
Diffusion-Inspired Quantum Noise Mitigation in Parameterized Quantum Circuits

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Parameterized Quantum Circuits (PQCs) have been acknowledged as a leading
2 strategy to utilize near-term quantum advantages in multiple problems, including
3 machine learning and combinatorial optimization. When applied to specific tasks,
4 the parameters in the quantum circuits are trained to minimize the target function.
5 Although there have been comprehensive studies to improve the performance of
6 the PQCs on practical tasks, the errors caused by the quantum noise downgrade
7 the performance when running on real quantum computers. In particular, when the
8 quantum state is transformed through multiple quantum circuit layers, the effect
9 of the quantum noise happens cumulatively and becomes closer to the maximally
10 mixed state or complete noise. This paper studies the relationship between the
11 quantum noise and the diffusion model. Then, we propose a novel diffusion-
12 inspired learning approach to mitigate the quantum noise in the PQCs and reduce
13 the error for specific tasks. Through our experiments, we illustrate the efficiency
14 of the learning strategy and achieve state-of-the-art performance on classification
15 tasks in the quantum noise scenarios.

1 Introduction

17 Quantum machine learning (QML) [4, 8, 21, 32] is an emerging and promising interdisciplinary
18 research direction in the fields of quantum computing and artificial intelligence. In this area, quantum
19 computers are expected to enhance machine learning algorithms through their inherent parallel
20 characteristics, thus demonstrating quantum advantages to solve some computational tasks out of
21 reach even of classical supercomputers [12]. With the increasing enormous efforts from academia and
22 industry, current quantum devices (usually acknowledged as the noisy intermediate-scale quantum
23 (NISQ) devices [28]) already can show quantum advantages on specific carefully designed tasks
24 [1, 45] despite their limitations in quantum circuit width and depth. Moreover, prior experiments
25 represent evidence for the utility of quantum computing on NISQ devices [17]. Thus, the NISQ
26 devices open a direction to explore the quantum advantages of quantum machine learning tasks
27 and leading strategies of hybrid classical-quantum algorithms, including parameterized quantum
28 circuits. Parameterized Quantum Circuits (PQCs) contain trainable parameters, offer a concrete
29 way to implement algorithms, and demonstrate quantum supremacy in the NISQ era. Even at low
30 circuit depth, some classes of PQCs are capable of generating highly non-trivial outputs [12, 3].
31 For example, classical resources cannot efficiently simulate the class of PQCs called instantaneous
32 quantum polynomial time under well-believed complexity-theoretic assumptions. However, the
33 behavior and impact of quantum noise remain critical questions for quantum computers today. The
34 NISQ devices still suffer from a high error rate of 10^{-2} to 10^{-4} , much higher than CPUs/GPUs with
35 an error rate of 10^{-18} . The quantum errors, unfortunately, introduce a detrimental influence on PQCs
36 accuracy.

37 Noise mitigation techniques [38, 43, 20, 10, 36] have been proposed to reduce the noise impact.
 38 However, these methods do not utilize the unique characteristics of PQCs and can only be applied
 39 to the inference process of the PQCs. Meanwhile, prior PQCs work [11, 3, 15] does not study the
 40 impact caused by the quantum noise.

41 Diffusion models are inspired by non-equilibrium thermodynamics [34, 13, 35]. They are defined as
 42 a Markov chain of diffusion steps to gradually add random noise to data and then learn to reverse
 43 the diffusion process to construct desired data from the noise. In quantum computing, random
 44 noise can be added into a quantum state by depolarizing until obtaining a maximally mixed state
 45 [18]. Inspired by the quantum depolarizing channel, some prior works propose diffusion models on
 46 quantum computing [5, 6]. However, more literature is needed to study the relationship between
 47 diffusion models and learnable quantum noise mitigation.

48 **Contributions of this Work:** In this work, we present the insights of diffusion-inspired modeling
 49 in the problem of PQCs. We show the relationship between diffusion and denoising processes in
 50 diffusion models and the forward and reverse processes in quantum states. The contributions of this
 51 paper are three-fold. First, we investigate the quantum noise properties in the PQCs and express
 52 the similarity between diffusion models and quantum computing in the noised PQCs. Then, we
 53 introduce a diffusion-inspired quantum noise mitigation framework for the PQCs. Second, from
 54 the diffusion-inspired quantum noise mitigation framework, we propose a novel loss function, i.e.,
 55 forward-backward quantum divergence loss, to learn the quantum noise model for mitigation. Finally,
 56 our proposed method is benchmarked on various specific tasks and achieves State-of-the-Art (SOTA)
 57 results compared to the prior methods.

58 2 Related Work

59 **Quantum Machine Learning.** There has been recent interest in studying the combination between
 60 quantum computing and machine learning. Early studies explored the quantum algorithms in linear
 61 machine learning, including clustering [21], principal component analysis [22], least-squares fitting
 62 [33, 16], and binary classification [30], to utilize the quantum speedup over classical machine learning
 63 algorithms. Prior work focused on the quantum neural networks using the framework of variational
 64 quantum algorithms or parameterized quantum circuits [27, 23]. Cong et al. [9] introduced the
 65 quantum convolutional neural network that extends the fundamental properties of classical CNNs to
 66 quantum computing while requiring less trainable parameters. Meanwhile, Bausch [2] proposed the
 67 quantum recurrent neural networks by utilizing the structure of the variational quantum eigensolver
 68 circuits. Huang et al. [14] presented hybrid quantum generative adversarial networks to generate data
 69 via quantum computers effectively. Romero et al. [31] introduced quantum autoencoders to reduce
 70 the dimensionality of quantum states.

71 **Quantum Noise Mitigation.** Prior studies on quantum noise in the quantum circuit and approaches
 72 to mitigate the errors caused by quantum noise exist. Li and Benjamin [20] and Temme et al. [38]
 73 introduced zero-noise extrapolation that tries to obtain zero-noise value by using data points at
 74 different circuit fault rates and computing the expected value at circuit fault rate zero. Temme et al.
 75 [38] also first presented the probabilistic error cancellation that reformulates the noise model as a
 76 linear combination and estimates this noise model via probability fitting with sampled quantum states.
 77 Czarnik et al. [10] and Strikis et al. [36] applied learning-based methods to obtain the error-mitigated
 78 expectation value using training circuits. QuantumNAT [40] was introduced to reduce the error
 79 in the PQCs via post-measurement processing to mitigate the difference between quantum feature
 80 distributions in noise-free and noisy cases. However, the mitigation processing on quantum circuit
 81 operations is not considered.

