

Distributionally Robust Variational Quantum Algorithms with Shifted Noise

ZICHANG HE¹ BO PENG², YURI ALEXEEV³, AND ZHENG ZHANG¹,

¹Department of Electrical and Computer Engineering, University of California, Santa Barbara, Santa Barbara, CA, 93106, USA

²Physical and Computational Sciences Directorate, Pacific Northwest National Laboratory, Richland, WA, 99352, USA

³Computational Science Division, Argonne National Laboratory, Lemont, IL, 60439, USA

Corresponding author: Zichang He (email: zichanghe@ucsb.edu).

Z.H. and Z.Z. are supported in part by NSF CCF # 1846476. B.P. is supported by the U.S. Department of Energy, Office of Science and National Quantum Information Science Research Centers, DOE Q-NEXT Center (Grant No. DOE 1F-60579). Y.A. is supported in part by funding from the Defense Advanced Research Projects Agency. This work used in part the resources of the Argonne Leadership Computing Facility, which is Department of Energy, Office of Science User Facility supported under Contract DE-AC02-06CH11357. The views, opinions and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense, the Department of Energy, or the U.S. Government.

ABSTRACT Given their potential to demonstrate near-term quantum advantage, variational quantum algorithms (VQAs) have been extensively studied. Although numerous techniques have been developed for VQA parameter optimization, it remains a significant challenge. A practical issue is that quantum noise is highly unstable and thus it is likely to shift in real time. This presents a critical problem as an optimized VQA ansatz may not perform effectively under a different noise environment. For the first time, we explore how to optimize VQA parameters to be robust against unknown shifted noise. We model the noise level as a random variable with an unknown probability density function (PDF), and we assume that the PDF may shift within an uncertainty set. This assumption guides us to formulate a distributionally robust optimization problem, with the goal of finding parameters that maintain effectiveness under shifted noise. We utilize a distributionally robust Bayesian optimization solver for our proposed formulation. This provides numerical evidence in both the Quantum Approximate Optimization Algorithm (QAOA) and the Variational Quantum Eigensolver (VQE) with hardware-efficient ansatz, indicating that we can identify parameters that perform more robustly under shifted noise. We regard this work as the first step towards improving the reliability of VQAs influenced by shifted noise from the parameter optimization perspective.

INDEX TERMS Variational Quantum Algorithms, Distributionally Robust Optimization, Noise Shift, Bayesian Optimization

I. INTRODUCTION

VARIATIONAL quantum algorithms (VQAs) [1] have the potential to demonstrate quantum advantage and have been applied in diverse fields, such as optimization [2], [3], finance [4]–[6], machine learning [7]–[9], quantum simulation [10]–[12], and chemistry [13]–[15]. However, parameter optimization is a substantial challenge for VQAs [16].

Numerous efforts have been made to optimize VQA parameters [17]–[20]. One critical challenge for VQA parameter optimization is quantum noise [21]–[23], which limits their capabilities and introduces additional complexities to parameter optimization. Modeling and mitigating hardware noise is a core part of Near-term Intermediate-scale Quantum (NISQ) algorithms [24]–[26]. Quantifying and improving the reliability and robustness of a VQA has been an important task and has gained increasing attention recently. To name a

few, machine learning methods have been used to estimate the reliability of a quantum circuit [27]; noise-aware ansatz design methodologies [28] and robust circuit realization from a lower-level abstraction [29], [30] have also been investigated.

A more challenging yet practical problem is the instability of quantum noise. Suppose we have an accurate model of the quantum noise as a reference. However, the quantum noise can change significantly under different environmental conditions in real time, making the reference noise model inaccurate. [31] has shown that the noise fluctuation is usually less disinclined. Some studies [32], [33] have considered the reproducibility and stability under different noise models. We refer to this phenomenon of noise change as “noise shift”. The optimization of variational quantum circuits and error mitigation under real-time noise has gained attention

recently [34]–[38].

In this paper, we ask a fundamental question: Can we optimize the VQA parameters such that they are robust to potentially shifted (unknown) noise? We assume that we have access to a fixed noise model, but the actual noise level is an unknown random variable with an unknown PDF. This fixed PDF represents our limited knowledge about the potential noise shift.

To optimize VQA parameters under such unknown noise, for the first time, we propose a new min-max optimization formulation. Such an optimization formulation is called distributionally robust optimization (DRO) in the classical operation research community [39]–[42]. DRO is an advanced optimization framework that aims to find solutions resilient against a range of possible probability distributions rather than a single expected distribution. In our context, we aim to optimize parameters against the worst-case distribution of noise levels. This task, while distinct, complements error mitigation efforts. Rather than attempting to reduce quantum noise, our approach focuses on optimizing parameters in the presence of potentially shifting noise. Furthermore, our method allows for seamless integration with various error mitigation techniques.

Paper contributions. In this work, we investigate the problem formulation, numerical solver, and validation of variational quantum algorithm training under unknown shifted noise. The overview is illustrated in Fig. 1. Our specific contributions include:

- To be robust against the shifted noise, we formulate the problem of optimizing VQA algorithms as a distributionally robust optimization that aims to optimize a targeted performance under the worst-case noise distribution. We characterize the quantum noise using a fixed noise model with uncertain and varying levels of strengths, where the noise level is a random variable with an unknown probability density function (PDF).
- To solve the distributionally robust optimization, we model the unknown PDF as a distributional uncertainty set that is defined by Maximum Mean Discrepancy (MMD). We then solve the min-max problem utilizing a distributionally robust Bayesian optimization (DRBO) method [43]–[45]. Recently, Bayesian optimization (BO) has attracted attention in the field of quantum algorithm optimization [46]–[56].
- We validate the proposed min-max formulation on two well-recognized VQAs, namely QAOA for the MaxCut problem and VQE with hardware-efficient ansatz for the one-dimensional Heisenberg model. Numerical results show that the proposed parameter optimization algorithm outperforms conventional methods under shifted noise conditions.

II. PROBLEM FORMULATION

Variational quantum algorithms (VQAs) are a class of algorithms in quantum computing that utilize a hybrid approach, combining classical and quantum computing resources to

solve computational problems. They are especially pertinent for use with Noisy Intermediate-Scale Quantum (NISQ) devices, which are the currently available quantum hardware.

The core idea of VQAs is to define a parameterized quantum circuit (ansatz) that manipulates the state of a quantum system in a way that depends on a set of parameters θ . These parameters are then optimized classically to minimize an objective function $\langle \psi(\theta) | \mathbf{O} | \psi(\theta) \rangle$, where $\psi(\theta)$ is the resulting noiseless quantum state from the parameterized ansatz, \mathbf{O} is an observable of interest and the objective function is evaluated by the quantum system. However, due to the hardware noise, the actual ansatz and the resulting noisy quantum state \mathbf{P} differ from the ideal ones.

The quantum noise \mathcal{N} for a quantum system with state \mathbf{P} is characterized as

$$\mathcal{N}(\mathbf{P}) = \sum_i \mathbf{E}_i \mathbf{P} \mathbf{E}_i^\dagger, \quad (1)$$

where \mathbf{E}_i are Kraus operators satisfying $\sum_i \mathbf{E}_i \mathbf{E}_i^\dagger = \mathbf{I}$. In this paper, we use amplitude and phase damping channels as the noise model because such a noise model has been shown to shift the optimal parameter of a VQA [57]. The amplitude damping noise describes the energy dissipation of quantum systems, whose Kraus operator formulation on a single qubit is

$$\mathcal{N}(\mathbf{P}) = \mathbf{E}_0 \mathbf{P} \mathbf{E}_0^\dagger + \mathbf{E}_1 \mathbf{P} \mathbf{E}_1^\dagger, \quad (2)$$

with $\mathbf{E}_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-p_{ad}} \end{pmatrix}$, $\mathbf{E}_1 = \begin{pmatrix} 0 & \sqrt{p_{ad}} \\ 0 & 0 \end{pmatrix}$. The phase damping noise describes the quantum information loss without the energy loss, whose Kraus operator formulation on a single qubit is

$$\mathcal{N}(\mathbf{P}) = \mathbf{E}_0 \mathbf{P} \mathbf{E}_0^\dagger + \mathbf{E}_1 \mathbf{P} \mathbf{E}_1^\dagger, \quad (3)$$

with $\mathbf{E}_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-p_{pd}} \end{pmatrix}$, $\mathbf{E}_1 = \begin{pmatrix} 0 & 0 \\ 0 & \sqrt{p_{pd}} \end{pmatrix}$. To integrate these two amplitude and phase damping noise channels, the combined Kraus operator is as follows

$$\mathcal{N}(\mathbf{P}) = \sum_{i=0}^2 \mathbf{E}_i \mathbf{P} \mathbf{E}_i^\dagger, \quad (4)$$

with $\mathbf{E}_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-p_{ad}} \sqrt{1-p_{pd}} \end{pmatrix}$, $\mathbf{E}_1 = \begin{pmatrix} 0 & \sqrt{p_{ad}} \\ 0 & 0 \end{pmatrix}$, $\mathbf{E}_2 = \begin{pmatrix} 0 & 0 \\ 0 & \sqrt{1-p_{ad}} \sqrt{p_{pd}} \end{pmatrix}$. The parameters p_{ad} and p_{pd} are strongly related to the T_1 and T_2 time of quantum hardware. In this paper, we assume an equal probability of two damping channels $p_{ad} = p_{pd} = p$ for simplicity. We do not expect such noise modeling to capture practical hardware noise accurately but only use it for proof of concept.

