pm 0.036	0.632 \pm 0.017	0.765\pm0.030	0.638\pm0.045	0.718\pm0.043
HeadCT (1)	0.597 \pm 0.022	0.698 \pm 0.092	0.758 \pm 0.038	0.768\pm0.031	0.804\pm0.010	0.781\pm0.007
Hyper-Kvasir (4)	0.498 \pm 0.100	0.668 \pm 0.004	0.445 \pm 0.040	0.682\pm0.038	0.712\pm0.010	0.768\pm0.015
MVTec AD (-)	0.834\pm0.007	0.792 \pm 0.014	0.744 \pm 0.019	0.812 \pm 0.049	0.904\pm0.033	0.914\pm0.029

Table 4: AUC results (mean \pm std) on 8 real-world AD datasets under **general** setting (noted by ♡). The best and second-best results and the third-best are respectively highlighted in **red** and **blue** and **bold**. ϖ is the number of anomaly classes. SAOE, FLOS, and MLEP results are from Ding, Pang, and Shen (2022).

Dataset	SAOE	FLOS	MLEP	DevNet	DRA	DRA + Qsco ($\ell = 2$)
Ten Training Anomaly Examples (Random) ◇						
Carpet ∞	0.762 \pm 0.073	0.761 \pm 0.012	0.751 \pm 0.023	0.847\pm0.017	0.920\pm0.039	0.927\pm0.044
Metal-nut ∞	0.855 \pm 0.016	0.922 \pm 0.014	0.878 \pm 0.058	0.965\pm0.011	0.949\pm0.021	0.956\pm0.020
AITEX ∞	0.724\pm0.032	0.635 \pm 0.043	0.626 \pm 0.041	0.683 \pm 0.032	0.700\pm0.049	0.709\pm0.038
ELPV ∞	0.683 \pm 0.047	0.646 \pm 0.032	0.745\pm0.020	0.702 \pm 0.023	0.705\pm0.050	0.731\pm0.042
Hyper-Kvasir ∞	0.698 \pm 0.021	0.786\pm0.021	0.571 \pm 0.014	0.822\pm0.019	0.764\pm0.047	0.750 \pm 0.031
One Training Anomaly Example (Random) ◇						
Carpet ∞	0.753 \pm 0.055	0.678 \pm 0.040	0.679 \pm 0.029	0.767\pm0.018	0.878\pm0.066	0.881\pm0.072
Metal-nut ∞	0.816 \pm 0.029	0.855 \pm 0.024	0.825 \pm 0.023	0.855\pm0.016	0.932\pm0.028	0.928\pm0.037
AITEX ∞	0.674\pm0.034	0.624 \pm 0.024	0.466 \pm 0.030	0.646 \pm 0.034	0.678\pm0.071	0.676\pm0.040
ELPV ∞	0.614 \pm 0.048	0.691\pm0.008	0.566 \pm 0.111	0.648\pm0.057	0.627\pm0.056	0.616 \pm 0.032
Hyper-Kvasir ∞	0.406 \pm 0.018	0.571 \pm 0.004	0.480 \pm 0.044	0.595\pm0.023	0.677\pm0.065	0.687\pm0.037

Table 5: AUC results (mean \pm std) on 8 real-world AD datasets under the **hard** setting (noted by ◇). The best and second-best results and the third-best are respectively highlighted in **red** and **blue** and **bold**. ∞ is the average of anomaly classes. SAOE, FLOS, MLEP, and DevNet results are from Ding, Pang, and Shen (2022).

