

Received 23 September 2022; revised 1 December 2022; accepted 12 December 2022; date of publication 16 December 2022; date of current version 31 January 2023.

Digital Object Identifier 10.1109/TQE.2022.3229747

Error-Mitigation-Aided Optimization of Parameterized Quantum Circuits: Convergence Analysis

SHARU THERESA JOSE¹  (Member, IEEE),
AND OSVALDO SIMEONE²  (Fellow, IEEE)

¹Department of Computer Science, University of Birmingham, Birmingham B15 2TT, U.K.

²Department of Engineering, King's College London, London WC2R 2LS, U.K.

Corresponding author: Sharu Theresa Jose (e-mail: s.t.jose@bham.ac.uk).

This work was supported by the European Research Council under the European Union's Horizon 2020 Research and Innovation Programme under Grant 725731. This work was partly done when Sharu Theresa Jose was a postdoctoral researcher at King's College London.

ABSTRACT Variational quantum algorithms (VQAs) offer the most promising path to obtaining quantum advantages via noisy intermediate-scale quantum (NISQ) processors. Such systems leverage classical optimization to tune the parameters of a parameterized quantum circuit (PQC). The goal is minimizing a cost function that depends on measurement outputs obtained from the PQC. Optimization is typically implemented via stochastic gradient descent (SGD). On NISQ computers, gate noise due to imperfections and decoherence affects the stochastic gradient estimates by introducing a bias. Quantum error mitigation (QEM) techniques can reduce the estimation bias without requiring any increase in the number of qubits, but they in turn cause an increase in the variance of the gradient estimates. This work studies the impact of quantum gate noise on the convergence of SGD for the variational eigensolver (VQE), a fundamental instance of VQAs. The main goal is ascertaining conditions under which QEM can enhance the performance of SGD for VQEs. It is shown that quantum gate noise induces a nonzero error-floor on the convergence error of SGD (evaluated with respect to a reference noiseless PQC), which depends on the number of noisy gates, the strength of the noise, as well as the eigenspectrum of the observable being measured and minimized. In contrast, with QEM, any arbitrarily small error can be obtained. Furthermore, for error levels attainable with or without QEM, QEM can reduce the number of required iterations, but only as long as the quantum noise level is sufficiently large, and a sufficiently large number of measurements is allowed at each SGD iteration. Numerical examples for a max-cut problem corroborate the main theoretical findings.

INDEX TERMS Quantum computing, quantum error mitigation, quantum gate noise, variational quantum algorithms.

I. INTRODUCTION

A. MOTIVATION

Current noisy intermediate-scale quantum (NISQ) processors are severely limited by the availability of a small number of qubits and by the unavoidable errors due to *quantum gate noise*. Quantum gate noise refers to unwanted interactions of qubits with the environment, leading to *decoherence*; as well as to the imperfect execution of user-specified quantum gates on the physical device, leading to *gate infidelity*. While fault-tolerant computation is not feasible on NISQ processors due to the small number of available qubits [1], forms of *quantum error*

mitigation (QEM), which do not require the addition of further qubits, are possible [2], [3]. Many promising applications of NISQ devices involve *parameterized quantum circuits* (PQCs), whose architecture—or *ansatz*—is compatible with NISQ processors. As illustrated in Fig. 1, VQAs tune the parameters θ defining the operation of a PQC via a classical optimizer that relies on measurements from the PQC [4], [5]. This article studies the impact of quantum gate noise on the performance of VQAs by focusing on the question of whether QEM can be effective in enhancing the performance of an optimized PQC and/or in reducing the optimization complexity.




FIGURE 1. Illustration of the VQA framework the model parameter vector θ , specifying the operation of a, generally noisy, PQC $\mathcal{U}(\theta)$, is optimized via SGD based on measurement outputs. The measurement outputs are used to produce an estimate \hat{g}_t of the gradient. The estimate \hat{g}_t is affected by the inherent randomness of quantum measurements, as well as by the noise introduced by the imperfections of the quantum gates within the PQC.

To address this question, we consider the most fundamental instance of VQAs, namely the *variational quantum eigensolver (VQE)* [6]. The goal of the VQE is to ensure that the output state of the PQC, denoted as $|\Psi(\theta)\rangle$, provides a good approximation of the ground state of a given observable H . Mathematically, we wish to approximately solve the problem

$$\theta^* = \arg \min_{\theta} \langle H \rangle_{|\Psi(\theta)\rangle} \quad (1)$$

of minimizing the expected value of observable H when evaluated for the output state $|\Psi(\theta)\rangle$. Assuming that the ansatz of the PQC is sufficiently expressive, the solution $|\Psi(\theta^*)\rangle$ obtained from problem (1) is a close approximation of the desired eigenstate. Applications of VQEs include the solution of quadratic unconstrained binary optimization [7], [8], as well as problems in quantum chemistry [9].

In a VQE, the optimization (1) of the parameters θ of the PQC is typically carried out in an iterative manner by means of *stochastic gradient descent (SGD)*. SGD estimates the gradient $\nabla_{\theta} \langle H \rangle_{|\Psi(\theta)\rangle}$ of the expected value $\langle H \rangle_{|\Psi(\theta)\rangle}$ via measurements from the output of the PQC [4], [5]. Such estimates are affected by the inherent randomness of quantum measurements, as well as by the noise introduced by the imperfections of the quantum gates within the PQC and by decoherence. As a result of such quantum gate noise, as illustrated in Fig. 3(a), the estimates of the gradient $\nabla_{\theta} \langle H \rangle_{|\Psi(\theta)\rangle}$ are *biased* [10].

QEM techniques provide an *algorithmic* approach to mitigate quantum gate noise in NISQ devices. Unlike quantum error correction codes [1], [11], [12], quantum error mitigation requires no additional qubit resources. Examples of quantum error mitigation techniques include *quasi-probabilistic QEM* [2], zero noise extrapolation [2], [13], randomized compiling [14], and Pauli-frame randomization [15]. Henceforth, we use QEM to denote quasi-probabilistic quantum error mitigation.

As illustrated in Fig. 2, QEM provides protection against quantum gate noise by running multiple, N_c , noisy quantum circuits, which are sampled from a set of circuits implementable on the NISQ device, and by postprocessing the measurement outputs. As sketched in Fig. 3(b), QEM can reduce the mentioned quantum noise-induced bias in the

TABLE 1. Main Notations Used in This Article

Description	Notation
Number of parameters of PQC	D
Number of measurements per iteration	$2DN_m$
Number of sampled noisy circuits	N_c
Noise level	$\gamma \in [0, 1]$

measurement of a quantum observable, while generally increasing the corresponding variance [16].

Using QEM, one can hence obtain a *less biased* estimate of the gradient $\nabla_{\theta} \langle H \rangle_{|\Psi(\theta)\rangle}$. This can potentially improve the convergence of SGD when applied to the VQE problem (1). However, the bias reduction should be weighted against the variance increase as a function of the number of quantum measurements that one can afford at any SGD iteration. Under which conditions is it advantageous to trade an increased variance for a decrease in the bias in VQEs?

B. MAIN CONTRIBUTIONS

This article aims at providing some theoretical insights on the potential advantages of QEM for VQE by studying the convergence of SGD for problem (1) with and without QEM. Assuming that gradients are estimated via the *parameter-shift rule* [4], the specific contributions are summarized in Table 2, where we report the derived upper bounds on the number of SGD iterations required to ensure convergence to a target error floor of order $\mathcal{O}(\delta)^1$ for any $\delta > 0$. We compare the performance of SGD with shot noise only; as well as with shot and gate noise with and without QEM. For all schemes, as per the notation detailed in Table 1, we fix the number of model parameters θ to integer D ; the number of measurements per-iteration as $2DN_m$ for some integer N_m ; and we quantify the noise level with parameter $\gamma \in [0, 1]$. Noise is modeled as an Pauli quantum channel that can account for spatial correlations across the qubits [18]. The ansatz and the observable H determine a constant μ (see Section III-C for details). The main results are described as follows.