82 3 Background

83 **Quantum Basics.** In general, quantum information is described by quantum states [26]. An n -
 84 qubit quantum state is mathematically represented by a density matrix $\rho \in \mathbb{C}^{2^n \times 2^n}$ with property
 85 $\text{Tr}(\rho) = 1$. If $\text{Rank}(\rho) = 1$, the quantum state ρ is a pure state; otherwise, it is a mixed state. A
 86 pure state can also be represented by a unit vector $|\psi\rangle \in \mathbb{C}^{2^n}$, where $\rho = |\psi\rangle\langle\psi|$ and $\langle\psi| = |\psi\rangle^\dagger$. A

Petz-Rényi Divergence between Quantum State and Maximally Mixed State




Figure 1: **The divergence between the quantum state and maximally mixed state when running operations on the IBM Quito quantum system.** After a number of quantum state operations with noise, the quantum state becomes closer to the full quantum noise, i.e., maximally mixed state. It shows the relation between the effect of quantum noise and the number of quantum operations in PQCs. Note that although the amplitude damping makes the quantum state close to the $|0\rangle$ state, the quantum state distribution still gets closer to the maximally mixed state in the first 1200 operations.

87 mixed state can be defined as a weighted sum of pure states and presented as in Eqn. (1).

$$\rho = \sum_i \lambda_i |\psi_i\rangle\langle\psi_i|, \quad \lambda_i \geq 0, \sum_i \lambda_i = 1 \quad (1)$$

88 Specifically, a mixed state whose density matrix is proportional to the identity matrix is called the
 89 maximally mixed state $\mathbb{I}_n = \frac{I}{2^n}$. Physically, it is a uniform mixture of states on an orthonormal basis.
 90 It means that all states occur with the same probability.

91 A quantum state ρ can be evolved to another state ρ' through a quantum circuit (or gate) mathemati-
 92 cally represented by a unitary matrix U , i.e., $\rho' = U\rho U^\dagger$. Typical single-qubit gates include Pauli
 93 gates:

$$\sigma_x = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad \sigma_y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \quad \sigma_z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad (2)$$

94 and their corresponding rotation gates $R_\sigma(\theta) = e^{-i\theta\sigma/2}$ with a parameter θ and $\sigma \in \{\sigma_x, \sigma_y, \sigma_z\}$. A
 95 multi-qubit gate can be either an individual gate (e.g., CNOT) or a tensor product of single-qubit gates.
 96 To get classical information from a quantum state ρ' , one needs to perform quantum measurements,
 97 e.g., calculating the expectation value $\langle H \rangle = \text{Tr}(H\rho')$ of a Hermitian matrix H , and we often call H
 98 an observable.

99 **Parameterized Quantum Circuits.** The parameterized quantum circuits (PQCs), also known as
 100 variational quantum circuits (VQCs) [3, 23], are a special kind of quantum circuit with parameters
 101 that can be optimized or learned iteratively. The PQCs are composed of three parts, including data
 102 encoding, parameterized layer, and quantum measurements.

103 PQCs use a hybrid quantum-classical procedure to optimize the trainable parameters iteratively. The
 104 popular optimization approaches include gradient descent [37], parameter-shift rule [42, 23], and
 105 gradient-free techniques [25, 7]. All learning methods take the training data as input and evaluate the
 106 model performance by comparing the predicted and ground-truth labels. Based on this evaluation,
 107 the methods update the model parameters for the next iteration and repeat the process until the model
 108 converges and achieves the desired performance. The hybrid method performs the evaluation and
 109 parameter optimization on a classical computer, while the model inference is processed on a quantum
 110 computer.

111 **Noise in Quantum Computing.** Noise refers to the multiple factors that can affect the accuracy of
 112 the calculations a quantum computer performs. Because of the noise in quantum computing, the
 113 transformation of the quantum state can cause errors. Pauli channel error is one of the common
 114 noise models that cause probabilistic error between Pauli operations defined as $\tilde{\rho} = \hat{\Lambda}(\rho) = \rho +$
 115 $\sum_{\sigma \in \mathcal{K}} \lambda_\sigma (\sigma \rho \sigma^\dagger - \rho)$, where \mathcal{K} is a set of Pauli operations and λ_σ is the probability that the error
 116 caused by σ . Here, we define ρ_i as a noise-free quantum state and $\tilde{\rho}_i$ as a quantum state in the noise
 117 scenario. These errors are cumulated via quantum circuit layers that make the quantum state closer to

the maximally mixed state, i.e., the fully noise quantum state as shown in Fig. 1. Moreover, the noise model $\tilde{\Lambda}(\rho)$ can be formulated to an alternative model $\Lambda(\rho)$:

$$\Lambda(\rho) = \prod_{\sigma \in \mathcal{K}} (w_\sigma \cdot \sigma + (1 - w_\sigma) \sigma \cdot \sigma^\dagger) \rho \quad (3)$$

where $w_\sigma = 2^{-1}(1 + e^{-2\lambda_\sigma})$. In the PQCs, the error caused by the noise can occur in the data encoding, parameterized layer, or even the quantum measurements. However, as the quantum operations are mostly in the parameterized layers, we focus on the quantum noise mitigation in this part of the PQCs.

4 Diffusion-inspired Modeling for Quantum Denoising

4.1 Diffusion Modeling Revisited

Diffusion models [13] are a class of latent variable models that learn a generative model to reverse a fixed probabilistic noising process $\mathbf{x}_0 \rightarrow \mathbf{x}_1 \rightarrow \dots \rightarrow \mathbf{x}_T$, where $\mathbf{x}_1, \dots, \mathbf{x}_T$ are latent variables of the same dimensionality of the data $\mathbf{x}_0 \sim q(\mathbf{x}_0)$. This probabilistic noising process gradually adds noise to clean data \mathbf{x}_0 until no information remains, i.e., pure noise $\mathbf{x}_T \sim p(\mathbf{x}_T)$. The divergence of data in this process can be formulated as:

$$D(q(\mathbf{x}_0) || \mathcal{N}_G) > D(q(\mathbf{x}_1) || \mathcal{N}_G) > \dots > D(q(\mathbf{x}_T) || \mathcal{N}_G) \quad (4)$$

where \mathcal{N}_G is the probability distribution of the pure Gaussian noise. In most cases, the Kullback-Leibler divergence is applied to compute the difference between distributions of classical data. For continuous data, the forward process is defined as a fixed Markov chain $q(\mathbf{x}_t | \mathbf{x}_{t-1})$ with Gaussian transitions. The Markov chain of the reverse process can be obtained by approximating the true posterior $q(\mathbf{x}_{t-1} | \mathbf{x}_0, \mathbf{x}_t)$ with a model $p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t)$. Therefore, sampling new data \mathbf{x}_0 from the modeled data distribution $p_\theta(\mathbf{x}_0) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t)$ is performed by starting from random noise $\mathbf{x}_T \sim p(\mathbf{x}_T)$ and gradually denoising it over T steps $\mathbf{x}_T \rightarrow \mathbf{x}_{T-1} \rightarrow \dots \rightarrow \mathbf{x}_0$.