Dataset (ϖ)	DRA + Qsco ($\ell = 2$)	DRA + Qsco ($\ell = 2$) ♣	DRA + Qsco ($\ell = 2$) ♠	DRA + Qsco ($\ell = 2$) *	DRA + Qsco ($\ell = 2$) \emptyset
Ten Training Anomaly Examples (Random) ♡					
AITEX (12)	0.875 \pm 0.029	0.869 \pm 0.040 (0.69%) \downarrow	0.778 \pm 0.032 (11.09%) \downarrow	0.801 \pm 0.0334 (8.46%) \downarrow	0.794 \pm 0.020 (9.26%) \downarrow
SDD (1)	0.980 \pm 0.007	0.966 \pm 0.020 (1.43%) \downarrow	0.971 \pm 0.006 (0.92%) \downarrow	0.968 \pm 0.010 (1.22%) \downarrow	0.972 \pm 0.012 (0.82%) \downarrow b
ELPV (2)	0.819 \pm 0.020	0.832 \pm 0.021 (1.59%) \uparrow b	0.823 \pm 0.018 (0.49%) \uparrow b	0.832 \pm 0.015 (1.59%) \uparrow b	0.840 \pm 0.020 (2.56%) \uparrow b
Optical (1)	0.969 \pm 0.008	0.971 \pm 0.012 (0.21%) \uparrow b	0.970 \pm 0.009 (0.10%) \uparrow b	0.973 \pm 0.017 (0.41%) \uparrow b	0.983 \pm 0.005 (1.44%) \uparrow b
BrainMRI (1)	0.972 \pm 0.017	0.878 \pm 0.018 (9.67%) \downarrow	0.893 \pm 0.019 (8.13%) \downarrow	0.866 \pm 0.020 (10.90%) \downarrow	0.871 \pm 0.022 (10.39%) \downarrow
HeadCT (1)	0.999 \pm 0.003	0.998 \pm 0.002 (0.10%) \downarrow b	0.990 \pm 0.007 (0.90%) \downarrow	0.982 \pm 0.010 (1.70%) \downarrow	0.982 \pm 0.010 (1.70%) \downarrow
Hyper-Kvasir (4)	0.813 \pm 0.019	0.852 \pm 0.023 (4.80%) \uparrow b	0.789 \pm 0.041 (2.95%) \downarrow	0.740 \pm 0.039 (9.00%) \downarrow	0.787 \pm 0.038 (3.20%) \downarrow
One Training Anomaly Example (Random) ♡					
AITEX (12)	0.708 \pm 0.027	0.734 \pm 0.023 (3.67%) \uparrow b	0.700 \pm 0.038 (1.13%) \downarrow	0.722 \pm 0.019 (2.00%) \uparrow b	0.708 \pm 0.034 \longleftrightarrow b
SDD (1)	0.876 \pm 0.032	0.848 \pm 0.045 (3.20%) \downarrow	0.908 \pm 0.026 (3.65%) \uparrow b	0.839 \pm 0.045 (4.22%) \downarrow	0.904 \pm 0.036 (3.20%) \uparrow b
ELPV (2)	0.749 \pm 0.048	0.621 \pm 0.046 (17.10%) \downarrow b	0.749 \pm 0.020 \longleftrightarrow b	0.770 \pm 0.035 (2.80%) \uparrow b	0.760 \pm 0.031 (1.47%) \uparrow b
Optical (1)	0.910 \pm 0.011	0.893 \pm 0.024 (1.87%) \downarrow	0.904 \pm 0.015 (0.66%) \downarrow b	0.890 \pm 0.017 (2.20%) \downarrow	0.909 \pm 0.017 (0.11%) \downarrow b
BrainMRI (1)	0.718 \pm 0.043	0.739 \pm 0.046 (2.92%) \uparrow b	0.742 \pm 0.040 (3.34%) \uparrow b	0.729 \pm 0.044 (1.53%) \uparrow b	0.703 \pm 0.048 (2.09%) \downarrow b
HeadCT (1)	0.781 \pm 0.007	0.871 \pm 0.061 (11.52%) \uparrow b	0.847 \pm 0.048 (8.45%) \uparrow b	0.710 \pm 0.100 (9.10%) \downarrow	0.716 \pm 0.085 (8.32%) \downarrow
Hyper-Kvasir (4)	0.768 \pm 0.015	0.804 \pm 0.057 (4.69%) \uparrow b	0.771 \pm 0.037 (0.39%) \uparrow b	0.804 \pm 0.035 (4.69%) \uparrow b	0.788 \pm 0.066 (2.60%) \uparrow b

Table 6: AUC results (mean \pm std) on 7 real-world AD datasets under **general** setting (noted by ♡). The increasing and the decreasing (absolute values) is compared with $\ell = 2$. ϖ is the number of anomaly classes. If the performance under the noise model is still better than the baseline (DRA), it is noted as \uparrow b. The amplitude damping noise model is ♣, the phase flipping noise is ♠, the depolarizing noise is *, and the bit flipping noise is \emptyset .

be found in **Appendix**.

In the One Training Anomaly Example scenario, we observe an overall improvement, which is likely due to the reduced number of abnormal samples in the training data, which enables the model to focus more on capturing and learning the characteristics of normal data rather than being influenced by the noise associated with abnormal samples. Consequently, when noise is introduced, the model demonstrates a robust capacity to withstand minor data disturbances. This fortifies its adaptability and resilience against unseen noise, enhancing performance outcomes.

Qubits	9	12
PennyLane	13.11 \pm 0.25	17.26 \pm 1.37
Qubits	15	18
PennyLane	58.55 \pm 1.24	182.19 \pm 2.35

Table 7: Time complexity (in seconds) in the scalability analysis of Qsco. Average of 100 tests was performed on an NVIDIA GeForce RTX 3060 with 48 GB of RAM.