- 1) We first quantify the impact of the inherent randomness of quantum measurements, also known as *shot noise*, as well as of the bias induced by *quantum gate noise*, on the convergence of SGD. We show that, while the impact of shot noise—quantified by the variance term V in Table 2—can be arbitrarily decreased by increasing the number N_m of quantum measurements, the presence of quantum gate noise induces a *nonzero error floor* on the convergence of SGD. The error floor is caused by the bias in the estimate of the gradient, as quantified by the term $B^{\mathcal{E}}$ in Table 2, which is proved to depend on the strength of the quantum gate noise γ , on the number D of free parameters of the PQC, and on the eigenspectrum of the observable H . Accordingly, the

¹In this work, we use the standard big-O notation $\mathcal{O}(\cdot)$, with notation $\tilde{\mathcal{O}}(\cdot)$ further hiding the poly-logarithmic factors [17].

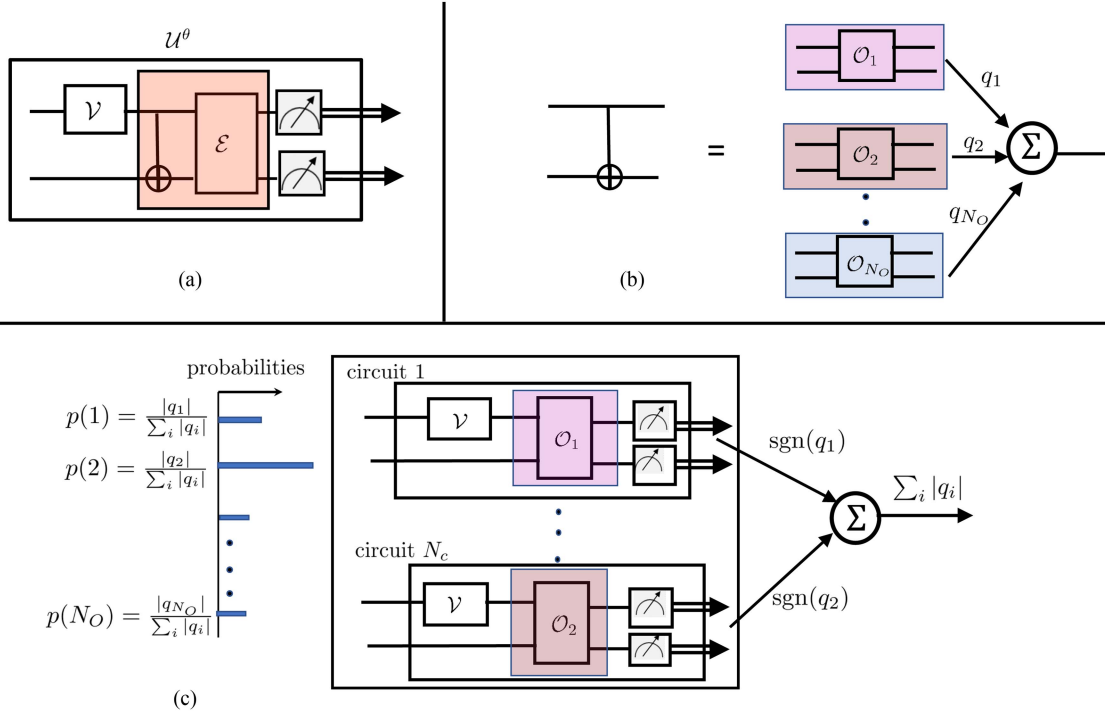


FIGURE 2. VQA aided by quasi-probabilistic quantum error mitigation (QEM). (a) Noisy PQC where gate noise channel ε acts on a CNOT gate. (b) quasi-probabilistic representation (QPR) of a CNOT gate as a linear combination of a set of operations $\{O_i\}$ implementable on the quantum computer. (c) QEM approach, whereby multiple circuits are sampled, and their outputs combined, in order to approximate the operation of the PQC in the absence of quantum gate noise.

TABLE 2. (Left) Upper Bounds on the Number of Iterations T Required to Ensure a Target Error Floor, $\mathbb{E}[L(\theta^T)] - L(\theta^*)$, of Order $\mathcal{O}(\delta)$ for Any $\delta > 0$ When Only Shot Noise Is Present, As Well As in the Presence of Both Shot and Gate Noise With and Without QEM. Parameters V , V^ε , and V^{QEM} quantify the respective variances of the gradient estimates; $\mu > 0$ is a positive constant dependent on the ansatz and on the observable H ; $B^\varepsilon > 0$ accounts for the bias due to quantum gate noise; and functions $c(\gamma)$, $c_1(\gamma)$, and $c_2(\gamma)$ are detailed in Section IV-B and Section V-C.

Schemes	Iteration Complexity
shot-noise only	$\tilde{\mathcal{O}}\left(\log \frac{1}{\delta} + \frac{V}{\mu\delta}\right)$
shot and gate noise	$\tilde{\mathcal{O}}\left(\log \frac{1}{\delta - B^\varepsilon/\mu} + \frac{V^\varepsilon}{\mu\delta}\right)$
shot and gate noise with QEM	$\tilde{\mathcal{O}}\left(\log \frac{1}{\delta} + \frac{V^{\text{QEM}}}{\mu\delta}\right)$

Parameters	Scaling
variance V	$\mathcal{O}(D/N_m)$
bias B^ε	$\mathcal{O}(D\gamma)$
variance V^ε	$\mathcal{O}(Dc(\gamma)/N_m)$
variance V^{QEM}	$\mathcal{O}(c_1(\gamma)V^\varepsilon) + \mathcal{O}(c_2(\gamma)D/N_c)$

derived upper bound on the number of SGD iterations in the presence of gate noise diverges for any error floor δ smaller than some level, B^ε/μ , dependent on the bias.

- To mitigate the error floor induced by quantum gate noise, we then study the impact of QEM on the convergence of SGD. As seen in Table 2, QEM can obtain any error $\delta > 0$ in a finite number of iterations thanks to the elimination of the bias on the gradient estimate caused by quantum gate noise. Therefore, in order to obtain an error smaller than B^ε/μ , QEM is necessary. For levels of the error larger than B^ε/μ , QEM can help reduce the number of required SGD iterations. According to the bounds derived in this article, this is the case if the noise level γ is sufficiently large and if the number N_c of circuits sampled at each iteration is large enough. In fact, as illustrated in Fig. 3, QEM causes an increase in the variance of the gradient estimator, as

quantified by the term V^{QEM} in Table 2. In particular, QEM adds a contribution to the variance that decreases as $1/N_c$, requiring a sufficiently large value of N_c to ensure that this term does not dominate the overall performance.

- We corroborate the theoretical findings with numerical experiments on the VQE problem of maximizing the weighted max-cut Hamiltonian [8].

C. RELATED WORKS

1) IMPACT OF QUANTUM GATE NOISE

Recent works [19], [20] provide numerical evidence for the detrimental impact of the quantum gate noise on the trainability of PQCs. Wang et al. [21] provided a theoretical explanation of this phenomenon by showing that local Pauli gate noise induces barren plateaus in the loss function. This means that the gradient of the cost function vanishes with a

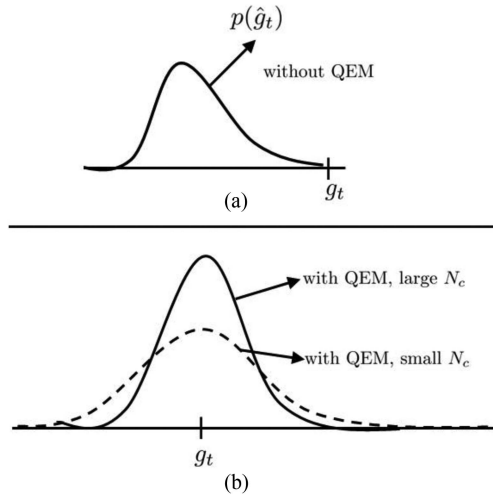


FIGURE 3. Illustration of the effect of QEM. (Top) Without QEM, the presence of quantum gate noise results in a biased gradient estimator. (Bottom) Employing QEM by sampling N_c circuits can decrease the bias, at the cost of increase in the variance, unless N_c is sufficiently large.

rate exponential in the number of qubits, as long as the circuit depth is linear in the number of qubits.