A. DISTRIBUTIONALLY ROBUST OPTIMIZATION FORMULATION OF VQAS

We assume that we have access to the fixed noise model, i.e., the fixed Kraus presentation (4), but do not know the precise

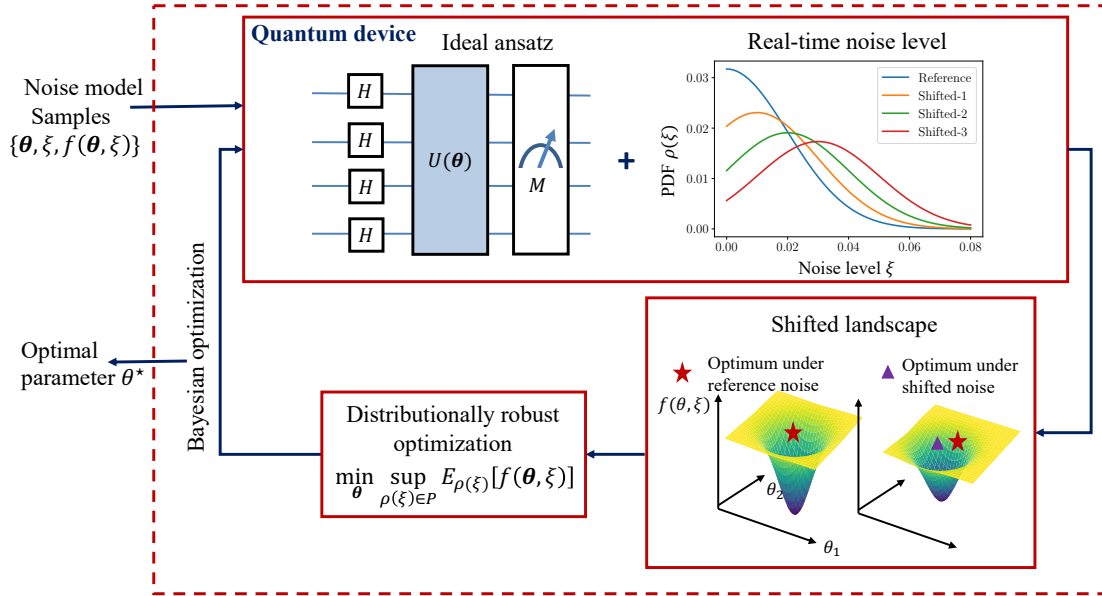


FIGURE 1. Overview of the distributionally robust variational quantum algorithms. Given an ideal ansatz and noise model, we assume the noise level is a random variable that can change in real time. We have samples of the noise level variable ξ from a reference distribution, ansatz parameter θ , and the corresponding VQA performance $f(\theta, \xi)$. With the shifted noise, the VQA landscape and its optimum θ can potentially change. Specifically, the optimal θ under a certain noise level may not perform well under another noise level. Likewise, an optimal θ under a reference noise level PDF may not perform well under another noise level PDF. To address the landscape shift, we reformulate the parameter optimization problem as a min-max formulation to find a robust parameter θ . In other words, we aim to optimize the performance under the worst-case noise level PDF. We use a distributionally robust Bayesian optimization solver to solve the new parameter optimization formulation, which is still handled by classical computers.

noise level p . We model p as a random variable, which follows a certain PDF $\xi \sim \rho(\xi)$. Let $f(\theta, \xi) = \text{Tr}[\mathbf{P}(\theta, \xi)\mathbf{O}]$ be the quantity of interest evaluated by an ansatz parametrized by θ under a noise level ξ , where $\mathbf{P}(\theta, \xi)$ is the resulting noisy quantum state. A standard parameter optimization of a VQA becomes stochastic programming:

$$\min_{\theta} \mathbb{E}_{\rho(\xi)}[f(\theta, \xi)]. \quad (5)$$

Note that we here consider the one-dimensional noise level for simplicity. It can be seamlessly extended to a high-dimensional case.

However, due to the real-time fluctuation of quantum noise, the actual PDF of the noise level can shift and become unknown. As a result, we assume that $\rho(\xi)$ is not exactly known and it can be any PDF inside a set \mathcal{P} , which makes it impossible to obtain a deterministic value of $\mathbb{E}_{\rho(\xi)}[f(\theta, \xi)]$. As a result, we try to optimize the worst-case value of $\mathbb{E}_{\rho(\xi)}[f(\theta, \xi)]$ by solving

$$\min_{\theta} \sup_{\rho(\xi) \in \mathcal{P}} \mathbb{E}_{\rho(\xi)}[f(\theta, \xi)]. \quad (6)$$

When the uncertainty set degenerates to $\mathcal{P} = \{\rho(\xi)\}$, problem (6) degenerates to the standard stochastic optimization problem in (5). On the other hand, the problem degenerates to a robust optimization under the worst noise level when the PDF of a noise level degenerates to a Dirac function.

The distributionally robust circuit optimization (6) may be intractable in practice because (a) \mathcal{P} may contain an infinite number of PDFs describing process variations; (b) the min-max problem is hard to solve by nature; (c) we do not have an analytical form for $f(\theta, \xi)$ under the presence of noise.

III. THE PROPOSED SOLVER

In this section, we properly define the PDF uncertainty set \mathcal{P} and solve problem (6) leveraging distributionally robust Bayesian optimization (DRBO) [43], [45], [58] developed recently in the machine learning community.

A. DISTRIBUTION UNCERTAINTY SET

We model the PDF uncertainty set \mathcal{P} as a ball whose center is the nominal distribution $\rho_0(\xi)$ of the noise level and radius ε is measured by a distribution divergence \mathcal{D} :

$$\mathcal{P} := \mathcal{B}(\rho_0) = \{\rho : \mathcal{D}(\rho_0, \rho) \leq \varepsilon\}. \quad (7)$$

There are many options for the divergence \mathcal{D} , including Maximum Mean Discrepancy, Wasserstein distance, φ -divergence, etc [39]. Here, we choose the Maximum Mean Discrepancy (MMD).

MMD aims to compare the means of samples drawn from two distributions in a high-dimensional reproducing kernel Hilbert space (RKHS) induced by a positive definite kernel function [59]. For the tractability of the problem,

we discretize the noise level in a finite space Ξ with n parts. Then, let \mathcal{H}_M be an RKHS with corresponding kernel $k_M : \Xi \times \Xi \rightarrow \mathbb{R}$, we can embed the distributions ρ_0 (similarly for ρ) into \mathcal{H}_M via the mean embedding:

$$m_{\rho_0} := \mathbb{E}_{\xi \sim \rho_0} [k_M(\xi, \cdot)], \quad \text{such that} \\ \langle m_{\rho_0}, k_M(\xi', \cdot) \rangle = \mathbb{E}_{\xi \sim \rho_0} [k_M(\xi', \xi)], \forall \xi \in \Xi.$$

Then the MMD between two distributions ρ_0 and ρ over Ξ is defined as

$$\mathcal{D}(\rho_0, \rho) := \|m_{\rho_0} - m_{\rho}\|_{\mathcal{H}}, \quad (8)$$

where $\|\cdot\|_{\mathcal{H}} = \sqrt{\langle \cdot, \cdot \rangle}$ is the Hilbert norm. Let $w_i = \rho_0(\xi_i)$ and $w'_i = \rho(\xi_i)$ be the density probability of two discrete distributions, if we replace the expectation with the empirical expectation, i.e., $m_{\rho} = \sum_{i=1}^n w_i k_M(\xi_i, \cdot)$ and $m_{\rho'} = \sum_{i=1}^n w'_i k_M(\xi_i, \cdot)$, Eq. (8) can be written as:

$$\mathcal{D}(\rho_0, \rho) = \sqrt{(\mathbf{w} - \mathbf{w}')^T \mathbf{M} (\mathbf{w} - \mathbf{w}')}, \quad (9)$$

where $M_{ij} = k_M(\xi_i, \xi_j)$ is the kernel matrix.

B. DRO MAIN WORKFLOW

By modeling the distribution uncertainty set defined via MMD, the DRO problem (6) becomes tractable. The main steps are summarized below.