4.2 Diffusion-inspired Quantum Noise Mitigation

Quantum computing has similar properties to diffusion models when transforming the quantum state through quantum circuits. In the following, we derive a diffusion-inspired learning for quantum noise mitigation. The two main components of this learning strategy are forward and reverse processes. The forward process transforms a quantum state into a next state through a quantum circuit layer and adds noise to the quantum state as the nature of PQCs in the noise scenario. As the resulting quantum state has noise, a denoising process is considered to reduce the error. In the general diffusion models, this process is solved by the reverse process to train the denoising module.

Given a quantum state ρ_{i-1} and a transformed quantum state ρ_i from ρ_{i-1} , without noise, as the Shannon entropy of the quantum state is unchanged by the unitary circuit $S(\rho) = S(U\rho U^\dagger)$, the divergence between ρ_i and the maximally mixed state \mathbb{I}_n is the same as the divergence between ρ_{i-1} and \mathbb{I}_n :

$$D(\rho_i || \mathbb{I}_n) = D(\rho_{i-1} || \mathbb{I}_n) \quad (5)$$

However, as shown in Fig. 1, in the noise case, the quantum noise makes the divergence of the noise quantum state $\tilde{\rho}_i$ closer to the maximally mixed state \mathbb{I}_n after the quantum transformation:

$$D(\tilde{\rho}_i || \mathbb{I}_n) < D(\tilde{\rho}_{i-1} || \mathbb{I}_n) \quad (6)$$

This phenomenon is similar to the forward process of the diffusion model when the classical data becomes closer to the Gaussian noise, as shown in Eqn. (4). Moreover, as shown in Eqn. (3), the quantum noise in the PQCs can be considered as a Markov chain analogous to the diffusion process.

On the other hand, as the quantum circuits are reversible, the quantum state ρ_i can be backward into the quantum state ρ_{i-1} of the previous quantum circuit layer in the noise-free scenario. In the quantum noise scenario, as the quantum state $\tilde{\rho}_i$ has noise, the backward process causes errors that the resulting quantum state is not similar to its original state. Hence, a noise mitigation module models and reduces the quantum noise, making the backward quantum state closer to the original state. Fig. 2 illustrates the forward and backward processes. From this observation, we proposed a diffusion-inspired learning method to model the quantum noise distribution for noise mitigation.

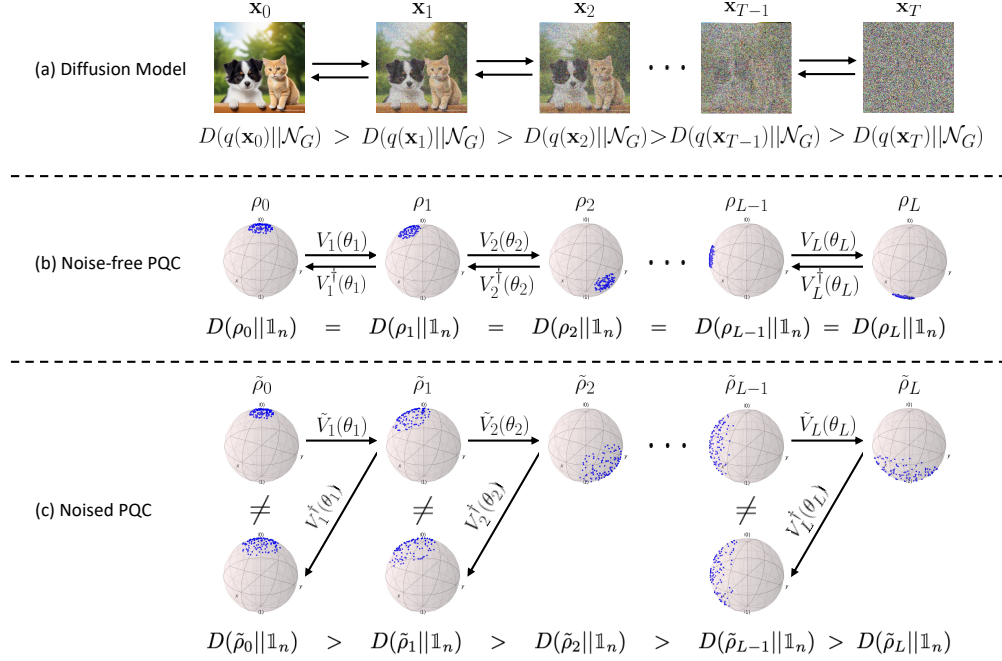


Figure 2: (a) In the traditional diffusion model, the distribution of the data $q(\mathbf{x}_i)$ becomes closer to the Gaussian noise \mathcal{N}_G through diffusion steps. (b) In the noise-free scenario, each quantum transformation $V_i(\theta_i)$ can be reversed and maintain the quantum information of the previous circuit layer via reverse operation $V_i^\dagger(\theta_i)$. (c) In the PQCs having quantum noise case, the quantum state transformations $\tilde{V}_i(\theta_i) = \Lambda_i \circ V_i(\theta_i)$ are affected by noise Λ_i that makes the reversed quantum state different from the original one. Motivated by this, we propose a learning-based method to mitigate the quantum noise via computing the divergence between the quantum state and its noisy forward-and-backward state.

5 The Proposed Method

To address the quantum noise in the PQCs, we first consider the general PQCs framework and their quantum noise mitigation module in the noise scenario. Then, we study the distribution of the quantum noise and quantum states. Finally, we introduce a novel Forward-backward Quantum Divergence Loss for quantum noise mitigation.

5.1 The Overall Framework

Fig. 3 illustrates the overall framework of the quantum noise mitigation in PQCs. Given classical data \mathbf{x} , a data encoder U is applied to encode the classical data into a quantum state ρ_0 . Then, the PQCs are trained to transform and measure the quantum state for a specific task. In the noise-free scenario, learnable unitary matrices $V_i(\theta_i)$ are applied to transform the quantum state:

$$\rho_i = V_i(\theta_i)\rho_{i-1}V_i^\dagger(\theta_i) \quad (7)$$

The design of the parameterized circuits is described in Section 6. Because of the noise in quantum computing, we can model the noise by a matrix Λ_i :

$$\tilde{\rho}_i = \tilde{V}_i(\theta_i)\tilde{\rho}_{i-1}\tilde{V}_i^\dagger(\theta_i), \quad \tilde{V}_i(\theta_i) = \Lambda_i \circ V_i(\theta_i) \quad (8)$$

It causes errors in the PQCs when training and inferring. Hence, quantum noise mitigation layers Λ_i^{-1} are applied to reduce these errors. To obtain these quantum noise mitigation layers, the probabilities of the quantum noise have to be computed. Therefore, in this work, we aim to learn the noise model Λ_i to mitigate the error caused by the quantum noise.