Even in the absence of barren plateaus, quantum gate noise can affect the convergence of SGD algorithms for the optimization of PQCs by causing the stochastic gradients to be biased estimates of the gradients. This aspect was studied from a theoretical standpoint in [10] for VQE by focusing on gradient estimators based on symmetric logarithmic derivative operators. The variance of the resulting gradient estimator was quantified in [10] via the quantum Fisher information, which was shown to decrease with increasing gate noise. In contrast, the focus of this work is on SGD schemes that leverage the more commonly used parameter-shift rule-based gradient estimators.

Also related is the work [22], which studies the impact of noisy gates on the convergence of SGD used for empirical risk minimization in supervised learning. However, Du et al. [22] do not explicitly characterize the nonzero error floor induced by the quantum gate noise. Furthermore, none of the above-mentioned works account for the impact of QEM.

2) QUANTUM ERROR MITIGATION

The detrimental effect of quantum gate noise observed in the references summarized above suggest that optimization of PQCs can indeed benefit from quantum error mitigation. Earlier works that attempt to investigate the advantage of quantum error mitigation include [23], which demonstrates enhanced performance in terms of accuracy for VQE with error mitigation via zero noise extrapolation on a superconducting quantum processor. Resorting to error mitigation via hidden inverses [24], Leyton-Ortega et al. [25] proves experimentally that error-mitigated VQE converges faster when applied to problems in quantum chemistry. Other studies focus

on developing error mitigation strategies that can efficiently mitigate the bias induced by quantum gate noise [26], [27]. In particular, Strikis et al. [28] proposed a learning-based quasi-probabilistic error mitigation strategy (LB-QEM) that does not require knowledge of noise model and gate decompositions. Strikis et al. [28] showed experimentally that the LB-QEM can reduce the bias induced by gate noise for VQE. None of the above-mentioned works provide a theoretical analysis of the observed enhanced convergence performance of SGD-based VQE due to QEM.

In a different line of research, the work in [21] and [29] addressed the question of whether error mitigation techniques can improve the resolvability of any two points on the cost landscape, thus tackling noise-induced barren plateaus. It is shown that QEM improves resolvability under global depolarizing noise, provided that the number of qubits is small. However, any improvement in resolvability due to QEM degrades exponentially for large problem sizes, in terms of the number of qubits and circuit depth, in the presence of local depolarizing noise.

Overall, to the best of authors' knowledge, no prior work has addressed the convergence performance of SGD-based VQE with the parameter-shift rule and in the presence of QEM.

D. ORGANIZATION

The rest of this article is organized as follows. In Section II, we detail the VQE problem (1), the solution method based on SGD, as well as the considered quantum gate noise model. Section III presents the convergence analysis of SGD for VQE when gate noise is negligible. The impact of gate noise on the convergence of SGD is studied in Section IV. In Section V, we analyze the convergence of SGD with QEM. Finally, Section VI presents numerical experiments that corroborate our analytical findings. Finally, Section VII concludes this article.

II. PROBLEM FORMULATION

The goal of the VQE is to prepare a quantum state, via a PQC that minimizes the expectation of an observable. This problem is also at the core of quantum machine learning algorithms, and is the focus of this article. This section first introduces the VQE problem, as well as the standard solution method based on SGD. Then, we describe the noise model assumed for the quantum circuit.

A. VARIATIONAL QUANTUM EIGENSOLVER

An ideal, noiseless, PQC implements a unitary transformation $U(\theta)$, operating on n qubits, that is, parameterized by a vector $\theta = (\theta_1, \dots, \theta_D)$, with $\theta_d \in \mathbb{R}$ for $d = 1, \dots, D$. A common architecture for the PQC, also known as ansatz, prescribes unitaries of the form [30]

$$U(\theta) = \prod_{d=1}^D U_d(\theta_d) V_d \quad (2)$$

where the d th parameter θ_d in vector θ determines the operation of the d th unitary $U_d(\theta_d)$. The parameterized gate is assumed to be given as [31]

$$U_d(\theta_d) = \exp\left(-i\frac{\theta_d}{2}G_d\right) \quad (3)$$

where $G_d \in \{I, X, Y, Z\}^{\otimes n}$ denotes the Pauli string generator. Furthermore, the unitary V_d in (2) is fixed, and not dependent on parameters θ .

The noiseless PQC is applied to n qubits that are initially in the ground state $|0\rangle$ to yield the parameterized quantum state

$$|\Psi(\theta)\rangle = U(\theta)|0\rangle. \quad (4)$$

The pure-state density matrix corresponding to the quantum state (4) is denoted as

$$\Psi(\theta) = |\Psi(\theta)\rangle\langle\Psi(\theta)|. \quad (5)$$

The VQE seeks to find the parameter vector $\theta \in \mathbb{R}^D$ that addresses the problem (1), which we restate here as

$$\min_{\theta \in \mathbb{R}^D} \{L(\theta) := \langle H \rangle_{|\Psi(\theta)\rangle}\} \quad (6)$$

where the loss function $L(\theta)$ is the expected value

$$\langle H \rangle_{|\Psi(\theta)\rangle} = \langle 0|U^\dagger(\theta)HU(\theta)|0\rangle = \text{Tr}(H\Psi(\theta)) \quad (7)$$

of an observable H for the state $|\Psi(\theta)\rangle$ produced by the PQC. We denote as

$$L^* = \min_{\theta \in \mathbb{R}^D} L(\theta) \quad (8)$$

the minimum value of the loss function.

The Hermitian matrix H describing the observable can be written via its eigendecomposition as

$$H = \sum_{y=1}^{N_h} h_y \Pi_y \quad (9)$$

where $\{h_y\}_{y=1}^{N_h}$ denote the $N_h \leq 2^n$ distinct eigenvalues of matrix H , and $\{\Pi_y\}_{y=1}^{N_h}$ are the projection operators onto the corresponding eigenspaces. Consequently, the loss function $L(\theta)$ in problem (6) can be expressed as

$$L(\theta) = \sum_{y=1}^{N_h} h_y \text{Tr}(\Pi_y \Psi(\theta)). \quad (10)$$

B. STOCHASTIC GRADIENT DESCENT

Assuming the ideal case in which the PQC is noiseless, problem (6) can be addressed by following a hybrid quantum-classical optimization approach. Accordingly, N_m measurements of the observable H are made for the output state $|\Psi(\theta)\rangle$ of PQC, producing samples $\{H_j\}_{j=1}^{N_m}$ with $H_j \in \{h_1, \dots, h_{N_h}\}$. From these samples, an estimate of the loss function $L(\theta)$ is obtained as

$$\hat{L}(\theta) = \widehat{\langle H \rangle}_{|\Psi(\theta)\rangle} = \frac{1}{N_m} \sum_{j=1}^{N_m} H_j \quad (11)$$

where the “hat” notation is used for empirical estimates. The empirical estimate (11) is used as input to classical optimization algorithm.

In this work, we focus on the standard implementation of VQE based on SGD optimization. Starting from an initialization $\theta^0 \in \mathbb{R}^D$, SGD performs the iterative update of the parameter θ as

$$\theta^{t+1} = \theta^t - \eta_t \hat{g}_t \quad (12)$$

for iterations $t = 1, \dots, T$, where \hat{g}_t represents a stochastic estimate of the gradient

$$\nabla L(\theta)|_{\theta=\theta^t} = \begin{bmatrix} \frac{\partial L(\theta)}{\partial \theta_1} \\ \vdots \\ \frac{\partial L(\theta)}{\partial \theta_D} \end{bmatrix}_{\theta=\theta^t} \quad (13)$$

of the loss function $L(\theta)$ at the current iterate $\theta = \theta^t$, and $\eta_t > 0$ denotes the learning rate at the t th iteration. We will detail how to obtain the estimate \hat{g}_t as a function of measurements of observable H of the form (11) in the following sections.

C. NOISY QUANTUM GATES

In the current era of NISQ hardware, the implementation of the unitary gates that determine the operation of the PQC in (2) is subject to quantum gate noise due to decoherence as well as gate infidelities. Therefore, a quantum computer can in practice only implement a given set Ω of, generally noisy, quantum operations.