Limitations and Future Work The data presented in Table. 7 shows that classically simulating multiple instances of Qsco results in an exponential increase in time complexity. This exponential growth significantly constrains the scalability of classical simulations of Qsco. However, with the advent of frameworks such as NVIDIA’s CUDA Quantum (Kim et al. 2023), which enables the simulation of quantum

algorithms on CUDA-enabled GPUs, there is potential to enhance the scalability and efficiency of our experiments by leveraging GPU-accelerated quantum simulations in the future. Despite these limitations, current findings demonstrate that a single Qsco surpasses the established baseline in the OSAD task. Meanwhile, we will continue to compare with the most SOTA OSAD methods (Zhu et al. 2024), striving to continuously improve Qsco and explore the potential of quantum computing in the OSAD domain.

Conclusion

This paper proposed a quantum scoring module Qsco for the OSAD task. This module can be embedded in any advanced model and learn the high-dimensional distribution of abnormal data through variational quantum circuits and entanglement between qubits. The proposed approach solves the problem of abnormal data not seen during the training process, which causes model performance degradation. Experimental results indicate that implementing Qsco markedly enhances the predicted AUC score and does not result in significant time loss. In this paper, we explored the application of quantum computing in anomaly detection by implementing Qsco. Drawing on results from a simulated quantum computing environment with noise models, we demonstrate a proof of concept highlighting the potential of integrating quantum computing and deep learning into open-set anomaly detection.