- **Step 1.** Characterize the nominal noise distribution ρ_0 .
- **Step 2.** Given a current θ , solve the inner problem to determine the worst-case PDF of ξ :

$$\sup_{\rho(\xi): \mathcal{D}(\rho_0, \rho) \leq \varepsilon} \mathbb{E}_{\rho(\xi)} [f(\theta, \xi)] = \\ \max_{\substack{\mathbf{w}': \|\mathbf{w}'\|_1 = 1, \\ 0 \leq w'_j \leq 1, \forall j \in [n], \\ \sqrt{(\mathbf{w} - \mathbf{w}')^T \mathbf{M} (\mathbf{w} - \mathbf{w}')} \leq \varepsilon}} \langle \mathbf{w}', f_{\theta} \rangle, \quad (10)$$

where $f_{\theta} := f(\theta, \cdot) \in \mathbb{R}^n$ is the output with a given parameter θ .

- **Step 3.** Solve the outer problem to update θ .

$$\min_{\theta} \langle \mathbf{w}', f_{\theta} \rangle \quad (11)$$

- **Step 4.** If not converge, go back to Step 2.

Specifically, Step 2 can be solved analytically via convex programming as it is second-order cone programming with respect to the worst-case distribution \mathbf{w}' . Step 3 can be solved via a numerical optimizer. However, one of the challenges in steps 2 and 3 is that we need to simulate multiple $f(\theta, \xi)$, which can be expensive in practice. To address the computational issue, we apply a Bayesian optimization solver to the workflow. The key idea is to sequentially learn a surrogate model of $f(\theta, \xi)$ and optimize it by iteratively adding informative samples.

Algorithm 1 Overall DRBO algorithm with GP

Input: Initial sample set $\mathcal{S}_0 = \{\theta^i, \xi^i, f(\theta^i, \xi^i)\}_{i=1}^M$, reference PDF of noise level $\rho_0(\xi)$ with $\rho_0(\xi_i) = w_i, \forall i = [n]$, uncertainty ball radius ε , maximum iteration T

Output: The optimal circuit design variables θ^*

- 1: **for** $t = 1, 2, \dots, T$ **do**
- 2: Construct a GP model as the probabilistic surrogate model $\hat{f}(\theta, \xi) = \mathcal{GP}(\theta, \xi)$ based on \mathcal{S}_{t-1}
- 3: Define $\text{LCB}(\theta, \xi) := \mu(\mathcal{GP}(\theta, \xi)) - \beta \cdot \sigma(\mathcal{GP}(\theta, \xi))$
- 4: Define the PDF of the worst-case distribution $\mathbf{w}' := \arg\max_{\mathbf{w}'} \langle \mathbf{w}', \text{LCB}(\theta, \xi) \rangle$ s.t. $\mathbf{w}' : \|\mathbf{w}'\|_1 = 1, 0 \leq w'_j \leq 1, \forall j \in [n]$, and $\|\mathbf{w}' - \mathbf{w}\|_M \leq \varepsilon$
- 5: Solve the robust parameter $\theta_t = \arg\min_{\theta} \langle \mathbf{w}', \text{LCB}(\theta, \xi) \rangle$
- 6: Sample K noise levels from the reference PDF $\xi_k \sim \rho_0$ and simulate $f(\theta_t, \xi_k)$, for $k = 1, 2, \dots, K$
- 7: Augment data set $\mathcal{S}_t \leftarrow \mathcal{S}_{t-1} \cup \{(\theta_t, \xi_t, f(\theta_t, \xi_k))\}_{k=1}^K$
- 8: **end for**
- 9: Return optimal θ^*

C. BO SOLVER FOR DRO PROBLEM

Next, we explain how to solve DRO via Bayesian optimization with a few quantum circuit simulations. Bayesian optimization sequentially builds a probabilistic surrogate model of $f(\theta, \xi)$ and explores the design space by minimizing an acquisition function. The overall DRBO algorithm is summarized in Algorithm 1.

We first construct a probabilistic surrogate model $\hat{f}(\theta, \xi)$, which can estimate both the output and its uncertainty given an input (θ, ξ) . Here, we use the Gaussian process regression model $\mathcal{GP}(\theta, \xi)$ as the surrogate $\hat{f}(\theta, \xi)$. Then we use its lower confidence bound (LCB) to replace the original objective function $f(\theta, \xi)$ in Eqs. (10) and (11) in steps 2-4

$$f(\theta, \xi) \rightarrow \text{LCB}(\theta, \xi) = \mu(\mathcal{GP}(\theta, \xi)) - \beta \cdot \sigma(\mathcal{GP}(\theta, \xi)), \quad (12)$$

where $\mu(\cdot)$ and $\sigma(\cdot)$ denote the estimated mean and standard deviation, and β is a parameter to balance the model exploitation and exploration.

Gaussian process surrogate. To build the Gaussian process regression (GPR) model, we need to predefine the mean function $m(\cdot)$ and the kernel function $k_{\mathcal{GP}}(\cdot, \cdot)$. Given a dataset $\mathbf{X} = \{\mathbf{x}^i\}_{i=1}^M = \{\theta^i, \xi^i\}_{i=1}^M$ and their simulation outputs $\mathbf{y} = \{f(\mathbf{x}^i) + \epsilon_j\}_{i=1}^M$, the GP model assumes that the simulation outputs follows a Gaussian distribution [60]:

$$\text{Prob}(\mathbf{y}) = \mathcal{N}(\mathbf{y} | \boldsymbol{\mu}, \mathbf{K}), \quad (13)$$

where $\boldsymbol{\mu} \in \mathbb{R}^M$ is the mean vector with $\mu_i = m(\mathbf{x})$, $\mathbf{K} \in \mathbb{R}^{M \times M}$ is the covariance matrix with $K_{i,j} = k_{\mathcal{GP}}(\mathbf{x}_i, \mathbf{x}_j)$. The measurement noise is characterized as a white noise ϵ in the simulation output.

Then, the GP model can offer a probabilistic prediction of a new data \mathbf{x}' , $\mathcal{GP}(\mathbf{x}') \sim N(\mu(\mathbf{x}'), \sigma^2(\mathbf{x}'))$, as follows:

$$\begin{aligned}\mu(\mathbf{x}') &= \mathbf{k}_{\mathcal{GP}}(\mathbf{x}', \mathbf{X})^T (\mathbf{K} + \epsilon^2 \mathbf{I})^{-1} \mathbf{y} \\ \sigma^2(\mathbf{x}') &= k_{\mathcal{GP}}(\mathbf{x}', \mathbf{x}') - \mathbf{k}_{\mathcal{GP}}(\mathbf{x}', \mathbf{X})^T (\mathbf{K} + \epsilon^2 \mathbf{I})^{-1} \mathbf{k}_{\mathcal{GP}}(\mathbf{X}, \theta')\end{aligned}$$

In our cases, we choose the prior mean as $m(\mathbf{x}) = 0$ and use the RBF kernel $k_{\mathcal{GP}}(\mathbf{x}_i, \mathbf{x}_j) = e^{-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2l^2}}$. Note that the kernel $k_{\mathcal{GP}}$ used in GPR differs from the kernel function k_M employed in MMD. Apart from the Gaussian process, there are various other surrogate models that can be utilized.

Optimal selection. Regarding the selection of an optimal solution, it turns out to be non-trivial. Since we want to estimate the expectation over the noise distribution, it will be too expensive to estimate with real quantum devices. Instead, we choose the solution with maximized model posterior, i.e., we choose the $\theta = \operatorname{argmax}_{\theta} \mathbb{E}_{\rho(\xi)} [\mu(\hat{f}(\theta, \xi))] = \langle \mathbf{w}, \mu(\hat{f}_{\theta}) \rangle$, where \hat{f}_{θ} is the mean prediction from \hat{f} with a given parameter θ . It is a common strategy for similar conditions [61], [62].

In addition, an accurate estimation over $\hat{f}(\theta, \cdot)$ will benefit our model output. For this motivation, at the end of each iteration in Algorithm 1, we can add a batch of samples of ξ from Ξ in step 3. It can help build a better probabilistic model and fasten the BO solver convergence.

Remarks. One possible further improvement is to treat the BO solver as a warm-start procedure. After returning a few high-quality solutions from BO, we can conduct the local numerical optimization by taking them as the initial. The local search step may introduce additional computational cost and need more calling of $f(\theta, \xi)$ instead of the surrogate model, but it can lead to a potentially better solution. The hybrid of different solvers is also a common strategy in VQA parameter optimization [49], [63].

The proposed distributionally robust optimization can easily degenerate into stochastic optimization or robust optimization. Stochastic optimization, namely Eq. (5), does not consider the real-time change of the noise. Robust optimization degenerates the PDF of the noise level $\rho(\xi)$ to a single scalar. This case can easily lead to over-conservative parameter optimization.