The set $\Omega = \{\mathcal{O}_i^\theta\}$ consists of a set of gates $\{\mathcal{O}_i^\theta\}$ that can be implemented on the quantum device. The notation \mathcal{O}_i^θ emphasizes the possible dependence of operation \mathcal{O} on a model parameter θ . Each gate \mathcal{O}_i^θ can be modeled as a *quantum channel*, that is, as a completely positive trace preserving map between density matrices. Accordingly, we can write the corresponding maps as $\mathcal{O}_i^\theta(\rho) = \rho'$, where ρ is the input density matrix and ρ' the output density matrix [32]. The operations in set Ω can be determined via quantum gate set tomography as detailed in [3].

Each unitary gate $U_d(\theta_d)V_d$ in the ideal, noiseless, PQC (2) corresponds to the noiseless quantum channel

$$\rho' = \mathcal{U}_d^{\theta_d}(\rho) = V_d U_d(\theta_d) \rho U_d(\theta_d)^\dagger V_d^\dagger. \quad (14)$$

Thus, the operation of the overall noiseless PQC $U(\theta)$ in (2) on the initial state $\rho_0 = |0\rangle\langle 0|$ can be expressed as the composition of the mappings

$$\mathcal{U}^\theta(\rho_0) = \mathcal{U}_D^{\theta_D} \circ \dots \circ \mathcal{U}_1^{\theta_1}(\rho_0) \quad (15)$$

where \circ indicates the composition operation.

In practice, the quantum gates in the ideal PQC (15) can be only approximately implemented on the quantum computer, in the sense that the set of feasible operations Ω includes only noisy versions of such gates. Formally, given the noiseless parameterized gate $\mathcal{U}_d^{\theta_d}(\cdot)$ in (14), we assume that the set Ω includes a noisy operation of the form

$$\mathcal{O}_i^{\theta_d}(\cdot) = \mathcal{E} \circ \mathcal{U}_d^{\theta_d}(\cdot) = \tilde{\mathcal{U}}_d^{\theta_d}(\cdot) \quad (16)$$

where $\mathcal{E}(\cdot)$ represents a quantum channel. Following (16), we will write as $\tilde{\mathcal{U}}_d^{\theta_d}(\cdot)$ the noisy quantum gate corresponding to the noiseless gate $\mathcal{U}_d^{\theta_d}(\cdot)$ in (14).

Throughout this work, we model gate noise via *Pauli quantum channels* acting on the n qubits of the form

$$\mathcal{E}(\rho) = (1 - \epsilon)\rho + \epsilon \sum_j E_j \rho E_j^\dagger \quad (17)$$

where $0 \leq \epsilon \leq 1$ is the error probability, and the sum in (17) runs over a subset of $4^n - 1$ n -length Pauli strings E_j , excluding the identity matrix. The noise model in (17) includes the standard model with independent channels across the n qubits, and it can also more generally account for spatial correlations across n qubits [18]. It is practically well justified when quantum gate set tomography returns a real, diagonal matrix for the Pauli transfer matrix representation [33].

We write P_s , with $s \in \{0, 1, 2, 3\}$, for a single-qubit Pauli operator, where $P_0 = I$, $P_1 = X$, $P_2 = Y$, and $P_3 = Z$. With this notation, the Kraus operators in (17) can be written as

$$E_j = \sqrt{p_j} P_{s_{j,0}} \otimes \dots \otimes P_{s_{j,n-1}} \quad (18)$$

where $s_{j,k} \in \{0, 1, 2, 3\}$ for $k = 0, 1, \dots, n-1$; we have the equality $\sum_j p_j = 1$; and the indices $s_{j,k}$ cannot be all equal to 0.

As an important example that we will assume in some of the theoretical derivations, the noise model (17) includes the *depolarizing noise channel*, which is obtained when the Kraus operators E_j include all $4^n - 1$ Pauli strings excluding the identity matrix, and $p_j = 1/(4^n - 1)$. Note that this corresponds to a situation in which the gate noise applies separately to each qubit.

Overall, when QEM is not applied, the actual implementation of PQC (2) on the quantum computer amounts to the following composition of noisy gates:

$$\tilde{U}^\theta(\cdot) = \tilde{U}_D^{\theta_D} \circ \dots \circ \tilde{U}_1^{\theta_1}(\cdot) \quad (19)$$

where each operation $\tilde{U}_d^{\theta_d}(\cdot)$, for $d = 1, \dots, D$, belongs to the family of quantum operations Ω . Accordingly, the PQC operation of (19) on the initial quantum state $\rho_0 = |0\rangle\langle 0|$ results in the output noisy quantum state

$$\rho^\mathcal{E}(\theta) = \tilde{U}^\theta(\rho_0). \quad (20)$$

D. FEASIBLE QUANTUM GATE SET

Apart from the noisy gates (16), in order to enable QEM, we assume that the set Ω of implementable operations also includes the noisy cascade of Pauli operations and ideal gates $U_d^{\theta_d}(\cdot)$. To denote an n -qubit Pauli operation, we, henceforth, use the following standard convention. Let $s = (s_0, \dots, s_{n-1})$ be a vector of indices with $s_k \in \{0, 1, 2, 3\}$ for $k = 0, \dots, n-1$. Using this notation, we denote an n -qubit Pauli operator as

$$\mathcal{P}_s(\cdot) = \mathcal{P}_{s_0}(\cdot) \otimes \dots \otimes \mathcal{P}_{s_{n-1}}(\cdot) \quad (21)$$

where

$$\mathcal{P}_{s_k}(\cdot) = P_{s_k}(\cdot) P_{s_k}^\dagger \quad (22)$$

for $k = 0, \dots, n-1$. The set of implementable operations Ω then includes all the noisy gates $\tilde{U}_d^{\theta_d}(\cdot)$, as well as the noisy concatenations $\mathcal{E} \circ \mathcal{P}_s \circ U_d^{\theta_d}$ of Pauli operation \mathcal{P}_s on the ideal unitary for all $s \in \{0, 1, 2, 3\}^{\otimes n}$.

III. IMPACT OF SHOT NOISE

We start our analysis of the SGD-based optimization for the VQE problem (6) by assuming that the quantum gate noise is negligible. This amounts to assuming that the set Ω of feasible quantum operations includes all gates $U_d^{\theta_d}(\cdot)$ in (15), and hence, we have $\tilde{U}_d^{\theta_d}(\cdot) = U_d^{\theta_d}(\cdot)$ in (16). Accordingly, the analysis in this section only accounts for the randomness due to shot noise, that is, due to the inherent stochasticity of quantum measurements. The impact of gate noise will be studied in the following section.

A. ESTIMATING THE GRADIENT

With no gate noise, the d th component of the gradient, $[g_t]_d = [\nabla L(\theta^t)]_d$, can be estimated via the *parameter shift-rule* as [34]

$$[\hat{g}_t]_d = \frac{1}{2} \left(\widehat{\langle H \rangle}_{|\Psi(\theta^t + \frac{\pi}{2} e_d)\rangle} - \widehat{\langle H \rangle}_{|\Psi(\theta^t - \frac{\pi}{2} e_d)\rangle} \right) \quad (23)$$

where e_d denotes a unit vector with all zero elements except in the d th position. In (23), the quantum state $|\Psi(\theta^t + \frac{\pi}{2} e_d)\rangle$ is obtained by shifting the phase of d th parameter θ_d by $\pi/2$; and $|\Psi(\theta^t - \frac{\pi}{2} e_d)\rangle$ is obtained by shifting the phase of d th parameter θ_d by $-\pi/2$. Furthermore, as in (11), the notation $\widehat{\langle H \rangle}_{|\Psi(\theta)\rangle}$ describes an empirical estimation of the expected observable H , obtained from N_m independent identically distributed (i.i.d.) measurements $\{H_1, \dots, H_{N_m}\}$ of the observable under state $|\Psi(\theta)\rangle$.

The measurements of observable H return $H_j = h_y$ with probability

$$p(y|\theta) = \text{Tr}(\Pi_y \Psi(\theta)), \quad y \in \{1, \dots, N_h\} \quad (24)$$

for all measurements $j = 1, \dots, N_m$. The resulting empirical estimator of each expected value in (23) evaluates as in (11). Therefore, estimating the full gradient vector (13) requires running the PQC a number of times equal to $2DN_m$.