References

- Abbaszade, M.; Salari, V.; Mousavi, S. S.; Zomorodi, M.; and Zhou, X. 2021. Application of quantum natural language processing for language translation. *IEEE Access*, 9: 130434–130448.
- Acsintoae, A.; Florescu, A.; Georgescu, M.-I.; Mare, T.; Sumedrea, P.; Ionescu, R. T.; Khan, F. S.; and Shah, M. 2022. Ubnormal: New benchmark for supervised open-set video anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 20143–20153.
- Akcay, S.; Atapour-Abarghouei, A.; and Breckon, T. P. 2019. Ganomaly: Semi-supervised anomaly detection via adversarial training. In *Computer Vision—ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part III 14*, 622–637. Springer.
- Alchieri, L.; Badalotti, D.; Bonardi, P.; and Bianco, S. 2021. An introduction to quantum machine learning: from quantum logic to quantum deep learning. *Quantum Machine Intelligence*, 3(2): 28.
- Bergholm, V.; Izaac, J.; Schuld, M.; Gogolin, C.; Ahmed, S.; Ajith, V.; Alam, M. S.; Alonso-Linaje, G.; AkashNarayanan, B.; Asadi, A.; Arrazola, J. M.; Azad, U.; Banning, S.; Blank, C.; Bromley, T. R.; Cordier, B. A.; Ceroni, J.; Delgado, A.; Matteo, O. D.; Dusko, A.; Garg, T.; Guala, D.; Hayes, A.; Hill, R.; Ijaz, A.; Isacsson, T.; Ittah, D.; Jahangiri, S.; Jain, P.; Jiang, E.; Khandelwal, A.; Kottmann, K.; Lang, R. A.; Lee, C.; Loke, T.; Lowe, A.; McKiernan, K.; Meyer, J. J.; Montañez-Barrera, J. A.; Moyard, R.; Niu, Z.; O’Riordan, L. J.; Oud, S.; Panigrahi, A.; Park, C.-Y.; Polatajko, D.; Quesada, N.; Roberts, C.; Sá, N.; Schoch, I.; Shi, B.; Shu, S.; Sim, S.; Singh, A.; Strandberg, I.; Soni, J.; Száva, A.; Thabet, S.; Vargas-Hernández, R. A.; Vincent, T.; Vitucci, N.; Weber, M.; Wierichs, D.; Wiersema, R.; Willmann, M.; Wong, V.; Zhang, S.; and Killoran, N. 2022. PennyLane: Automatic differentiation of hybrid quantum-classical computations. arXiv:1811.04968.
- Bergmann, P.; Fauser, M.; Sattlegger, D.; and Steger, C. 2019. MVTEC AD — A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9584–9592.
- Borgli, H.; Thambawita, V.; Smedsrud, P. H.; Hicks, S.; Jha, D.; Eskeland, S. L.; Randel, K. R.; Pogorelov, K.; Lux, M.; Nguyen, D. T. D.; et al. 2020. HyperKvasir, a comprehensive multi-class image and video dataset for gastrointestinal endoscopy. *Scientific data*, 7(1): 283.
- Bozorgtabar, B.; and Mahapatra, D. 2023. Attention-conditioned augmentations for self-supervised anomaly detection and localization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 14720–14728.
- Caro, M. C.; Huang, H.-Y.; Cerezo, M.; Sharma, K.; Sornborger, A.; Cincio, L.; and Coles, P. J. 2022. Generalization in quantum machine learning from few training data. *Nature communications*, 13(1): 4919.
- Chen, Y.; Liu, Z.; Zhang, B.; Fok, W.; Qi, X.; and Wu, Y.-C. 2023. Mgn: Magnitude-contrastive glance-and-focus network for weakly-supervised video anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 387–395.
- Cho, M.; Kim, M.; Hwang, S.; Park, C.; Lee, K.; and Lee, S. 2023. Look around for anomalies: weakly-supervised anomaly detection via context-motion relational learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12137–12146.
- Deutsch, S.; Christlein, V.; Berger, S.; Buerhop-Lutz, C.; Maier, A.; Gallwitz, F.; and Riess, C. 2019. Automatic classification of defective photovoltaic module cells in electroluminescence images. *Solar Energy*, 185: 455–468.
- Ding, C.; Pang, G.; and Shen, C. 2022. Catching Both Gray and Black Swans: Open-set Supervised Anomaly Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Farhi, E.; and Neven, H. 2018. Classification with Quantum Neural Networks on Near Term Processors. arXiv: *Quantum Physics*.
- Georgescu, M.-I.; Barbalau, A.; Ionescu, R. T.; Khan, F. S.; Popescu, M.; and Shah, M. 2021. Anomaly detection in video via self-supervised and multi-task learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12742–12752.
- Görnitz, N.; Kloft, M.; Rieck, K.; and Brefeld, U. 2013. Toward supervised anomaly detection. *Journal of Artificial Intelligence Research*, 46: 235–262.
- Guo, M.; Liu, H.; Li, Y.; Li, W.; Gao, F.; Qin, S.; and Wen, Q. 2022. Quantum algorithms for anomaly detection using amplitude estimation. *Physica A: Statistical Mechanics and its Applications*, 604: 127936.
- Guo, M.; Pan, S.; Li, W.; Gao, F.; Qin, S.; Yu, X.; Zhang, X.; and Wen, Q. 2023. Quantum algorithm for unsupervised anomaly detection. *Physica A: Statistical Mechanics and its Applications*, 625: 129018.
- Han, S.; Hu, X.; Huang, H.; Jiang, M.; and Zhao, Y. 2022. Adbench: Anomaly detection benchmark. *Advances in Neural Information Processing Systems*, 35: 32142–32159.
- Kim, J.-S.; McCaskey, A.; Heim, B.; Modani, M.; Stanwyck, S.; and Costa, T. 2023. CUDA Quantum: The Platform for Integrated Quantum-Classical Computing. In *2023 60th ACM/IEEE Design Automation Conference (DAC)*, 1–4.
- Kong, S.; and Ramanan, D. 2021. Opengan: Open-set recognition via open data generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 813–822.
- Kottmann, K.; Metz, F.; Fraxanet, J.; and Baldelli, N. 2021. Variational quantum anomaly detection: Unsupervised mapping of phase diagrams on a physical quantum computer. *Physical Review Research*, 3(4): 043184.
- Krizhevsky, A. 2009. Learning Multiple Layers of Features from Tiny Images.
- LeCun, Y.; Cortes, C.; Burges, C.; et al. 2010. MNIST handwritten digit database.