The current distribution uncertainty set modeling of Maximum Mean Discrepancy has the great power of capturing the worst-case distribution. However, we need to discretize the noise level PDF in order to estimate the MMD efficiently. There exist some other approaches to modeling the uncertainty set that do not discretize the noise level, such as f -divergence modeling [45]. Meanwhile, some alternative uncertainty set modeling could potentially reduce the computational overhead of estimating worst-case distribution (10) [45], [64], [65] iteratively. However, they may not perform well in our experiments due to the unfitted modeling of shifted noise level distribution.

IV. NUMERICAL EXPERIMENTS

We validate the distributionally robust formulation of optimizing VQA parameters (6) in two widely used VQA applications: one is using QAOA for MaxCut and the other one is using VQE for a one-dimensional Heisenberg model.

Here, we conduct the numerical experiments on a simulator in order to adjust the noise level easily and correspondingly to validate the method. To apply the distributional robustness formulation in hardware experiments, the estimation of the noise model and noise level is another challenge, which is out of the scope of this work.

Baselines. We compared the proposed DRBO solver to two standard Bayesian optimization methods for solving stochastic optimization, one is with a lower confidence bound acquisition function (BO-LCB) and the other one is with an expectation improvement acquisition function (BO-EI) [66], and robust Bayesian optimization (BO-Stable) [67].

In BO-LCB, we target problem 5 with a fixed reference distribution of noise level $\rho_0(\xi)$ using a Bayesian optimization approach. We use the same GP surrogate model and its LCB as Eq. (12), but without solving the outer problem (11). Specifically, the lines 4 and 5 of Algorithm 1 are combined as solving $\min_{\theta} \langle \mathbf{w}, \text{LCB}(\theta, \xi) \rangle$.

In BO-EI, we replace the above LCB function with the EI acquisition function. Let θ^- be the best sample with the smallest value $f(\theta^-, \xi)$ so far. The EI acquisition function is defined as:

$$\text{EI}(\theta, \xi) = \Phi(z) (f(\theta^-, \xi) - \mu(\mathcal{GP}(\theta, \xi))) + \phi(z) \sigma(\mathcal{GP}(\theta, \xi)), \quad (14)$$

$$\text{where } z = \begin{cases} \frac{f(\theta^-, \xi) - \mu(\mathcal{GP}(\theta, \xi))}{\sigma(\mathcal{GP}(\theta, \xi))} & \sigma(\mathcal{GP}(\theta, \xi)) > 0 \\ 0 & \sigma(\mathcal{GP}(\theta, \xi)) = 0 \end{cases}, \quad \Phi(\cdot)$$

and $\phi(\cdot)$ are the cumulative distribution function and the probability density function of the standard normal distribution. In the literature of applying BO for learning VQAs, there exists the usage of variant kernel functions and acquisition functions under different noisy environments. However, they can all be categorized as standard BO as they consider a fixed noise model only.

In BO-Stable, we target a shift-aware problem but only focus on the worst noise level instead of the worst distributional noise level. We use the same GP surrogate model and its LCB as Eq. (12). Differing from using DRBO for problem (6), the lines 4 and 5 of Algorithm 1 are replaced with solving $\min_{\theta} \text{LCB}(\theta, \xi^*)$ where given a θ , the worst ξ^* is defined as $\xi^* := \operatorname{argmax}_{\xi} \text{LCB}(\theta, \xi)$.

To compare these different VQA parameter optimization methods, we first obtain their different optimized parameters and evaluate them under different levels of noise $\rho(\xi)$. In the simulation, we used the qiskit_aer noisy simulator with the statevector backend, whose simulation algorithm is the Monte Carlo trajectory approach. In one circuit trajectory, a Kraus operator is randomly applied on an ideal gate with probability that is defined by the noise channel. The multi-qubit noise channel is defined as the tensor product of single-qubit noise ones (4).

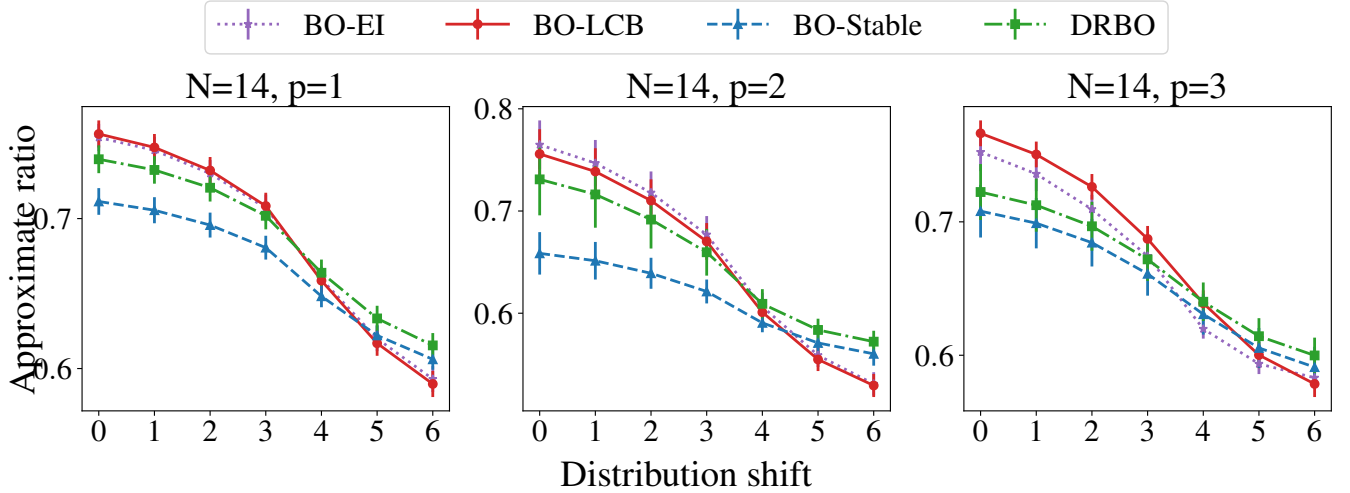


FIGURE 2. The results for solving $N = 14$ 3-regular graph MaxCut problems via QAOA. The x axis denotes the significance of noise shift, where the noise level PDF is the reference one at $x = 0$. The y axis is the expectation of the approximation ratio of the QAOA solution evaluated at different noise PDFs. We report the average result over 10 non-isomorphic graphs. As we can see, the standard BO-LCB and BO-EI solutions have the best performance under the reference noise. However, under an increasingly shifted noise, the DRBO solution begins to outperform the standard BO solutions. Meanwhile, the BO-Stable solution is over-conservative with respect to the noise. It significantly sacrifices the performance under the reference PDF and the slightly shifted PDFs to gain an improvement under significant shifts. These observations are consistent in the experiments with different QAOA depths.

A. NUMERICAL EXPERIMENTS ON QAOA

QAOA is a leading variational quantum algorithm for combinatorial optimization problems. It alternatively applies two operators, a phase-separation operator and a mixer operator, to drive a quantum system to the target solution state. A noiseless QAOA solution is denoted as $\psi(\theta) = e^{-i\beta_p \mathbf{H}_M} e^{-i\gamma_p \mathbf{H}_P} \dots e^{-i\beta_1 \mathbf{H}_M} e^{-i\gamma_1 \mathbf{H}_P} |\psi_0\rangle$.

We will take the MaxCut problem as a case study of QAOA. Given a graph $G = (V, E)$ with vertices V and edges E , the MaxCut problem aims to find a cut that partitions the graph vertices into two sets with the largest number of edges. Its cost function is written as

$$C = \sum_{(i,j) \in E} 1 - s_i s_j, \quad (15)$$

where s_i and s_j are binary variables associated to the vertices in V , which assume value 1 or -1 depending on which of the two partitions defined by the cut are assigned. Its cost Hamiltonian is defined as $\mathbf{H}_C = \sum_{(u,v) \in E} \frac{1}{2} (\mathbf{I} - \mathbf{Z}_u \mathbf{Z}_v)$, where \mathbf{Z}_i denotes a Pauli-Z operator.

In applying QAOA for solving the MaxCut problem, given a noise model with noise level ξ , we aim to optimize the QAOA parameters $\theta = (\gamma, \beta)$ such that the resulting quantum state $\mathbf{P}(\theta, \xi)$ has minimal energy $f(\theta, \xi) = \text{Tr}[\mathbf{P}(\theta, \xi) \mathbf{H}_C]$. Considering the uncertainty and fluctuation of noise level, we aim to find parameters θ that make QAOA performance more robust towards the shifted noise by solving the DRO problem (6).

Here, we discretize the noise level into 20 bins in $[0, 0.08]$ evenly. We assume the reference noise follows a truncated Gaussian distribution, with the noise fluctuation shifting its mean to a larger value. We first generate a truncated Gaussian

distribution with mean -0.01 and standard deviation 0.01 . We estimate the probability density at each discretized level and do the normalization to obtain the reference PDF of the noise level. We follow a similar procedure to generate the PDF of shifted noise by shifting the mean of the initial truncated Gaussian distribution.