Upon obtaining a solution θ^T after T iterations, the final outcome of the VQE is evaluated as $\langle H \rangle_{|\Psi(\theta^T)\rangle}$, where the expectation is taken over a large number of measurements.

B. PROPERTIES OF THE GRADIENT ESTIMATE

The estimate in (23) is unbiased, i.e., we have

$$\mathbb{E}[\hat{g}_t] = g_t = \nabla L(\theta^t) \quad (25)$$

where the expectation $\mathbb{E}[\cdot]$ is taken with respect to the $2DN_m$ measurements used in (23). Therefore, we can decompose the gradient estimator (23) as

$$\hat{g}_t = g_t + \xi_t \quad (26)$$

with noise ξ_t satisfying the conditions $\mathbb{E}[\xi_t] = 0$ and $\text{var}(\xi_t) = \mathbb{E}[\|\hat{g}_t - g_t\|^2]$. The following lemma gives an upper bound on the variance $\text{var}(\xi_t)$. To this end, we first rewrite the empirical estimate (11) as

$$\widehat{\langle H \rangle}_{|\Psi(\theta)\rangle} = \frac{1}{N_m} \sum_{j=1}^{N_m} \sum_{y=1}^{N_h} h_y \mathbb{I}\{Y_j = y\} \quad (27)$$

where $Y_j \in \{1, \dots, N_h\}$ is the random variable defining the index corresponding to the j th measurement so that we have

the equality $H_j = h_{Y_j}$. Furthermore, we introduce the variance of the Bernoulli quantum measurement $\mathbb{I}\{Y_j = y\}$ as

$$\nu(p(y|\theta)) = p(y|\theta)(1 - p(y|\theta)) \quad (28)$$

where $p(y|\theta)$ is as defined in (24).

Lemma III.1: The variance of the unbiased gradient estimator \hat{g}_t in (23) can be upper bounded as

$$\text{var}(\xi_t) \leq \frac{\nu N_h D \text{Tr}(H^2)}{2N_m} := V \quad (29)$$

where $\nu \in [0, 0.25]$ is

$$\nu = \max_{\theta \in \mathbb{R}^D} \max_{y \in \{1, \dots, N_h\}} \nu(p(y|\theta)) \quad (30)$$

with $\nu(p(y|\theta))$ in (28).

Proof: The proof is in Appendix A. \blacksquare

Lemma III.1 shows that the variance of the gradient estimate (26) decreases (at least) as $1/(2N_m)$ with respect to the number of per-parameter measurements $2N_m$, while being proportional to the number of parameters, D . The term ν in (30) captures the maximum variance of the Bernoulli measurement indicator variables $\mathbb{I}\{Y_j = y\}$, hence accounting for the degree of randomness of the measurements. Finally, the term $\text{Tr}(H^2)$ is a scale parameter dependent on the numerical values $\{h_y\}_{y=1}^{N_h}$ of the observable H .

C. CONVERGENCE OF SGD

We now study the convergence properties of SGD in (12) with the unbiased estimator (26) in the case under study of noiseless gates. Toward this goal, we make the following assumptions on the loss function as in [10] and [30]. These assumptions are further discussed in Appendix B.

Assumption III.1: The loss function $L(\theta)$ is \mathcal{L} -smooth, i.e., we have the inequality

$$L(\theta) \leq L(\theta') + \nabla L(\theta')^T (\theta - \theta') + \frac{\mathcal{L}}{2} \|\theta - \theta'\|^2 \quad (31)$$

for all $\theta, \theta' \in \Theta$. Furthermore, it satisfies the μ -Polyak-Lojasiewicz (PL) condition, i.e., there exists a constant $\mu > 0$ and $\mu \leq \mathcal{L}$ such that the inequality

$$\|\nabla L(\theta)\|^2 \geq 2\mu(L(\theta) - L^*) \quad (32)$$

holds for all $\theta \in \Theta$.

In practice, the constant μ depends on the number n of qubits and the number of gate parameters D , since a larger circuit is typically characterized by smaller-gradient norms [35].

The following result gives a bound on the optimality of the SGD output θ^T after T iteration. The proof of the result is based on the classical convergence analysis of SGD, and can be found in [17, Th. 6].

Theorem III.1: Under Assumption III.1, for any given initial point θ^0 , the following bound holds for any fixed learning rate $\eta_t = \eta \leq 1/\mathcal{L}$:

$$\mathbb{E}[L(\theta^T)] - L^* \leq (1 - \eta\mu)^T (\mathbb{E}[L(\theta^0)] - L^*) + \frac{1}{2} \left[\frac{\eta\mathcal{L}V}{\mu} \right] \quad (33)$$

where V is as defined in (29), and the expectation is taken over the distribution of the measurement outputs. Furthermore, given some target error level $\delta > 0$, for learning rate $\eta = \eta^{\text{shot-noise}} \leq \min\{\frac{1}{\mathcal{L}}, \frac{\delta\mu}{\mathcal{L}V}\}$, a number of iterations

$$T^{\text{shot-noise}} = \tilde{\mathcal{O}} \left(\log \frac{1}{\delta} + \frac{V}{\delta\mu} \right) \frac{\mathcal{L}}{\mu} \quad (34)$$

is sufficient to ensure an error $\mathbb{E}[L(\theta^T)] - L^* = \mathcal{O}(\delta)$.

The result in (33) shows that SGD with the unbiased gradient estimator (26) converges up to an error floor of the order $\mathcal{O}(\eta\mathcal{L}V)$. By choosing the learning rate η sufficiently small, this error can be made arbitrarily small, thereby guaranteeing convergence in at most $T^{\text{shot-noise}}$ iterations.

IV. IMPACT OF GATE NOISE

In this section, we study the impact of gate noise. To this end, we assume that all quantum gate $\tilde{\mathcal{U}}_d^{\theta_d}(\cdot) \in \Omega$ in the PQC (15) can be written as in (16) with quantum noise (17). As we will see, the presence of noisy gates induces a bias in the gradient estimator (26) studied in the previous section. The main result of this section quantifies the impact of the bias in the gradient estimate on the convergence of SGD in (12).

A. ESTIMATING THE GRADIENT

To start, in a manner similar to (11), let us denote as $\widehat{\langle H \rangle}_{\rho^\mathcal{E}(\theta)}$, the empirical estimator of the observable H obtained from measurements of the state $\rho^\mathcal{E}(\theta)$ in (20) produced by the noisy PQC. This estimate is obtained from N_m i.i.d. measurements $\{H_1^\mathcal{E}, \dots, H_{N_m}^\mathcal{E}\}$ of observable H under the noisy quantum state $\rho^\mathcal{E}(\theta)$. We also write as $\{Y_1^\mathcal{E}, \dots, Y_{N_m}^\mathcal{E}\}$, with $Y_j^\mathcal{E} \in \{1, \dots, N_h\}$, the corresponding indices of the measurement outcomes, so that we have $H_j^\mathcal{E} = h_{Y_j^\mathcal{E}}$. Accordingly, the output $Y_j^\mathcal{E} = y$ is produced with the probability

$$p^\mathcal{E}(y|\theta) = \text{Tr}(\Pi_y \rho^\mathcal{E}(\theta)). \quad (35)$$

At t th iteration, following the parameter shift rule as in (23), we adopt the stochastic gradient estimator $\hat{g}_t^\mathcal{E}$ whose d th component is obtained as

$$[\hat{g}_t^\mathcal{E}]_d = \frac{1}{2} \left(\widehat{\langle H \rangle}_{\rho^\mathcal{E}(\theta^t + \frac{\pi}{2} e_d)} - \widehat{\langle H \rangle}_{\rho^\mathcal{E}(\theta^t - \frac{\pi}{2} e_d)} \right). \quad (36)$$

The estimate (36) involves the noisy quantum states $\rho^\mathcal{E}(\theta^t \pm \frac{\pi}{2} e_d)$ in (20) that are produced by the noisy PQC with d th parameter θ_d phase-shifted by $\pi/2$ or $-\pi/2$. In a manner similar to (28), we define the variance of the Bernoulli quantum measurement $\mathbb{I}\{Y_j^\mathcal{E} = y\}$ as

$$\nu(p^\mathcal{E}(y|\theta)) = p^\mathcal{E}(y|\theta)(1 - p^\mathcal{E}(y|\theta)). \quad (37)$$

B. PROPERTIES OF THE GRADIENT ESTIMATE

The stochastic gradient estimator in (36) is affected by two sources of randomness, namely the shot noise due to quantum measurements as studied in the previous section, and the distortion caused by the noisy quantum gates. While shot noise does not cause a bias in the gradient estimate (36), the quantum gate noise causes the estimate (36) to be biased,

i.e., $\mathbb{E}[\hat{g}_t^\mathcal{E}] \neq g_t$ where $g_t = \nabla L(\theta^t)$ is the exact gradient at iteration t as in (25).