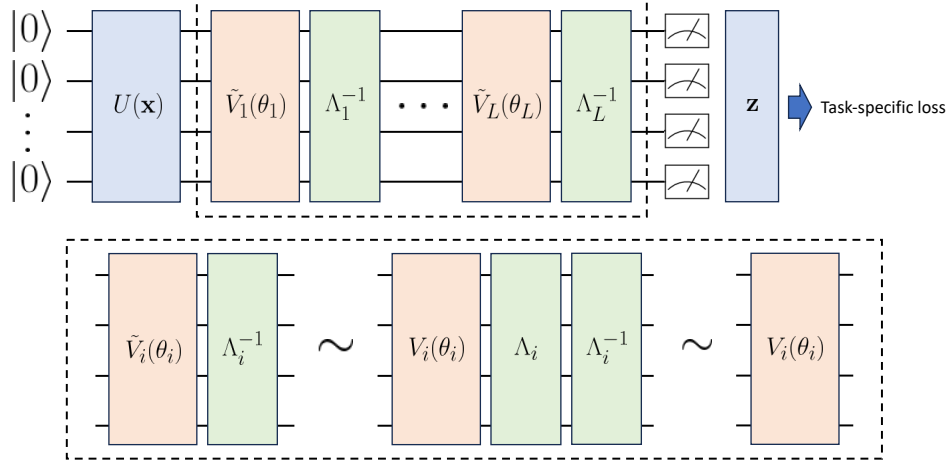


Figure 3: An overall framework of PQCs with quantum noise mitigation. For every quantum circuit layer $\tilde{V}_i(\theta_i) = \Lambda_i \circ V_i(\theta_i)$, a quantum noise mitigation layer Λ_i^{-1} is applied to reduce the error.

5.2 Quantum Noise Distribution Learning

As we want to learn the quantum noise model of Λ_i , we define a probability density function $p_\omega(\Lambda_i)$ where ω are learnable parameters of the overall model. Moreover, to train the PQCs for the specific task, we also learn the distribution of the quantum states $\tilde{\rho}_i$ where $i \in \{0, 1, \dots, L-1\}$ to represent the data for specific tasks. Since Λ_i and $\tilde{\rho}_{i-1}$ are independent, the objective of ω can be defined as:

$$\omega^* = \arg \min_{\omega} \sum_{i=1}^L (-\log p_\omega(\Lambda_i) - \log p_\omega(\tilde{\rho}_{i-1})) = \arg \min_{\omega} \sum_{i=1}^L -\log p_\omega(\Lambda_i, \tilde{\rho}_{i-1}) \quad (9)$$

We also define the forward process from the quantum state $\tilde{\rho}_{i-1}$ to the next quantum states $\tilde{\rho}_i$ as $q(\tilde{\rho}_i | \tilde{\rho}_{i-1})$. To optimize the objective function Eqn. (9), we optimize the upper bound on negative log-likelihood via Jensen's inequality:

$$\sum_{i=1}^L -\log p_\omega(\Lambda_i, \tilde{\rho}_{i-1}) \leq \mathbb{E}_q \left[-\sum_{i=1}^L \log \frac{p_\omega(\Lambda_i, \tilde{\rho}_{i-1}, \tilde{\rho}_i)}{q(\tilde{\rho}_i | \tilde{\rho}_{i-1})} \right] := \mathcal{L} \quad (10)$$

Then \mathcal{L} can be expanded as:

$$\mathcal{L} = \mathbb{E}_q \left[\underbrace{\sum_{i=1}^L D_{KL}(p_\omega(\tilde{\rho}_{i-1} | \Lambda_i, \tilde{\rho}_i) || q(\tilde{\rho}_{i-1} | \tilde{\rho}_i))}_{\mathcal{L}_{fb}} - \underbrace{\sum_{i=1}^L \log p_\omega(\Lambda_i, \tilde{\rho}_i) - \log \frac{q(\tilde{\rho}_0)}{q(\tilde{\rho}_L)}}_{\mathcal{L}_{task}} \right] \quad (11)$$

(See Supplementary for details of Eqn. (10) and Eqn. (11).) The Kullback-Leibler divergence in Eqn. (11) compares the original quantum state and its forward-backward in the quantum circuit layer. Thus, we propose a learning method to model the quantum noise via quantum divergence.

5.3 Forward-backward Quantum Divergence Loss

In quantum information theory, to measure the similarity between two quantum states, fidelity is defined as the probability that one state will pass a test to identify as the other. In this work, we utilize the fidelity of quantum states to define a forward-backward quantum divergence loss. Let $F(\rho, \sigma) \in [0, 1]$ be the fidelity between quantum states ρ and σ , then $F(\rho, \sigma) = 1$ when ρ and σ are completely similar and vice versa.

Given a quantum state $\tilde{\rho}_{i-1}$ and its forwarded quantum state $\tilde{\rho}_i$ as Eqn. (8), we define a learnable noise mitigation layer $\Lambda_i^{-1}(\cdot, \omega_i)$ to reduce the noise from quantum state $\tilde{\rho}_i$.

$$\hat{\rho}_i = \Lambda_i^{-1}(\tilde{\rho}_i, \omega_i) \quad (12)$$




Figure 4: **The forward-backward quantum divergence loss in a quantum circuit.** For each layer $V_i(\theta_i)$ of the quantum circuit, the current quantum state $\tilde{\rho}_i$ is denoised by a learnable noise mitigation layer Λ_i^{-1} .