To begin with, for a depth- p QAOA ansatz, we initialize the sampling set $\mathcal{S}_0 = \{\theta^i, \xi^i, f(\theta^i, \xi^i)\}_{i=1}^M$ by taking $M = 20p$, where θ^i is drawn from the design space based on a Latin hypercube approach [68], and the noise level samples ξ^i are drawn from the reference distribution $\rho_0(\xi)$. We set the maximum BO iterations as $T = 20p$.

As shown in Fig. 2, we evaluate different BO-based parameter optimization results on 10 graphs with degree-3 and graph size $N = 14$. We report the average approximation ratio results under different shifted noise levels. The x -axis denotes the index of the levels of noise shift, with a higher one denoting a more significant shift, and index-0 denotes the reference noise. Since we solve the optimal θ under shifted noise, the DRBO-solved QAOA is expected to perform worse than the one solved from a standard BO solver under the reference noise. However, as the noise shift becomes more and more significant, the DRBO solution begins to show its advantages. Notably, BO-Stable performs better than BO-LCB and BO-EI under a significantly shifted noise as well. However, it is also over-conservative under the reference noise since it only considers a single worst noise level. The results and observations are consistent over different QAOA depths.

We plot the solution during the BO iterations in Fig. 3. During the iteration, the performance is evaluated under the optimal solution selected from the maximum posterior rather




FIGURE 3. One example of evolving of solution in different BO algorithms. The x axis is the iterations in a BO algorithm, and the y axis is the expectation of cost function evaluated over noise level at a θ . The evaluated θ at one iteration is obtained by maximizing the model posterior, which is unnecessary to be the explored θ at that iteration. Under the reference noise PDF, the BO-LCB algorithm converges to a better solution, while the DRBO converges to a better solution under the shifted noise. The rightmost figure shows the example PDFs of the reference noise level, shifted noise level, and the estimated worst-case noise level from the DRBO algorithm.




FIGURE 4. An example of explored θ in $p = 1$ noisy QAOA cost landscape in a $N = 8$ MaxCut problem. The optimum θ differs from the noiseless optimum. Compared to the BO-LCB and BO-Stable, the DRBO explores the parameter space that performs well under the shifted noise.

than from the solution of an acquisition function. We show the performance evaluated under the reference noise and the shifted noise. During the iterations, DRBO consistently converges to a shifted-noise preferred solution while the LCB converges to a reference-noise preferred one. We also show the PDFs of the reference noise, shifted noise, and the worst-case that is estimated by the DRBO algorithm. We can see that the MMD approach successfully captures the shifted noise under the worse-case distribution, enabling the DRBO to explore the parameters space that performs better under shifted noise.

More results on MaxCut with graph size $N = 8, 10, 12$ and QAOA depth $p = 1, 2, 3$ are shown in Fig. 5. The results are consistent with the ones in Fig. 2. The DRBO solution

performs better than the baselines under significantly shifted noise, which demonstrates that our method could optimize the VQA parameters that are more robust to the shifted noise. One example of the newly sampled θ with $p = 1$ is plotted in Fig. 4. The DRBO algorithm explores the parameter space toward the optimal one under a shifted noise, while the other algorithms exploit the space surrounding the optimal parameter under the reference noise. Therefore, the DRBO could find a parameter that performs better under shifted noise.

B. NUMERICAL EXPERIMENTS ON VQE

Variational quantum eigensolver (VQE) is another popular VQA, specifically designed to simulate quantum systems and find the ground state energy of quantum systems. We use the VQE algorithm with a hardware-efficient ansatz [69] for simulating the ground energy of a one-dimensional Heisenberg model defined as:

$$\mathbf{H} = J \sum_i \mathbf{X}_i \mathbf{X}_{i+1} + \mathbf{Y}_i \mathbf{Y}_{i+1} + \mathbf{Z}_i \mathbf{Z}_{i+1} + B \sum_i \mathbf{Z}_i, \quad (16)$$

where J is the strength of the spin-spin interaction and B is the magnetic field along the Z direction. Here, we use a hardware-efficient ansatz to implement the VQE algorithm (see more details in the supplementary). Given an ansatz parameterized by θ and under a noise model with noise level ξ , we denote the state as $\mathbf{P}(\theta, \xi)$. The noisy VQE algorithm cost function is defined as $f(\theta, \xi) = \text{Tr}[\mathbf{P}(\theta, \xi)\mathbf{H}]$. Here, we aim to find the ansatz parameters that lead to robust performance under shifted noise by solving the DRO problem (6).

Here, for a two-layer hardware-efficient ansatz, we discretize the noise level into 20 bins in $[0, 0.08]$ evenly. We assume the reference noise follows a truncated Gaussian distribution, and the changing noise shifts its mean to a larger

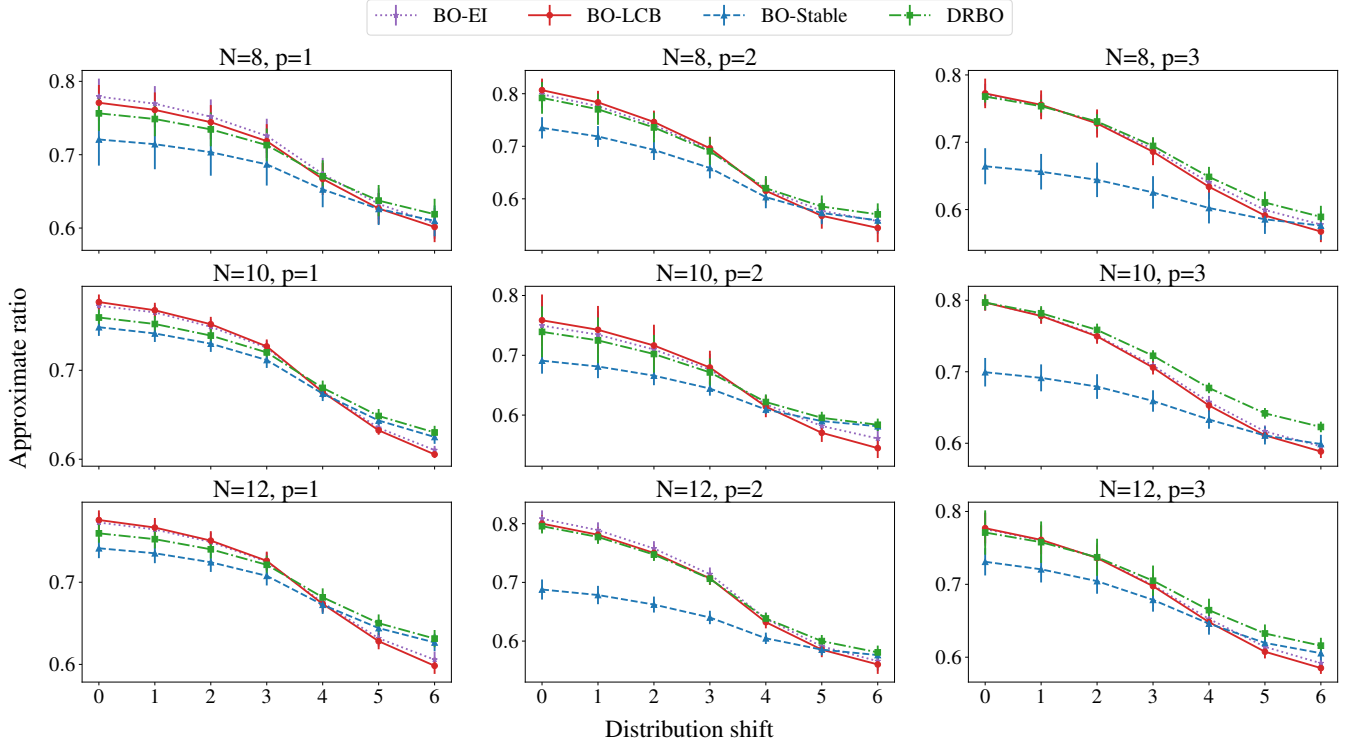


FIGURE 5. More results on the Max-Cut experiments with graph sizes $N = 8, 10, 12$ and QAOA depth $p = 1, 2, 3$. While potentially sacrificing the performance under the reference noise a little, the DRBO solution performs better than standard BO methods under the significantly shifted noise. Meanwhile, BO-Stable solutions are over-conservative.

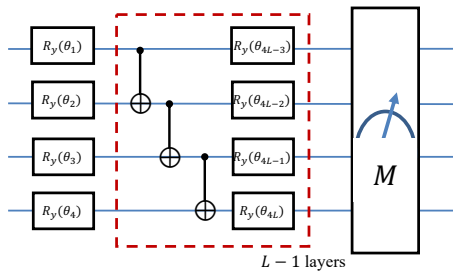


FIGURE 6. The schematic of the hardware-efficient ansatz for variational quantum eigensolver. The layer of two-qubit entanglement and one-qubit rotation gates are repeated for $L - 1$ times.

value. The hardware-efficiency ansatz is set up as shown in Fig. 6. The number of parameters grows quickly and becomes challenging for a BO solver. For simplicity, we only optimize the last N parameters and fix the others, similar to the idea of layer-wise optimization in Ref. [70]. For the demonstration purpose, the fixed parameters are obtained through a multi-start classical optimization routine. We follow the same procedure as QAOA to set up the noise level distribution of both the reference and the shifted ones.