To elaborate on this point, we decompose the biased gradient estimator in (36) as

$$\hat{g}_t^\mathcal{E} = g_t + \underbrace{(g_t^\mathcal{E} - g_t)}_{\text{bias}} + \xi_t^\mathcal{E} \quad (38)$$

where the noise term is zero mean, i.e., $\mathbb{E}[\xi_t^\mathcal{E}] = 0$, and we have defined

$$[g_t^\mathcal{E}]_d = \frac{1}{2} \left(\langle H \rangle_{\rho^\mathcal{E}(\theta^t + \frac{\pi}{2} e_d)} - \langle H \rangle_{\rho^\mathcal{E}(\theta^t - \frac{\pi}{2} e_d)} \right). \quad (39)$$

The gradient (39) represents the average of the estimate (36), where the expectation is taken over the quantum measurements. Consequently, the difference $(g_t^\mathcal{E} - g_t)$ in (39) captures the bias due to gate noise.

The following lemma presents bounds on the bias and on the variance of the measurement noise of the gradient estimator in (36). To proceed, we write the noisy quantum state in (20) in terms of the following decomposition [36]:

$$\rho^\mathcal{E}(\theta) = (1 - \gamma)\Psi(\theta) + \gamma\tilde{\rho}(\theta) \quad (40)$$

where parameter $\gamma = 1 - (1 - \epsilon)^D$ describes the level of noise. In particular, we have $\gamma = 0$ if there is no gate noise ($\epsilon = 0$), and $\gamma = 1$ when the noise level is maximal. The decomposition (40) amounts to a convex combination of the ideal state density matrix $\Psi(\theta) = |\Psi(\theta)\rangle\langle\Psi(\theta)|$ and of the error density matrix $\tilde{\rho}(\theta)$ defined as

$$\tilde{\rho}(\theta) = \frac{1}{\gamma} (\rho^\mathcal{E}(\theta) - (1 - \gamma)\Psi(\theta)). \quad (41)$$

Note that, under the noise model of (17), matrix $\tilde{\rho}(\theta)$ is indeed a valid density operator for $\gamma > 0$ [36].

Let $\tilde{Y} \in \{1, \dots, N_h\}$ denote the index of the quantum measurement $\tilde{H} = h_{\tilde{Y}}$ of the observable H under the error density matrix $\tilde{\rho}(\theta)$ in (40). The probability of observing $\tilde{Y} = y \in \{1, \dots, N_h\}$ is given as

$$\tilde{p}(y|\theta) = \text{Tr}(\Pi_y \tilde{\rho}(\theta)). \quad (42)$$

In a manner similar to (28), we define the variance of the Bernoulli quantum measurement $\mathbb{I}\{\tilde{Y} = y\}$ as

$$v(\tilde{p}(y|\theta)) = \tilde{p}(y|\theta)(1 - \tilde{p}(y|\theta)).$$

Lemma IV.1: The following upper bound holds on the variance of the stochastic gradient estimator $\hat{g}_t^\mathcal{E}$

$$\text{var}(\xi_t^\mathcal{E}) \leq \frac{DN_h \text{Tr}(H^2)}{2N_m} c(\gamma) := V^\mathcal{E} \quad (43)$$

where $c(\gamma) = \max_{y \in \{1, \dots, N_h\}} \max_{\theta \in \mathbb{R}^D} v(p^\mathcal{E}(y|\theta))$, with

$$v(p^\mathcal{E}(y|\theta)) = \gamma v(\tilde{p}(y|\theta)) + v((1 - \gamma)(p(y|\theta) - \tilde{p}(y|\theta))^2 + (1 - \gamma)v(p(y|\theta))) \quad (44)$$

and $\gamma = 1 - (1 - \epsilon)^D$; $p(y|\theta)$ as in (24); and $\tilde{p}(y|\theta)$ in (42). Furthermore, the norm of the bias in (38) can be bounded as

$$\|\text{bias}\|^2 \leq 4D\|H\|_\infty^2 \gamma := B^\mathcal{E}. \quad (45)$$

Proof: Proof is included in Appendix C. ■

The bound (45) on the bias of the estimator $\hat{g}_t^\mathcal{E}$ can be seen to be increasing with the noise level γ ; to be proportional to

the number of parameters D ; and to depend on the spectrum of the problem Hamiltonian. In contrast, the upper bound $V^\mathcal{E}$ in (43) quantifies the impact of shot noise and gate noise on the variance of the gradient estimator in (36). By (43), the variance of the gradient estimator decreases as $1/2N_m$ with respect to the total number of per-parameter measurements. Furthermore, unlike the bias, the variance $V^\mathcal{E}$ is not necessarily increasing with the noise level γ .

To see this, note that for a fixed number N_m of measurements, the impact of the gate noise on the variance is captured by the term $v(p^\mathcal{E}(y|\theta))$ in (44). In the absence of gate noise, i.e., when $\epsilon = \gamma = 0$, we have the equality $v(p^\mathcal{E}(y|\theta)) = v(p(y|\theta))$, whereby the variance $V^\mathcal{E}$ reduces to the variance V in (29) caused solely by shot noise. The variance $v(p^\mathcal{E}(y|\theta))$ is not necessarily monotonically increasing with the noise γ . In fact, as shown in Appendix D, there exists a $\gamma^*(y, \theta) \in [0, 1]$ such that the variance $v(p^\mathcal{E}(y|\theta))$ is a concave function of γ , increasing in the range $\gamma \in [0, \gamma^*(y, \theta)]$ and decreasing in the range $\gamma \in (\gamma^*(y, \theta), 1]$.

C. CONVERGENCE OF SGD

Using the results in the previous section, we now study the convergence properties of the SGD with the biased estimator $\hat{g}_t^\mathcal{E}$. As in the previous section, we evaluate the performance of the obtained solution θ^T after T iterations of the SGD update (26) by assuming the availability of a noiseless PQC for testing. In other words, we evaluate the final performance in terms of the loss $L(\theta^T)$ in (6). In practice, this requires the application of QEM for final testing (but not during the optimization phase). The following theorem illustrates the convergence properties of the biased gradient estimator.

Theorem IV.1: Under Assumption III.1, for any given initial point θ^0 , the following bound on the optimality gap holds for the SGD with the biased gradient estimator in (36), given any fixed learning rate $\eta_t = \eta \leq 1/\mathcal{L}$:

$$\mathbb{E}[L(\theta^T)] - L^* \leq (1 - \eta\mu)^T (\mathbb{E}[L(\theta^0)] - L^*) + \frac{1}{2} \left[\frac{B^\mathcal{E} + \eta\mathcal{L}V^\mathcal{E}}{\mu} \right] \quad (46)$$

where the expectations are taken over quantum measurements, and $V^\mathcal{E}$ and $B^\mathcal{E}$ are defined as shown in (43) and (45), respectively. Furthermore, given some target error level $\delta > B^\mathcal{E}/\mu > 0$, for learning rate $\eta = \eta^{\text{gate-noise}} \leq \min\{\frac{1}{\mathcal{L}}, \frac{\delta\mu}{\mathcal{L}V^\mathcal{E}}\}$, a number of iterations

$$T^{\text{gate-noise}} = \tilde{\mathcal{O}} \left(\log \frac{1}{\delta - \frac{B^\mathcal{E}}{\mu}} + \frac{V^\mathcal{E}}{\delta\mu} \right) \frac{\mathcal{L}}{\mu} \quad (47)$$

is sufficient to ensure the error floor $\mathbb{E}[L(\theta^T)] - L^* = \mathcal{O}(\delta)$.