As the quantum circuits are reversible, the quantum state $\hat{\rho}_i$ can be reversed to the previous state via the quantum circuit $V_i^\dagger(\theta_i)$.

$$\hat{\rho}_{i-1} = V_i^\dagger(\theta_i) \hat{\rho}_i V_i(\theta_i) \quad (13)$$

Then, we define a forward-backward quantum divergence loss to compute the similarity between quantum states $\tilde{\rho}_{i-1}$ and $\hat{\rho}_{i-1}$ as:

$$\mathcal{L}_{fb}(\tilde{\rho}_{i-1}, \hat{\rho}_{i-1}) = -\log F(\tilde{\rho}_{i-1}, \hat{\rho}_{i-1}) \quad (14)$$

where $F(\rho, \sigma)$ is the fidelity between two quantum states ρ and σ . As the forward-backward procedure can be processed in L quantum circuit layers, the overall forward-backward quantum divergence loss is computed as:

$$\mathcal{L}_{fb} = \frac{1}{L} \sum_{i=1}^L \mathcal{L}_{fb}(\tilde{\rho}_{i-1}, \hat{\rho}_{i-1}) = -\frac{1}{L} \sum_{i=1}^L \log F(\tilde{\rho}_{i-1}, \hat{\rho}_{i-1}) \quad (15)$$

The forward-backward quantum divergence loss process is illustrated in Fig. 4. To compute the fidelity between quantum states, a widely used metric is the quantum Rényi divergence [24]:

$$F(\rho, \sigma) = D(\rho || \sigma) = 2 \log \text{Tr} \left[\sqrt{\sqrt{\sigma} \rho \sqrt{\sigma}} \right] \quad (16)$$

5.4 Task-specific Training

In addition to learning the quantum noise model, the PQCs are also trained for specific tasks, such as classification and clustering. As shown in Eqn. (11), the PQCs can be trained to obtain a desired distribution of quantum states while learning the quantum noise model. Thus, a task-specific loss \mathcal{L}_{task} is defined for the PQCs. Let H be the observable of the PQC, a classical information obtained from the PQCs is measured as $\langle H \rangle = \text{Tr}(H \hat{\rho}_L)$. Given n qubits for the PQC, we define n observables $H_i = I^{\otimes i} \otimes \sigma_z \otimes I^{\otimes(n-i-1)}$ where $I \in \mathbb{R}^{2 \times 2}$ is the identity matrix and σ_z is the Pauli-Z matrix. A classical information vector $\mathbf{z} = \{\langle H_i \rangle\}_{i=0}^{n-1}$ is defined for inference and task-specific loss computing. In this work, the PQCs are applied for the classification tasks. Then the task-specific loss \mathcal{L}_{task} is defined as:

$$\mathcal{L}_{task} = - \sum_{i=0}^c \hat{y}_i \log y_i \quad (17)$$

where c is the number of classes, $\hat{y}_i \in \{0, 1\}$ is the classification ground-truth, and $\mathbf{y} = f(\mathbf{z})$ is the prediction from classical information vector \mathbf{z} . Finally, the total loss of the quantum noise mitigation training for the PQCs is computed as:

$$\mathcal{L}_{total} = \alpha_{fb} \mathcal{L}_{fb} + \alpha_{task} \mathcal{L}_{task} \quad (18)$$

where $\alpha_{fb}, \alpha_{task} \in \mathbb{R}$ is the training hyperparameter.

6 Experimental Results

In this section, we evaluate the proposed method in the quantum noise context. We first describe the experiment setups, including datasets, implementation, and evaluation protocol. Then, we present the ablation studies to illustrate the effectiveness of our proposed method. Finally, we illustrate the numerical results of the proposed method and compare our approach with prior quantum noise mitigation methods.

6.1 Experiment Setups

Datasets. Following [40], we evaluate the proposed method on four classification tasks, including MNIST [19] 4-class (0, 1, 2, 3) and 2-class (3, 6); and Fashion [44] 4-class (t-shirt/top, trouser, pullover, dress) and 2-class (dress, shirt). The images are resized into 8×8 and encoded via phase encoding using multiple rotation circuits.

Implementation. This work uses the quantum simulation for the PQCs and the quantum noise. For each layer, different learnable circuits are applied, including RX, U2 (RX + RY), and U3 (RX + RY + RZ). The efficiency of each circuit is shown in the ablation studies. Then, a non-learnable circuit, i.e., controlled-NOT (CNOT), is used. We use four qubits for the experiment and run simulations of the IBM quantum systems via Qiskit SDK [29] on a Quadro RTX 8000 GPU. The PQCs model training and testing is implemented based on the TorchQuantum library [41].

Evaluation Protocol. For evaluation, we compare the proposed method with other settings, i.e., training and testing on noise-free PQCs, training on noise-free PQCs and testing on noised PQCs, and training and testing on noised PQCs as a baseline. We also reimplement the training and evaluation of the prior quantum noise mitigation methods, i.e., QuantumNAT [40] normalizing and quantizing the measurement to match the output distribution with the noise-free case, and Van Den Berg et al. [39] learning the probabilistic quantum noise model by sampling random quantum states. The average accuracy metric is used in our experiments. The experiments are processed five times for each setting, and the standard deviation is used to compute the variance of the results.

6.2 Ablation Studies

Our ablative experiments study the effectiveness of our proposed method on the performance of the PQCs on the MNIST-4 benchmark.

Effectiveness of the Forward-backward Quantum Divergence Loss. We evaluate the impact of Forward-backward Quantum Divergence Loss (\mathcal{L}_{fb}) in improving the performance of the PQCs in the noise scenario. To demonstrate the efficiency of our proposed approach with diverse quantum circuit designs, we evaluate it with three different circuit designs: RX + CNOT, U2 + CNOT, and U3 + CNOT. As shown in Table 1, the forward-backward loss \mathcal{L}_{fb} has significant improvements compared to using the noise mitigation with the task-specific loss \mathcal{L}_{task} only. In particular, the accuracy of the three circuit designs, i.e., RX + CNOT, U2 + CNOT, and U3 + CNOT, has been improved from 42.61% to 45.89%, from 43.17% to 46.44%, and from 43.37% to 46.37%, respectively. It shows

Circuits	\mathcal{L}_{fb}	Step size		
		4	2	1
RX + CNOT	✓	41.76 ± 1.46 44.30 ± 1.65	41.90 ± 1.53 45.55 ± 1.46	42.61 ± 1.37 45.89 ± 1.43
U2 + CNOT	✓	40.58 ± 1.63 44.42 ± 1.69	42.91 ± 1.78 45.50 ± 1.73	43.17 ± 1.57 46.44 ± 1.60
U3 + CNOT	✓	42.51 ± 1.46 45.36 ± 1.29	43.18 ± 1.47 46.20 ± 1.45	43.37 ± 1.39 46.37 ± 1.20

Table 1: **Effectiveness of our approach on the MNIST-4 benchmark.** We compute the mean accuracies (%) and their standard deviation to evaluate the approach with different circuit designs, i.e., RX + CNOT, U2 + CNOT, and U3 + CNOT, with different noise mitigation step sizes, and without or with the forward-backward quantum divergence loss \mathcal{L}_{fb} .




Figure 5: Ablation studies on different numbers of circuit layers.