In Fig. 7, we show the ground energy solved for a 6-spin system with $J = 1, B = 0.2$, whose ground state is highly entangled. Aiming to optimize the last layer parameters, we initialize the sampling set with $M = 40$ and set the




FIGURE 7. The results for solving the ground energy of a 6-spin, $J = 1, B = 0.2$ one-dimensional Heisenberg model via VQE with two-layer hardware efficient ansatz. The x axis denotes the significance of noise shift, where the noise level PDF is the reference one at $x = 0$. We first obtained the optimal parameter θ_0 in a noiseless simulation, which solves the problem perfectly with $f(\theta_0) = -4.8$. Then we report the relative improvement of the energy $\frac{\mathbb{E}_{\rho(\xi)}[f(\theta, \xi)] - \mathbb{E}_{\rho(\xi)}[f(\theta_0, \xi)]}{\mathbb{E}_{\rho(\xi)}[f(\theta_0, \xi)]}$. Under all the shifted distributions, BO-LCB performs close to θ_0 . DRBO sacrifices limited performance under mild noise and performs much better than the BO-LCB, BO-EI, and noiseless optimal θ_0 in significantly shifted noise. BO-LCB and BO-EI are almost overlapping since they both have solutions close to θ_0 . While the BO-Stable can also find the robust parameter under the shifted noise, it does not perform as well as DRBO, especially when the noise shift is mild.

maximum BO iteration $T = 40$ as well. We reported the results from a two-layer hardware-efficiency ansatz with a more detailed setup and results in Supplementary. Similar to the QAOA results, DRBO performs better than BO-LCB and BO-EI under a significantly shifted noise. Furthermore, the




FIGURE 8. Example of energy landscapes under different noise models. The left heatmap is the depth-1 QAOA landscape for the MaxCut problem. The color denotes the solved energy. The optimal point in this landscape is highlighted as a triangle. The middle heatmap is the landscape under simple Pauli errors, which has been shown not to change the VQA optimal and uniformly flatten the landscape. The right heatmap is the landscape under the phase and amplitude damping noise, where the optimum is shifted and the energy landscape has a different shape. Under the noise, both the middle and the right ones have worse energies than the left noise-free landscape.

DRBO solution minimally compromises the reference noise performance, whereas Bo-Stable tends to produce overly conservative results.

V. DISCUSSION

A. RELATED WORKS OF BAYESIAN OPTIMIZATION

Bayesian optimization has been a prominent technique in addressing VQA learning tasks. For instance, [48] implemented a parallel optimization scheme, enhancing the efficiency of optimizing VQA parameters across multiple similar problems by leveraging information sharing within Bayesian optimization. [47] used Bayesian inference to reduce the redundancy of the parameterized circuit, ending with a shallower and more robust VQA ansatz. [46] used the standard Bayesian optimization techniques to optimize parameters for noisy VQE. [49] proposed a hybrid approach, utilizing Bayesian optimization solutions as warm starts and transitioning to multi-start local search from optimized Bayesian samples. [50] specifically applied Bayesian optimization to optimize QAOA parameters. [51] demonstrated the efficiency of Bayesian optimization in optimizing parameterized quantum annealing schedules. [52] introduced stochastic gradient line Bayesian optimization, leveraging Bayesian optimization to adjust step sizes in stochastic gradient descent, thereby reducing measurement-shot costs in optimizing VQA parameters. [53] extended standard QAOA with two mixers and optimized circuit parameters through Bayesian optimization. [54] utilized a VQE-kernel to construct a Gaussian process model, leveraging specific circuit properties as physics-informed priors and introducing a novel acquisition function to exploit the inductive bias of the kernel. [55] aimed to initialize good ansatz by fully exploring the Clifford parameter space through Bayesian optimization, where all simulations can be performed classically. [56] introduced a novel Bayesian optimization approach, optimizing QAOA parameters by constructing a surrogate model with constraints derived from two adaptive regions.

Among all the above literature, it is hard to justify the best Bayesian optimization techniques since they use different surrogate modeling techniques, acquisition functions under

different noise models, and target different applications. However, they can be all viewed as standard BO setups as they all aim to optimize either the parameter or ansatz architecture of VQA under a fixed noise environment. None of them shared the same shift-aware problem setup as us.

B. LANDSCAPE SHIFT

Ref. [57] has shown that the optimal variational parameters are unaffected by a broad class of noise models, such as measurement noise, gate noise, and Pauli channel noise. This phenomenon is called optimal parameter resilience. Meanwhile, some noise can shift the location of minima. A rich of work has studied how quantum noise can influence the VQA landscape [71]–[73]. We highlight that the shift location of optimal parameters motivates our work, i.e., given that the optimal parameter will change under different (shifted) noise, we aim to find optimal parameters with robust performance under the shifted noise environment.

In our simulation, we use the phase and amplitude damping noise model, which has been shown to change the values of optimal parameters. The landscape with a changed or unchanged optimum is illustrated in Fig. 8.

C. VARIANT PROBLEM FORMULATION

Radius varying formulation. Beyond optimization under varying noise, another case where the distributionally robust optimization can be applied is to calibrate the noise estimation. Assume we do not have a precise enough estimation of the real noise level distribution as the reference distribution. The key idea is that as we collect more data on the noise level, we can have a more accurate estimation of its PDF. Therefore, as the iteration continues, we can gradually refine the center $\rho_0(\xi)$ of the uncertainty ball and reduce its radius ε .

Gate error modeling. Beyond modeling hardware noise, another possible modeling is on the gate error $f(\theta + \xi)$ of a parameterized quantum circuit, which assumes that the gate parameters are not exactly implemented but suffer from some coherent errors. In such a formulation, under different error levels of ξ , the optimal θ will clearly have different values. The DRO formulation can optimize the VQA to find parameters that are robust to the shifted gate errors. A robustness analysis of such a formulation is discussed in [74]. It is also of great interest to connect the gate error with a more detailed control error and apply the proposed shift-aware optimization at the physical level, but it is out of the scope of this manuscript.

VI. CONCLUSION AND OPPORTUNITIES

Quantum noise has been a major obstacle to the practical applications of near-term quantum computers, particularly in variational quantum algorithms (VQAs). Despite various error mitigation techniques have been intensively studied, the dynamic nature of quantum noise presents a formidable challenge. Optimized VQA parameters may perform suboptimally when exposed to different noise environments.

In this paper, we have presented a distributionally robust optimization formulation designed to enhance the robustness of VQA parameters against varying quantum noise conditions. Our approach leverages a distributionally robust Bayesian optimization (DRBO) solver, efficiently tackling the proposed formulation. We validate the proposed method in two widely recognized VQA benchmarks: QAOA for MaxCut and VQE with hardware-efficient ansatz for an one-dimensional Heisenberg model. The proposed distributionally robust optimization formulation does not aim to mitigate inherent quantum noise directly. Instead, it addresses the noise at the algorithmic level. It can be potentially integrated with various error mitigation techniques to further improve the VQA robustness.

Our formulation can be more impactful in scenarios involving large-size problems. As we scale up to larger problems or deeper VQA ansatz, even minor fluctuations in noise levels can significantly affect performance. For instance, implementing QAOA for larger problems necessitates deeper circuit depths, amplifying the influence of even slight shifts in noise levels on VQA performance. Therefore, optimization VQA parameters under such shifts becomes increasingly critical.

To integrate the proposed distributionally robust formulation into more practical use cases, a better knowledge of noise models is highly desired. More specifically, in this paper, we characterize quantum noise by a fixed noise model, such as Eq. (4), and model the noise level as a random variable with an unknown PDF. Such mathematical modeling may not capture the actual hardware noise well since the noise characterization is in principle challenging. Another point is that we also need to model the uncertainty ball carefully. To make sure the optimized parameters perform well under the actual shifted noise, we need to tune the radius of uncertainty ball ε such that the shifted noise is lying with the ball and is close to the worst case within the ball. Otherwise, the proposed DRBO may become over-conservative as shown in Figs. 2 5 7.

To improve the solver of distributionally robust optimization, some better techniques that do not need to discretize the noise level or efficiently handle high-dimensional parameter optimization of θ can be developed.

We also have applied a similar distributionally robust optimization formulation for classical circuit optimization [75], where we identified the shifts of process variations. In this paper, we focus on handling the parameter optimization of noisy variation quantum algorithms. Differing from [75], the noise source and modeling in this paper are distinct, characterized by a fixed noise model with different levels of strength. We used MMD to define the uncertainty ball such that the shift-aware problem is solved in a two-step process. Additionally, we study its quantum applications from the energy landscape perspective, which is much less studied in classical circuit applications. There is a great thread studying the energy landscape of VQAs like Refs [76]–[78]. We believe our shift-aware optimization is interesting and could

be inspiring to the community.