As in the bound (33), which holds for $\epsilon = 0$, the contribution due to the variance term $V^\mathcal{E}$ in (46) can be made arbitrarily small by keeping the learning rate η sufficiently small. In contrast, the bias due to gate noise, which is quantified by the term $B^\mathcal{E}$ in (45), prevents bound (46) from vanishing. Therefore, according to the bound (46), the noise-induced bias causes a floor equal to $B^\mathcal{E}/\mu$ on the achievable error

δ . For error levels $\delta > B^{\mathcal{E}}/\mu$, by (47), the induced bias $B^{\mathcal{E}}$ entails a larger number of SGD iterations to converge.

V. IMPACT OF QUANTUM ERROR MITIGATION

As seen in the previous section, the quantum gate noise induces a bias in the gradient estimator (36). In this section, we first introduce the quasi-probabilistic error mitigation (QEM) technique introduced in [2]. We then study the convergence of SGD (12) when QEM is employed in evaluating the stochastic gradient estimator on the noisy PQC.

A. QUASI-PROBABILISTIC ERROR MITIGATION

QEM aims to recover the ideal expected value $\langle H \rangle_{|\Psi(\theta)\rangle} = \text{Tr}(HU^\theta(\rho_0))$ of observable H under the quantum state $U^\theta(\rho_0)$ produced by the noiseless PQC. As illustrated in Fig. 2, this is done via a quasi-probabilistic combination of the implementable operations in the set Ω (see Section II-C). To explain QEM, consider the d th ideal unitary map $U_d^{\theta_d}(\cdot)$. QEM expresses the ideal unitary operation $U_d^{\theta_d}(\cdot)$ as a linear combination of implementable operations in set Ω as

$$U_d^{\theta_d}(\cdot) = \sum_s q_{d,s} \mathcal{O}_s^{\theta_d}(\cdot), \quad \text{with } \mathcal{O}_s^{\theta_d} \in \Omega \quad (48)$$

where the *quasi-probabilities* $q_{d,s}$ are real numbers. Specifically, with the noise model (16) and (17) described in Section II-C, decomposition (48) applies with the operations $\mathcal{O}_s^{\theta_d}(\cdot) = \mathcal{E} \circ \mathcal{P}_s \circ U_d^{\theta_d}(\cdot)$, for $s = 0, \dots, 4^n - 1$, where \mathcal{P}_s is the string of Pauli operations as defined in (21) [2]. Note that we have identified for convenience the Pauli string \mathcal{P}_s with an integer $s \in \{0, 1, \dots, 4^n - 1\}$, rather than with a binary string $s \in \{0, 1, 2, 3\}^n$ as shown in Section II-D. This mapping between bit strings and integers is standard [37, Sec. 3.4].

With the noise model (16) and (17), (48) can be then expressed equivalently via the following *quasi-probabilistic representation* (QPR) [2]

$$U_d^{\theta_d}(\cdot) = Z_d \sum_{s=0}^{4^n-1} \text{sgn}(q_{d,s}) p_{d,s} \mathcal{O}_s^{\theta_d}(\cdot) \quad (49)$$

where $\mathcal{O}_s^{\theta_d}(\cdot) = \mathcal{E} \circ \mathcal{P}_s \circ U_d^{\theta_d}(\cdot)$ and

$$Z_d = \sum_{s=0}^{4^n-1} |q_{d,s}| \quad (50)$$

is the l_1 -norm of the vector q_d collecting all quasi-probabilities for the d th gate; and

$$p_{d,s} = \frac{|q_{d,s}|}{Z_d}, \quad \text{for } s = 0, \dots, 4^n - 1 \quad (51)$$

is a probability mass function satisfying $\sum_{s=0}^{4^n-1} p_{d,s} = 1$. By the trace-preserving property of the ideal and noisy operations, the quasi-probabilities satisfy the equality $\sum_{s=0}^{4^n-1} q_{d,s} = 1$ and the inequality $Z_d \geq 1$.

Using (49), QEM can ideally recover the overall unitary (15) implemented by the noiseless PQC as

$$U^\theta(\cdot) = Z \sum_{s^D} p_{s^D} \text{sgn}(q_{s^D}) \mathcal{O}_{s^D}^\theta(\cdot) \quad (52)$$

Algorithm 1: QEM via MC Sampling.

Input: Ideal unitary gates $\{U_d^{\theta_d}\}_{d=1}^D$, set of implementable operations Ω , and number of circuits N_c

Output: $\{\text{sgn}(q_{s_l^D})\}_{l=1}^{N_c}$, $\{\mathcal{O}_{s_l^D}^\theta\}_{l=1}^{N_c}$ and $\langle H \rangle_{\rho_{s_l^D}(\theta)}$

- 1: Using (48) obtain the quasi-probability vector $q_d = [q_{d,s}]_{s=0}^{4^n-1}$, for $d = 1, \dots, D$.
 - 2: Compute Z_d and Z using (50), and $p_{d,s}$ using (51)
 - 3: **for** $l = 1, \dots, N_c$ **do**
 - 4: **for** $d = 1, \dots, D$ **do**
 - 5: Sample $s_d \sim p_{d,s}$ to choose the parameterized gate $\mathcal{O}_{s_d}^{\theta_d}$
 - 6: **end for**
 - 7: Compute $\text{sgn}(q_{s_l^D}) = \prod_{d=1}^D \text{sgn}(q_{d,s_d})$
 - 8: Compute $\langle H_{\text{sgn}} \rangle_{\rho_{s_l^D}} = \text{sgn}(q_{s_l^D}) \text{Tr}(H \mathcal{O}_{s_l^D}^\theta(\rho_0))$
 - 9: **end for**
 - 10: **Return** $\{\text{sgn}(q_{s_l^D})\}_{l=1}^{N_c}$, $\{\mathcal{O}_{s_l^D}^\theta\}_{l=1}^{N_c}$ and $\langle H \rangle_{\rho_{s_l^D}(\theta)} = \frac{Z}{N_c} \sum_l \langle H_{\text{sgn}} \rangle_{\rho_{s_l^D}}$
-

where $s^D = (s_1, \dots, s_D)$ is a D -dimensional vector with $s_j \in \{0, \dots, 4^n - 1\}$ for $j = 1, \dots, D$; $\mathcal{O}_{s^D}^\theta(\cdot) = \mathcal{O}_{s_D}^{\theta_D} \circ \dots \circ \mathcal{O}_{s_1}^{\theta_1}(\cdot)$ is the corresponding noisy circuit composed of feasible operations from set Ω ; $Z = \prod_{d=1}^D Z_d$ is the product of the normalizing constants; $p_{s^D} = \prod_{d=1}^D p_{d,s_d}$ is the probability of choosing the s^D th circuit implementing operation $\mathcal{O}_{s^D}^\theta(\cdot)$; and we have $\text{sgn}(q_{s^D}) = \prod_{d=1}^D \text{sgn}(q_{d,s_d})$ with sgn being the sign function. Note that the sum in (52) is over all 4^{nD} values of string s^D . As illustrated in Fig. 2, QEM can thus exactly recover the true expected observable as

$$\text{Tr}(HU^\theta(\rho_0)) = Z \sum_{s^D} p_{s^D} \text{sgn}(q_{s^D}) \langle H \rangle_{\rho_{s^D}(\theta)} \quad (53)$$

where

$$\rho_{s^D}(\theta) = \mathcal{O}_{s^D}^\theta(\rho_0) \quad (54)$$

is the quantum state obtained by applying the s^D th noisy circuit $\mathcal{O}_{s^D}^\theta(\cdot)$; and $\langle H \rangle_{\rho_{s^D}(\theta)} = \text{Tr}(H \rho_{s^D}(\theta))$ is the expected observable under the s^D th quantum state.