Method	MNIST-4	MNIST-2	Fashion-4	Fashion-2
Noise-free testing	49.87 \pm 0.31	92.99 \pm 0.19	49.53 \pm 0.35	74.60 \pm 0.47
Noise-free training	37.24 \pm 1.88	79.07 \pm 8.22	39.08 \pm 1.88	66.35 \pm 3.39
Baseline	43.06 \pm 1.06	82.37 \pm 4.62	43.70 \pm 1.43	68.50 \pm 3.12
QuantumNAT [40]	43.59 \pm 1.15	82.32 \pm 4.68	43.60 \pm 1.40	68.55 \pm 3.11
Van Den Berg et al. [39]	45.15 \pm 1.45	84.32 \pm 4.30	45.25 \pm 1.89	69.15 \pm 2.44
Ours	46.44 \pm 1.60	85.32 \pm 4.35	46.83 \pm 1.81	70.50 \pm 2.57

Table 2: **Experimental accuracies (%) on 2- and 4-class benchmarks.** A noise-free testing is defined as training and testing on a noise-free PQC while the noise-free trained model is tested on the quantum noise scenario. We show the mean accuracies and their standard deviation for variances of the results.

that the forward-backward loss \mathcal{L}_{fb} helps to learn the noise distribution better for quantum noise mitigation.

Effectiveness on Different Noise Mitigation Step Sizes. We investigate the impact of the quantum noise mitigation step size on the performance. To achieve this, we conduct experiments on the MNIST-4 benchmark with a 4-layer PQCs model using three different step sizes, i.e., 4, 2, and 1. In this case, step size 4 means we forward the quantum state $\tilde{\rho}_0$ through 4 layers of the PQCs and backward the resulting state $\tilde{\rho}_4$ to the initial state $\tilde{\rho}_0$ for the forward-backward quantum divergence loss \mathcal{L}_{fb} . Similarly, step size 2 means we forward the quantum state $\tilde{\rho}_i$ through 2 layers and backward the resulting state $\tilde{\rho}_{i+2}$ for the forward-backward loss \mathcal{L}_{fb} , and step size 1 means we compute the loss for every layer. As depicted in Table 1, the performance gradually increases when the step size is decreased for all different circuit designs. It shows that quantum noise mitigation works effectively when the number of operation circuits for each process is negligible.

Effectiveness on Different Numbers of Circuit Layers. We evaluate the robustness of the proposed quantum noise mitigation approach on different numbers of circuit layers of the PQCs. As illustrated in Fig. 5, when the number of layers is increased, the performance of the baseline is dramatically dropped. Meanwhile, our proposed approach shows a slight decline in accuracy. Hence, the proposed approach can mitigate the quantum noise effectively even if the PQCs have many operation circuits.

6.3 Evaluation Results

As shown in Table 2, our method outperforms previous methods evaluated on the two datasets, i.e., MNIST and Fashion. In particular, in the MNIST dataset, our method achieves the accuracy of 46.44% and 85.32% on the MNIST-4 and MNIST-2 benchmarks, respectively, which shows better than the previous methods. Meanwhile, the results for the Fashion dataset are 46.83% and 70.50% for the Fashion-4 and Fashion-2 benchmarks.

7 Conclusions

This paper has studied the quantum noise in PQCs and introduces a novel, learnable quantum noise mitigation approach to improve their robustness when running on NISQ devices. By revisiting the diffusion model and investigating the relationship between the diffusion and denoising processes and the forward and reverse processes in quantum states, we have shown the quantum noise distribution learning ability while training the PQCs for specific tasks. Hence, a novel forward-backward quantum divergence loss function has been introduced to learn the quantum noise model for quantum noise mitigation. The experimental results on various benchmarks have shown our state-of-the-art method.

Limitations: Our paper has chosen specific configurations of quantum systems and hyperparameters to support our hypothesis theoretically in the simulation. However, other aspects, such as non-Markovian noise models, computational complexity, or learning hyperparameters, have yet to be fully investigated.

Broader Impact: This work studies a diffusion-inspired approach to quantum noise model learning for noise mitigation in the PQCs. Our contributions emphasize the importance of quantum noise mitigation in quantum machine learning and provide a solution to reduce the error and increase the robustness of the PQCs when running in real quantum computers.

References

- [1] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.
- [2] J. Bausch. Recurrent quantum neural networks. *Advances in neural information processing systems*, 33:1368–1379, 2020.
- [3] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini. Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4):043001, 2019.
- [4] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd. Quantum machine learning. *Nature*, 549(7671):195–202, 2017.
- [5] A. Cacioppo, L. Colantonio, S. Bordoni, and S. Giagu. Quantum diffusion models. *arXiv preprint arXiv:2311.15444*, 2023.
- [6] C. Chen and Q. Zhao. Quantum generative diffusion model. *arXiv preprint arXiv:2401.07039*, 2024.
- [7] S. Y.-C. Chen, C.-M. Huang, C.-W. Hsing, H.-S. Goan, and Y.-J. Kao. Variational quantum reinforcement learning via evolutionary optimization. *Machine Learning: Science and Technology*, 3(1):015025, 2022.
- [8] C. Ciliberto, M. Herbster, A. D. Ialongo, M. Pontil, A. Rocchetto, S. Severini, and L. Wossnig. Quantum machine learning: a classical perspective. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 474(2209):20170551, 2018.
- [9] I. Cong, S. Choi, and M. D. Lukin. Quantum convolutional neural networks. *Nature Physics*, 15(12):1273–1278, 2019.
- [10] P. Czarnik, A. Arrasmith, P. J. Coles, and L. Cincio. Error mitigation with clifford quantum-circuit data. *Quantum*, 5:592, 2021.
- [11] E. Farhi and H. Neven. Classification with quantum neural networks on near term processors. *arXiv preprint arXiv:1802.06002*, 2018.
- [12] A. W. Harrow and A. Montanaro. Quantum computational supremacy. *Nature*, 549(7671):203–209, 2017.
- [13] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [14] H.-L. Huang, Y. Du, M. Gong, Y. Zhao, Y. Wu, C. Wang, S. Li, F. Liang, J. Lin, Y. Xu, et al. Experimental quantum generative adversarial networks for image generation. *Physical Review Applied*, 16(2):024051, 2021.
- [15] W. Jiang, J. Xiong, and Y. Shi. A co-design framework of neural networks and quantum circuits towards quantum advantage. *Nature communications*, 12(1):579, 2021.
- [16] I. Kerenidis and A. Prakash. Quantum gradient descent for linear systems and least squares. *Physical Review A*, 101(2):022316, 2020.
- [17] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. Van Den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, et al. Evidence for the utility of quantum computing before fault tolerance. *Nature*, 618(7965):500–505, 2023.
- [18] C. King. The capacity of the quantum depolarizing channel. *IEEE Transactions on Information Theory*, 49(1):221–229, 2003.
- [19] Y. LeCun, C. Cortes, and C. Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
- [20] Y. Li and S. C. Benjamin. Efficient variational quantum simulator incorporating active error minimization. *Physical Review X*, 7(2):021050, 2017.