REFERENCES

- [1] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio et al., “Variational quantum algorithms,” *Nature Reviews Physics*, vol. 3, no. 9, pp. 625–644, 2021.
- [2] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn, A. Kandala, A. Mezzacapo, P. Müller, W. Riess, G. Salis, J. Smolin, I. Tavernelli, and K. Temme, “Quantum optimization using variational algorithms on near-term quantum devices,” *Quantum Science and Technology*, vol. 3, no. 3, p. 030503, jun 2018.
- [3] D. J. Egger, J. Mareček, and S. Woerner, “Warm-starting quantum optimization,” *Quantum*, vol. 5, p. 479, 2021.
- [4] Z. He, R. Shaydulin, S. Chakrabarti, D. Herman, C. Li, Y. Sun, and M. Pistoia, “Alignment between initial state and mixer improves qaoa performance for constrained optimization,” *npj Quantum Information*, vol. 9, no. 1, p. 121, 2023.
- [5] R. Orús, S. Mugel, and E. Lizaso, “Quantum computing for finance: Overview and prospects,” *Reviews in Physics*, vol. 4, p. 100028, 2019.
- [6] D. Herman, C. Googin, X. Liu, A. Galda, I. Safro, Y. Sun, M. Pistoia, and Y. Alexeev, “A survey of quantum computing for finance,” arXiv preprint arXiv:2201.02773, 2022.
- [7] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, “Quantum machine learning,” *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.
- [8] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, and P. J. Coles, “Challenges and opportunities in quantum machine learning,” *Nature Computational Science*, vol. 2, no. 9, pp. 567–576, 2022.
- [9] J. Liu, M. Liu, J.-P. Liu, Z. Ye, Y. Alexeev, J. Eisert, and L. Jiang, “Towards provably efficient quantum algorithms for large-scale machine-learning models,” arXiv preprint arXiv:2303.03428, 2023.
- [10] A. Miessen, P. J. Ollitrault, F. Tacchino, and I. Tavernelli, “Quantum algorithms for quantum dynamics,” *Nature Computational Science*, vol. 3, no. 1, pp. 25–37, 2023.
- [11] B. Peng, S. Gulania, Y. Alexeev, and N. Govind, “Quantum time dynamics employing the Yang-Baxter equation for circuit compression,” *Physical Review A*, vol. 106, no. 1, p. 012412, 2022.
- [12] S. Gulania, Z. He, B. Peng, N. Govind, and Y. Alexeev, “QuYBE-an algebraic compiler for quantum circuit compression,” in *2022 IEEE/ACM 7th Symposium on Edge Computing (SEC)*. IEEE, 2022, pp. 406–410.
- [13] D. A. Fedorov, B. Peng, N. Govind, and Y. Alexeev, “VQE method: a short survey and recent developments,” *Materials Theory*, vol. 6, no. 1, pp. 1–21, 2022.
- [14] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, “Quantum computational chemistry,” *Reviews of Modern Physics*, vol. 92, no. 1, p. 015003, 2020.
- [15] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. Sawaya et al., “Quantum chemistry in the age of quantum computing,” *Chemical reviews*, vol. 119, no. 19, pp. 10 856–10 915, 2019.
- [16] L. Bittel and M. Kliesch, “Training variational quantum algorithms is NP-hard,” *Physical review letters*, vol. 127, no. 12, p. 120502, 2021.
- [17] K. J. Sung, J. Yao, M. P. Harrigan, N. C. Rubin, Z. Jiang, L. Lin, R. Babbush, and J. R. McClean, “Using models to improve optimizers for variational quantum algorithms,” *Quantum Science and Technology*, vol. 5, no. 4, p. 044008, 2020.
- [18] X. Bonet-Monroig, H. Wang, D. Vermetten, B. Senjean, C. Moussa, T. Bäck, V. Dunjko, and T. E. O’Brien, “Performance comparison of optimization methods on variational quantum algorithms,” *Physical Review A*, vol. 107, no. 3, p. 032407, 2023.
- [19] A. Gilyén, S. Arunachalam, and N. Wiebe, “Optimizing quantum optimization algorithms via faster quantum gradient computation,” in *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2019, pp. 1425–1444.
- [20] R. Shaydulin, P. C. Lotshaw, J. Larson, J. Ostrowski, and T. S. Humble, “Parameter transfer for quantum approximate optimization of weighted maxcut,” *ACM Transactions on Quantum Computing*, vol. 4, no. 3, pp. 1–15, 2023.
- [21] D. Stilck França and R. Garcia-Patron, “Limitations of optimization algorithms on noisy quantum devices,” *Nature Physics*, vol. 17, no. 11, pp. 1221–1227, 2021.