- Li, C.-L.; Sohn, K.; Yoon, J.; and Pfister, T. 2021a. Cut-paste: Self-supervised learning for anomaly detection and localization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9664–9674.
- Li, C.-L.; Sohn, K.; Yoon, J.; and Pfister, T. 2021b. Cut-paste: Self-supervised learning for anomaly detection and localization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9664–9674.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Liu, N.; and Rebentrost, P. 2018. Quantum machine learning for quantum anomaly detection. *Physical Review A*, 97(4): 042315.
- Liu, W.; Luo, W.; Li, Z.; Zhao, P.; Gao, S.; et al. 2019. Margin Learning Embedded Prediction for Video Anomaly Detection with A Few Anomalies. In *IJCAI*, volume 3, 023–3.
- Liznerski, P.; Ruff, L.; Vandermeulen, R. A.; Franks, B. J.; Kloft, M.; and Müller, K.-R. 2020. Explainable deep one-class classification. *arXiv preprint arXiv:2007.01760*.
- Moro, L.; and Prati, E. 2023. Anomaly detection speed-up by quantum restricted Boltzmann machines. *Communications Physics*, 6(1): 269.
- Nielsen, M. A.; and Chuang, I. L. 2010a. *Quantum computation and quantum information*. Cambridge university press.
- Nielsen, M. A.; and Chuang, I. L. 2010b. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press.
- Pan, H.; Zhu, X.; Atici, S. F.; and Cetin, A. 2023. A hybrid quantum-classical approach based on the hadamard transform for the convolutional layer. In *International Conference on Machine Learning*, 26891–26903. PMLR.
- Pang, G.; Ding, C.; Shen, C.; and Hengel, A. v. d. 2021. Explainable Deep Few-shot Anomaly Detection with Deviation Networks. *arXiv preprint arXiv:2108.00462*.
- Pang, G.; Shen, C.; and Van Den Hengel, A. 2019. Deep anomaly detection with deviation networks. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 353–362.
- Peng, Y.; Li, X.; Liang, Z.; and Wang, Y. 2024. HyQ2: A Hybrid Quantum Neural Network for NextG Vulnerability Detection. *IEEE Transactions on Quantum Engineering*.
- Preskill, J. 2018. Quantum computing in the NISQ era and beyond. *Quantum*, 2: 79.
- Ruff, L.; Vandermeulen, R. A.; Görnitz, N.; Binder, A.; Müller, E.; Müller, K.-R.; and Kloft, M. 2019. Deep semi-supervised anomaly detection. *arXiv preprint arXiv:1906.02694*.
- Salehi, M.; Sadjadi, N.; Baselizadeh, S.; Rohban, M. H.; and Rabiee, H. R. 2021. Multiresolution knowledge distillation for anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14902–14912.
- Silvestre-Blanes, J.; Albero-Albero, T.; Miralles, I.; Pérez-Lloréns, R.; and Moreno, J. 2019. A Public Fabric Database for Defect Detection Methods and Results. *Autex Research Journal*, 19: 363 – 374.
- Sultani, W.; Chen, C.; and Shah, M. 2018. Real-world anomaly detection in surveillance videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6479–6488.
- Sun, X.; Yang, Z.; Zhang, C.; Ling, K.-V.; and Peng, G. 2020. Conditional gaussian distribution learning for open set recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13480–13489.
- Tabernik, D.; Šela, S.; Skvarč, J.; and Skočaj, D. 2019. Segmentation-based deep-learning approach for surface-defect detection. *Journal of Intelligent Manufacturing*, 31(3): 759–776.
- Tack, J.; Mo, S.; Jeong, J.; and Shin, J. 2020. Csi: Novelty detection via contrastive learning on distributionally shifted instances. *Advances in neural information processing systems*, 33: 11839–11852.
- Tang, J.; Hua, F.; Gao, Z.; Zhao, P.; and Li, J. 2024. Gadbench: Revisiting and benchmarking supervised graph anomaly detection. *Advances in Neural Information Processing Systems*, 36.
- Tian, Y.; Pang, G.; Chen, Y.; Singh, R.; Verjans, J. W.; and Carneiro, G. 2021. Weakly-supervised video anomaly detection with robust temporal feature magnitude learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4975–4986.
- Wang, M.; Huang, A.; Liu, Y.; Yi, X.; Wu, J.; and Wang, S. 2023. A quantum-classical hybrid solution for deep anomaly detection. *Entropy*, 25(3): 427.
- Wei, L.; Liu, H.; Xu, J.; Shi, L.; Shan, Z.; Zhao, B.; and Gao, Y. 2023. Quantum machine learning in medical image analysis: A survey. *Neurocomputing*, 525: 42–53.
- Wieler, M.; Hahn, T.; and Hamprecht, F. A. 2023. Weakly Supervised Learning for Industrial Optical Inspection.
- Yao, X.; Li, R.; Zhang, J.; Sun, J.; and Zhang, C. 2023. Explicit Boundary Guided Semi-Push-Pull Contrastive Learning for Supervised Anomaly Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 24490–24499.
- Yoshihashi, R.; Shao, W.; Kawakami, R.; You, S.; Iida, M.; and Naemura, T. 2019. Classification-reconstruction learning for open-set recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4016–4025.
- Zhang, H.; Li, A.; Guo, J.; and Guo, Y. 2020. Hybrid models for open set recognition. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, 102–117. Springer.
- Zhu, J.; Ding, C.; Tian, Y.; and Pang, G. 2024. Anomaly heterogeneity learning for open-set supervised anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17616–17626.