- [21] S. Lloyd, M. Mohseni, and P. Rebentrost. Quantum algorithms for supervised and unsupervised machine learning. *arXiv preprint arXiv:1307.0411*, 2013.
- [22] S. Lloyd, M. Mohseni, and P. Rebentrost. Quantum principal component analysis. *Nature Physics*, 10(9):631–633, 2014.
- [23] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii. Quantum circuit learning. *Physical Review A*, 98(3):032309, 2018.
- [24] M. Müller-Lennert, F. Dupuis, O. Szehr, S. Fehr, and M. Tomamichel. On quantum rényi entropies: A new generalization and some properties. *Journal of Mathematical Physics*, 54(12), 2013.
- [25] G. Nannicini. Performance of hybrid quantum-classical variational heuristics for combinatorial optimization. *Physical Review E*, 99(1):013304, 2019.
- [26] M. A. Nielsen and I. L. Chuang. *Quantum computation and quantum information*, volume 2. Cambridge university press Cambridge, 2001.
- [27] M. Panella and G. Martinelli. Neural networks with quantum architecture and quantum learning. *International Journal of Circuit Theory and Applications*, 39(1):61–77, 2011.
- [28] J. Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.
- [29] Qiskit contributors. Qiskit: An open-source framework for quantum computing, 2023.
- [30] P. Rebentrost, M. Mohseni, and S. Lloyd. Quantum support vector machine for big data classification. *Physical review letters*, 113(13):130503, 2014.
- [31] J. Romero, J. P. Olson, and A. Aspuru-Guzik. Quantum autoencoders for efficient compression of quantum data. *Quantum Science and Technology*, 2(4):045001, 2017.
- [32] M. Schuld, I. Sinayskiy, and F. Petruccione. An introduction to quantum machine learning. *Contemporary Physics*, 56(2):172–185, 2015.
- [33] M. Schuld, I. Sinayskiy, and F. Petruccione. Prediction by linear regression on a quantum computer. *Physical Review A*, 94(2):022342, 2016.
- [34] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015.
- [35] J. Song, C. Meng, and S. Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [36] A. Strikis, D. Qin, Y. Chen, S. C. Benjamin, and Y. Li. Learning-based quantum error mitigation. *PRX Quantum*, 2(4):040330, 2021.
- [37] R. Sweke, F. Wilde, J. Meyer, M. Schuld, P. K. Fährmann, B. Meynard-Piganeau, and J. Eisert. Stochastic gradient descent for hybrid quantum-classical optimization. *Quantum*, 4:314, 2020.
- [38] K. Temme, S. Bravyi, and J. M. Gambetta. Error mitigation for short-depth quantum circuits. *Physical review letters*, 119(18):180509, 2017.
- [39] E. Van Den Berg, Z. K. Mineev, A. Kandala, and K. Temme. Probabilistic error cancellation with sparse pauli-lindblad models on noisy quantum processors. *Nature Physics*, 19(8):1116–1121, 2023.
- [40] H. Wang, J. Gu, Y. Ding, Z. Li, F. T. Chong, D. Z. Pan, and S. Han. Quantumnat: quantum noise-aware training with noise injection, quantization and normalization. In *Proceedings of the 59th ACM/IEEE design automation conference*, pages 1–6, 2022.
- [41] H. Wang, Z. Liang, J. Gu, Z. Li, Y. Ding, W. Jiang, Y. Shi, D. Z. Pan, F. T. Chong, and S. Han. Torchquantum case study for robust quantum circuits. In *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design*, pages 1–9, 2022.

- 403 [42] D. Wierichs, J. Izaac, C. Wang, and C. Y.-Y. Lin. General parameter-shift rules for quantum
404 gradients. *Quantum*, 6:677, 2022.
- 405 [43] R. Wille, L. Burgholzer, and A. Zulehner. Mapping quantum circuits to ibm qx architectures
406 using the minimal number of swap and h operations. In *Proceedings of the 56th Annual Design*
407 *Automation Conference 2019*, pages 1–6, 2019.
- 408 [44] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking
409 machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- 410 [45] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding,
411 Y. Hu, et al. Quantum computational advantage using photons. *Science*, 370(6523):1460–1463,
412 2020.

A Appendix / supplemental material

A.1 Quantum Noise Mitigation Module

In this work, we take account of the Pauli-Lindblad noise model for quantum noise mitigation. As shown in Eqn. (3), for each Pauli channel, the noise model is defined as $(w_\sigma \rho + (1 - w_\sigma) \sigma \rho \sigma)$ where $w_\sigma = 2^{-1}(1 + e^{-2\lambda_\sigma})$. Then, the inverse of the Pauli channel can be written as $(2w_\sigma - 1)^{-1}(w_\sigma \tilde{\rho} - (1 - w_\sigma) \sigma \tilde{\rho} \sigma)$. In overall, the quantum noise $\Lambda(\cdot)$ can be mitigated from the quantum noise mitigation layer $\Lambda^{-1}(\cdot)$ as:

$$\Lambda^{-1}(\tilde{\rho}) = \gamma \prod_{\sigma \in \mathcal{K}} (w_\sigma \cdot - (1 - w_\sigma) \sigma \cdot \sigma^\dagger) \tilde{\rho} \quad (19)$$

where $\gamma = \prod_{\sigma \in \mathcal{K}} (2w_\sigma - 1)^{-1} = \exp(\sum_{\sigma \in \mathcal{K}} 2\lambda_\sigma)$ is *sampling overhead*. In this case, $\{\lambda_\sigma\}_{\sigma \in \mathcal{K}}$ are learnable parameters for the quantum noise mitigation.

A.2 Quantum Noise Mitigation Learning Algorithm

The quantum noise mitigation learning process can be described in Algorithm 1.