- [22] G. González-García, R. Trivedi, and J. I. Cirac, “Error propagation in NISQ devices for solving classical optimization problems,” *PRX Quantum*, vol. 3, no. 4, p. 040326, 2022.
- [23] G. De Palma, M. Marvian, C. Rouzé, and D. S. França, “Limitations of variational quantum algorithms: a quantum optimal transport approach,” *PRX Quantum*, vol. 4, no. 1, p. 010309, 2023.
- [24] R. Harper, S. T. Flammia, and J. J. Wallman, “Efficient learning of quantum noise,” *Nature Physics*, vol. 16, no. 12, pp. 1184–1188, 2020.
- [25] S. Endo, S. C. Benjamin, and Y. Li, “Practical quantum error mitigation for near-future applications,” *Physical Review X*, vol. 8, no. 3, p. 031027, 2018.
- [26] Y. Suzuki, S. Endo, K. Fujii, and Y. Tokunaga, “Quantum error mitigation as a universal error reduction technique: applications from the NISQ to the fault-tolerant quantum computing eras,” *PRX Quantum*, vol. 3, no. 1, p. 010345, 2022.
- [27] J. Liu and H. Zhou, “Reliability modeling of NISQ-era quantum computers,” in 2020 IEEE international symposium on workload characterization (IISWC). IEEE, 2020, pp. 94–105.
- [28] H. Wang, Y. Ding, J. Gu, Y. Lin, D. Z. Pan, F. T. Chong, and S. Han, “QuantumNAS: Noise-adaptive search for robust quantum circuits,” in 2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA). IEEE, 2022, pp. 692–708.
- [29] A. B. Magann, C. Arenz, M. D. Grace, T.-S. Ho, R. L. Kosut, J. R. McClean, H. A. Rabitz, and M. Sarovar, “From pulses to circuits and back again: A quantum optimal control perspective on variational quantum algorithms,” *PRX Quantum*, vol. 2, no. 1, p. 010101, 2021.
- [30] Z. Liang, Z. Song, J. Cheng, Z. He, J. Liu, H. Wang, R. Qin, Y. Wang, S. Han, X. Qian et al., “Hybrid gate-pulse model for variational quantum algorithms,” arXiv preprint arXiv:2212.00661, 2022.
- [31] J. J. Burnett, A. Bengtsson, M. Scigliuzzo, D. Niepce, M. Kudra, P. Delsing, and J. Bylander, “Decoherence benchmarking of superconducting qubits,” *npj Quantum Information*, vol. 5, no. 1, p. 54, 2019.
- [32] S. Dasgupta and T. S. Humble, “Assessing the stability of noisy quantum computation,” in *Quantum Communications and Quantum Imaging XX*, vol. 12238. SPIE, 2022, pp. 44–49.
- [33] —, “Characterizing the reproducibility of noisy quantum circuits,” *Entropy*, vol. 24, no. 2, p. 244, 2022.
- [34] Z. Hu, R. Wolle, M. Tian, Q. Guan, T. Humble, and W. Jiang, “Toward consistent high-fidelity quantum learning on unstable devices via efficient in-situ calibration,” in 2023 IEEE International Conference on Quantum Computing and Engineering (QCE), vol. 1. IEEE, 2023, pp. 848–858.
- [35] J. Zhang, H. Wang, G. S. Ravi, F. T. Chong, S. Han, F. Mueller, and Y. Chen, “Disq: Dynamic iteration skipping for variational quantum algorithms,” in 2023 IEEE International Conference on Quantum Computing and Engineering (QCE), vol. 1. IEEE, 2023, pp. 1062–1073.
- [36] S. Dasgupta and T. S. Humble, “Reliable devices yield stable quantum computations,” arXiv preprint arXiv:2307.05381, 2023.
- [37] S. Dasgupta, T. S. Humble, and A. Danageozian, “Adaptive mitigation of time-varying quantum noise,” in 2023 IEEE International Conference on Quantum Computing and Engineering (QCE), vol. 1. IEEE, 2023, pp. 99–110.
- [38] G. S. Ravi, K. Smith, J. M. Baker, T. Kannan, N. Earnest, A. Javadi-Abhari, H. Hoffmann, and F. T. Chong, “Navigating the dynamic noise landscape of variational quantum algorithms with QISMET,” in *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2, 2023*, pp. 515–529.
- [39] H. Rahimian and S. Mehrotra, “Distributionally robust optimization: A review,” arXiv preprint arXiv:1908.05659, 2019.
- [40] F. Lin, X. Fang, and Z. Gao, “Distributionally robust optimization: A review on theory and applications,” *Numerical Algebra, Control and Optimization*, vol. 12, no. 1, pp. 159–212, 2022.
- [41] D. Kuhn, P. M. Esfahani, V. A. Nguyen, and S. Shafieezadeh-Abadeh, “Wasserstein distributionally robust optimization: Theory and applications in machine learning,” in *Operations research & management science in the age of analytics*. Informs, 2019, pp. 130–166.
- [42] E. Delage and Y. Ye, “Distributionally robust optimization under moment uncertainty with application to data-driven problems,” *Operations research*, vol. 58, no. 3, pp. 595–612, 2010.
- [43] J. Kirschner, I. Bogunovic, S. Jegelka, and A. Krause, “Distributionally robust Bayesian optimization,” in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 2174–2184.
- [44] S. S. Tay, C. S. Foo, U. Daisuke, R. Leong, and B. K. H. Low, “Efficient distributionally robust Bayesian optimization with worst-case sensitivity,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 21180–21204.
- [45] H. Husain, V. Nguyen, and A. v. d. Hengel, “Distributionally robust Bayesian optimization with ϕ -divergences,” arXiv preprint arXiv:2203.02128, 2022.
- [46] G. Iannelli and K. Jansen, “Noisy Bayesian optimization for variational quantum eigensolvers,” arXiv preprint arXiv:2112.00426, 2021.
- [47] S. Duffield, M. Benedetti, and M. Rosenkranz, “Bayesian learning of parameterised quantum circuits,” *Machine Learning: Science and Technology*, vol. 4, no. 2, p. 025007, 2023.
- [48] C. N. Self, K. E. Khosla, A. W. Smith, F. Sauvage, P. D. Haynes, J. Knolle, F. Mintert, and M. Kim, “Variational quantum algorithm with information sharing,” *npj Quantum Information*, vol. 7, no. 1, pp. 1–7, 2021.
- [49] J. Müller, W. Lavrijsen, C. Iancu, and W. de Jong, “Accelerating noisy VQE optimization with Gaussian processes,” in 2022 IEEE International Conference on Quantum Computing and Engineering (QCE). IEEE, 2022, pp. 215–225.
- [50] S. Tibaldi, D. Vodola, E. Tignone, and E. Ercolessi, “Bayesian optimization for QAOA,” *IEEE Transactions on Quantum Engineering*, 2023.
- [51] J. R. Finžgar, M. J. Schuetz, J. K. Brubaker, H. Nishimori, and H. G. Katzgraber, “Designing quantum annealing schedules using Bayesian optimization,” *Physical Review Research*, vol. 6, no. 2, p. 023063, 2024.
- [52] S. Tamiya and H. Yamasaki, “Stochastic gradient line Bayesian optimization for efficient noise-robust optimization of parameterized quantum circuits,” *npj Quantum Information*, vol. 8, no. 1, p. 90, 2022.
- [53] J. E. Kim and Y. Wang, “Quantum approximate Bayesian optimization algorithms with two mixers and uncertainty quantification,” *IEEE Transactions on Quantum Engineering*, 2023.
- [54] A. Benítez-Buenache and Q. Portell-Montserrat, “Bayesian parameterized quantum circuit optimization (BPQCO): A task and hardware-dependent approach,” arXiv preprint arXiv:2404.11253, 2024.
- [55] G. S. Ravi, P. Gokhale, Y. Ding, W. Kirby, K. Smith, J. M. Baker, P. J. Love, H. Hoffmann, K. R. Brown, and F. T. Chong, “CAFQA: A classical simulation bootstrap for variational quantum algorithms,” in *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1, 2022*, pp. 15–29.
- [56] L. Cheng, Y.-Q. Chen, S.-X. Zhang, and S. Zhang, “Quantum approximate optimization via learning-based adaptive optimization,” *Communications Physics*, vol. 7, no. 1, p. 83, 2024.
- [57] K. Sharma, S. Khatri, M. Cerezo, and P. J. Coles, “Noise resilience of variational quantum compiling,” *New Journal of Physics*, vol. 22, no. 4, p. 043006, 2020.
- [58] T. Nguyen, S. Gupta, H. Ha, S. Rana, and S. Venkatesh, “Distributionally robust Bayesian quadrature optimization,” in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 1921–1931.
- [59] K. Muandet, K. Fukumizu, B. Sriperumbudur, B. Schölkopf et al., “Kernel mean embedding of distributions: A review and beyond,” *Foundations and Trends® in Machine Learning*, vol. 10, no. 1-2, pp. 1–141, 2017.
- [60] C. E. Rasmussen, “Gaussian processes in machine learning,” in *Summer school on machine learning*. Springer, 2003, pp. 63–71.
- [61] P. I. Frazier, “A tutorial on Bayesian optimization,” arXiv preprint arXiv:1807.02811, 2018.
- [62] S. Cakmak, R. Astudillo Marban, P. Frazier, and E. Zhou, “Bayesian optimization of risk measures,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 20130–20141, 2020.
- [63] W. Lavrijsen, A. Tudor, J. Müller, C. Iancu, and W. De Jong, “Classical optimizers for noisy intermediate-scale quantum devices,” in 2020 IEEE international conference on quantum computing and engineering (QCE). IEEE, 2020, pp. 267–277.
- [64] F. Farokhi, “Distributionally-robust optimization with noisy data for discrete uncertainties using total variation distance,” *IEEE Control Systems Letters*, 2023.
- [65] P. Mohajerin Esfahani and D. Kuhn, “Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations,” *Mathematical Programming*, vol. 171, no. 1-2, pp. 115–166, 2018.
- [66] N. Srinivas, A. Krause, S. Kakade, and M. Seeger, “Gaussian process optimization in the bandit setting: No regret and experimental design,” in *International Conference on Machine Learning*. Madison, WI, USA: Omnipress, 2010, p. 1015–1022.
- [67] I. Bogunovic, J. Scarlett, S. Jegelka, and V. Cevher, “Adversarially robust optimization with Gaussian processes,” *Advances in neural information processing systems*, vol. 31, 2018.

- [68] K. Q. Ye, "Orthogonal column Latin hypercubes and their application in computer experiments," *Journal of the American Statistical Association*, vol. 93, no. 444, pp. 1430–1439, 1998.
- [69] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *nature*, vol. 549, no. 7671, pp. 242–246, 2017.
- [70] X. Liu, A. Angone, R. Shaydulin, I. Safro, Y. Alexeev, and L. Cincio, "Layer VQE: A variational approach for combinatorial optimization on noisy quantum computers," *IEEE Transactions on Quantum Engineering*, vol. 3, pp. 1–20, 2022.
- [71] S. Wang, P. Czarnik, A. Arrasmith, M. Cerezo, L. Cincio, and P. J. Coles, "Can error mitigation improve trainability of noisy variational quantum algorithms?" *arXiv preprint arXiv:2109.01051*, 2021.
- [72] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, "Noise-induced barren plateaus in variational quantum algorithms," *Nature communications*, vol. 12, no. 1, p. 6961, 2021.
- [73] E. Fontana, M. Cerezo, A. Arrasmith, I. Rungger, and P. J. Coles, "Non-trivial symmetries in quantum landscapes and their resilience to quantum noise," *Quantum*, vol. 6, p. 804, 2022.
- [74] D. Rabinovich, E. Campos, S. Adhikary, E. Pankovets, D. Vinichenko, and J. Biamonte, "On the gate-error robustness of variational quantum algorithms," *arXiv preprint arXiv:2301.00048*, 2022.
- [75] Y. Pan, Z. He, N. Guo, and Z. Zhang, "Distributionally robust circuit design optimization under variation shifts," in *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*. IEEE, 2023, pp. 1–8.
- [76] T. Hao, K. Liu, and S. Tannu, "Enabling high performance debugging for variational quantum algorithms using compressed sensing," in *Proceedings of the 50th Annual International Symposium on Computer Architecture*, 2023, pp. 1–13.
- [77] A. Pérez-Salinas, H. Wang, and X. Bonet-Monroig, "Analyzing variational quantum landscapes with information content," *npj Quantum Information*, vol. 10, no. 1, p. 27, 2024.
- [78] T. Hao, Z. He, R. Shaydulin, M. Pistoia, and S. Tannu, "Variational quantum algorithm landscape reconstruction by low-rank tensor completion," *arXiv preprint arXiv:2405.10941*, 2024.

• • •