Algorithm 1: Pseudo-code for the implementation of Quantum Noise Mitigation

Data:

$\{\mathbf{x}_i\}_{i=1}^N$: a set of N classical inputs

$\{\hat{y}_i\}_{i=1}^N \in \mathbb{R}^N$: a set of N labels

n : the number of qubits

L : the number of layers in the PQCs

$\alpha_{fb}, \alpha_{task}$: training hyperparameters

while not convergent do

$\tilde{\rho}_0 \leftarrow U(\mathbf{x})$ // Encode the classical data into a quantum state

$\mathcal{L}_{fb} \leftarrow 0$ // Initialize the forward-backward quantum divergence loss to zero

for $i \in [1..L]$ **do**

$\tilde{\rho}_i \leftarrow \tilde{V}_i(\theta_i) \tilde{\rho}_{i-1} \tilde{V}_i^\dagger(\theta_i)$ // Forward the quantum state in the noised case

$\hat{\rho}_i \leftarrow \Lambda_i^{-1}(\tilde{\rho}_i)$ // Reduce the quantum noise

$\hat{\rho}_{i-1} \leftarrow V_i^\dagger(\theta_i) \hat{\rho}_i V_i(\theta_i)$ // Backward the quantum state

$\mathcal{L}_{fb} \leftarrow \log F(\tilde{\rho}_{i-1}, \hat{\rho}_{i-1})$ // Compute the forward-backward loss

end

$\mathcal{L}_{fb} \leftarrow \frac{1}{L} \mathcal{L}_{fb}$ // Compute the average of the loss

$\mathbf{z} \leftarrow \{\langle H_i \rangle\}_{i=0}^{n-1}$ // Compute the classical information from the quantum state $\hat{\rho}_L$

$\mathbf{y} \leftarrow f(\mathbf{z})$ // Compute the task-specific output

$\mathcal{L}_{task} \leftarrow -\sum_{i=0}^c \hat{y}_i \log y_i$ // Compute the task-specific loss

$\mathcal{L}_{total} \leftarrow \alpha_{fb} \mathcal{L}_{fb} + \alpha_{task} \mathcal{L}_{task}$ // Compute the total loss

$\theta \leftarrow \theta - \lambda \nabla_\theta \mathcal{L}_{total}$ // Do backpropagation

end

A.3 Proof of Eqn. (10)

The upper bound of the objective function Eqn.(9) can be derived via Jensen's inequality as follows:

$$\begin{aligned} \sum_{i=1}^L -\log p_\omega(\Lambda_i, \tilde{\rho}_{i-1}) &= -\sum_{i=1}^L \log \int p_\omega(\Lambda_i, \tilde{\rho}_{i-1}, \tilde{\rho}_i) d\tilde{\rho}_i \\ &= -\sum_{i=1}^L \log \int p_\omega(\Lambda_i, \tilde{\rho}_{i-1}, \tilde{\rho}_i) \frac{q(\tilde{\rho}_i | \tilde{\rho}_{i-1})}{q(\tilde{\rho}_i | \tilde{\rho}_{i-1})} d\tilde{\rho}_i \\ &= -\sum_{i=1}^L \log \mathbb{E}_q \frac{p_\omega(\Lambda_i, \tilde{\rho}_{i-1}, \tilde{\rho}_i)}{q(\tilde{\rho}_i | \tilde{\rho}_{i-1})} \\ &\leq \mathbb{E}_q \left[-\sum_{i=1}^L \log \frac{p_\omega(\Lambda_i, \tilde{\rho}_{i-1}, \tilde{\rho}_i)}{q(\tilde{\rho}_i | \tilde{\rho}_{i-1})} \right] := \mathcal{L} \end{aligned} \quad (20)$$

425 **A.4 Proof of Eqn. (11)**

$$\begin{aligned}
\mathcal{L} &= \mathbb{E}_q \left[- \sum_{i=1}^L \log \frac{p_\omega(\Lambda_i, \tilde{\rho}_{i-1}, \tilde{\rho}_i)}{q(\tilde{\rho}_i | \tilde{\rho}_{i-1})} \right] \\
&= \mathbb{E}_q \left[- \sum_{i=1}^L \log \frac{p_\omega(\tilde{\rho}_{i-1} | \Lambda_i, \tilde{\rho}_i)}{q(\tilde{\rho}_{i-1} | \tilde{\rho}_i)} p_\omega(\Lambda_i, \tilde{\rho}_i) \frac{q(\tilde{\rho}_{i-1})}{q(\tilde{\rho}_i)} \right] \\
&= \mathbb{E}_q \left[- \sum_{i=1}^L \log \frac{p_\omega(\tilde{\rho}_{i-1} | \Lambda_i, \tilde{\rho}_i)}{q(\tilde{\rho}_{i-1} | \tilde{\rho}_i)} - \sum_{i=1}^L \log p_\omega(\Lambda_i, \tilde{\rho}_i) - \log \frac{q(\tilde{\rho}_0)}{q(\tilde{\rho}_L)} \right] \\
&= \mathbb{E}_q \left[\sum_{i=1}^L D_{KL} (p_\omega(\tilde{\rho}_{i-1} | \Lambda_i, \tilde{\rho}_i) || q(\tilde{\rho}_{i-1} | \tilde{\rho}_i)) - \sum_{i=1}^L \log p_\omega(\Lambda_i, \tilde{\rho}_i) - \log \frac{q(\tilde{\rho}_0)}{q(\tilde{\rho}_L)} \right]
\end{aligned} \tag{21}$$

426 It should be noted that the fraction $\frac{q(\tilde{\rho}_0)}{q(\tilde{\rho}_L)}$ could be considered as constants. Thus, it should be ignored during the
427 optimization process.

428 **A.5 Parameterized Layer Diagram**

429 Fig. 6 shows the diagram of the actual i -th parameterized layer in the U2 + CNOT design. The learnable
430 parameters $\theta_{i,j:j+1}$ contain two parameters for X-axis and Y-axis rotation. The RX + CNOT and U3 + CNOT
431 designs have similar diagrams.




Figure 6: **The parameterized layer diagram.** The learnable parameters $\theta_{i,j:j+1}$ contain two parameters for X-axis and Y-axis rotation.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: Please refer to the Contributions at the end of the Introduction section

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: Please refer to the Limitations in the Conclusions section.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: Please refer to the Subsection 5.2 and the proof in the Supplementary.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.

- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Please refer to the Subsection 6.1 and Supplementary.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: The experiments are processed with public datasets and the code is not provided due to the institutional policies.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.

- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: Please refer to the Subsection 6.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[Yes\]](#)

Justification: Please refer to the Section 6.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: Please refer to the Subsection 6.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.

- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The authors have reviewed and conformed to the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Please refer to the Broader Impacts in the Conclusions section.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks of misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The assets are cited in the Subsection 6.1.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not research with human subjects.

Guidelines:

- 709 • The answer NA means that the paper does not involve crowdsourcing nor research with human
710 subjects.
- 711 • Depending on the country in which research is conducted, IRB approval (or equivalent) may be
712 required for any human subjects research. If you obtained IRB approval, you should clearly state
713 this in the paper.
- 714 • We recognize that the procedures for this may vary significantly between institutions and
715 locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for
716 their institution.
- 717 • For initial submissions, do not include any information that would break anonymity (if applica-
718 ble), such as the institution conducting the review.