Evaluating (53) becomes practically infeasible as the number n of qubits or the number D of unitary maps grows, since the number of terms in the sum (52) grow as 4^{nD} . Consequently, in practice, quasi-probabilistic error mitigation is implemented via a Monte Carlo sampling of N_c circuits $\{s_1^D, \dots, s_{N_c}^D\} := s_{1:N_c}^D$ in an i.i.d manner from the distribution p_{s^D} [2], [38]. This, in turn, gives an unbiased estimator of the QEM-mitigated expected observable (53) as

$$\langle H \rangle_{\rho_{s_{1:N_c}^D}(\theta)} = Z \sum_{l=1}^{N_c} \frac{1}{N_c} \text{sgn}(q_{s_l^D}) \langle H \rangle_{\rho_{s_l^D}(\theta)}. \quad (55)$$

We refer to Algorithm 1 for a summary of QEM.

B. QEM-BASED GRADIENT ESTIMATOR

In this section, we describe the QEM-based gradient estimator used to mitigate the bias due to gate noise. At each iteration t , the QEM-based gradient estimator \hat{g}_t^{QEM} is obtained in two steps. In the first step, the gradient estimator samples N_c noisy circuits $\{\mathcal{O}_{s_l^D}^{\theta^{t-1}}\}_{l=1}^{N_c}$, with the current parameter vector θ^{t-1} . This is done by using the distribution $p_{s,D}$ as described in Algorithm 1. For each l th sampled noisy circuit, the gradient estimator then applies the parameter-shift rule as

$$[\hat{g}_{t,l}^{\text{QEM}}]_d = \frac{1}{2} \left(\widehat{\langle H \rangle}_{\rho_{s,D}(\theta^{t-1} + \frac{\pi}{2} e_d)} - \widehat{\langle H \rangle}_{\rho_{s,D}(\theta^{t-1} - \frac{\pi}{2} e_d)} \right) \quad (56)$$

with $l = 1, \dots, N_c$. In (56), the term $\widehat{\langle H \rangle}_{\rho_{s,D}(\theta)}$ denotes an empirical estimator of the expected observable $\langle H \rangle_{\rho_{s,D}(\theta)}$, obtained from N_{QEM} i.i.d. measurements $\{H_1^D, \dots, H_{N_{\text{QEM}}}^D\}$ of observable H under the sampled noisy quantum state $\rho_{s,D}(\theta)$. This yields the output $H_j^D = h_y$ with probability

$$p_{s,D}(y|\theta) = \text{Tr}(\Pi_y \rho_{s,D}(\theta)). \quad (57)$$

Finally, the d th component of the gradient estimator, $[\hat{g}_t^{\text{QEM}}]_d$ is obtained as

$$[\hat{g}_t^{\text{QEM}}]_d = \frac{Z}{N_c} \sum_{l=1}^{N_c} \text{sgn}(q_{s_l^D}) [\hat{g}_{t,l}^{\text{QEM}}]_d \quad (58)$$

by averaging the product of the per-circuit gradient estimation $[\hat{g}_{t,l}^{\text{QEM}}]_d$ in (56) over the N_c sampled circuits, and multiplying by the normalizing constant Z . The QEM-based gradient estimator is described in Algorithm 2.

To enable ease of comparison of the convergence behavior of the QEM-based gradient estimator with the gradient estimators described in previous sections, we fix the measurement budget to N_m shots for the total of N_c circuits. In other words, we fix $N_{\text{QEM}} = N_m/N_c$ number of measurement shots on each sampled circuit.

C. PROPERTIES OF THE GRADIENT ESTIMATE

The QEM-based gradient estimator of (58) has two sources of randomness, namely the randomness arising from sampling noisy circuits $\{\mathcal{O}_{s_l^D}^{\theta}\}_{l=1}^{N_c}$, as well as that arising from the quantum measurements. Accordingly, we can decompose the QEM-based gradient estimator in (58) as

$$\hat{g}_t^{\text{QEM}} = g_t + \xi_t^{\text{QEM}} \quad (59)$$

where ξ_t^{QEM} denotes the total noise due to shot noise and circuit sampling noise that satisfies $\mathbb{E}[\xi_t^{\text{QEM}}] = 0$ and $\text{var}(\xi_t^{\text{QEM}}) = \mathbb{E}[\|\hat{g}_t^{\text{QEM}} - g_t\|^2]$, with the expectation taken over the randomly sampled noisy circuits as well as over quantum measurements. It can be seen that the following equality holds (see proof in Appendix E)

$$\text{var}(\xi_t^{\text{QEM}}) = \mathbb{E}[\|\hat{g}_t^{\text{QEM}} - g_t^{\text{circ}}\|^2] + \mathbb{E}[\|g_t^{\text{circ}} - g_t\|^2] \quad (60)$$

Algorithm 2: QEM-Based SGD.

Input: Initialization θ^0 , number of iterations T , number of circuit samples N_c , learning rate $\{\eta_t\}$

Output: Final iterate θ^T

```

1: Set  $t = 1$ 
2: while  $t \leq T$  do
3:   Set  $\theta = \theta^{t-1}$ 
4:   Get  $\{\text{sgn}(q_{s_l^D})\}_{l=1}^{N_c}, \{\mathcal{O}_{s_l^D}^{\theta}\}_{l=1}^{N_c}$  from Algorithm 1
5:   for  $d = 1, \dots, D$  do
6:     Set  $g_d = [\cdot]$ 
7:     for  $l = 1, \dots, N_c$  do
8:       Implement  $\mathcal{O}_{s_l^D}^{\theta \pm \frac{\pi}{2} e_d}(\rho_0)$ 
9:       Compute  $[\hat{g}_{t,l}^{\text{QEM}}]_d$  using (56)
10:      Set  $g_d[l] = \text{sgn}(q_{s_l^D})[\hat{g}_{t,l}^{\text{QEM}}]_d$ 
11:    end for
12:    Compute  $[\hat{g}_t^{\text{QEM}}]_d = \frac{Z}{N_c} \sum_{l=1}^{N_c} g_d[l]$ 
13:  end for
14:  Update parameter vector as  $\theta^t \leftarrow \theta - \eta_t \hat{g}_t^{\text{QEM}}$ 
15:  Update  $t \leftarrow t + 1$ 
16: end while
17: Return final parameter iterate  $\theta^T$ 

```

where

$$[g_t^{\text{circ}}]_d = \frac{Z}{2N_c} \sum_{l=1}^{N_c} \text{sgn}(q_{s_l^D}) \times \left(\langle H \rangle_{\rho_{s,D}(\theta^t + \frac{\pi}{2} e_d)} - \langle H \rangle_{\rho_{s,D}(\theta^t - \frac{\pi}{2} e_d)} \right) \quad (61)$$

is the shot-noise free estimate of the gradient obtained via N_c noisy sampled circuits. While the first term of (60) captures the impact of finite number, $N_{\text{QEM}} = N_m/N_c$, of measurements made per sampled noisy circuit, the second term captures the impact of sampling finite number of noisy circuits on the variance of the QEM-based gradient estimator. With this insight, the following theorem provides bounds on the variance of the QEM-based gradient estimator.

Theorem V.1: The following inequality holds for the variance of the QEM-based gradient estimator:

$$c_1(\gamma) \text{var}(\xi_t^{\mathcal{E}}) \leq \text{var}(\xi_t^{\text{QEM}}) \leq V^{\text{QEM}} \quad (62)$$

where $\text{var}(\xi_t^{\mathcal{E}})$ defined as in (43); we have $c_1(\gamma) \geq 0$; and

$$V^{\text{QEM}} = \frac{N_h D Z^2 \text{Tr}(H^2)}{2N_m} \sup_{\theta \in \mathbb{R}^D, y \in \{1, \dots, N_h\}} \mathbb{E}_{s,D}[v(p_{s,D}(y|\theta))] + \frac{Z^2 D \|H\|_{\infty}^2}{N_c} \quad (63)$$

satisfies the inequality

$$V^{\text{QEM}} \geq c_1(\gamma) V^{\mathcal{E}} + \frac{c_2(\gamma) D \|H\|_{\infty}^2}{N_c} \quad (64)$$

where $c_2(\gamma) \geq 1$ and $p_{s,D}(y|\theta)$ is defined as in (57). Furthermore, in the special case of the depolarizing channel, the

inequality $c_1(\gamma) \geq 1$ is satisfied, and both $c_1(\gamma)$ and $c_2(\gamma)$ are nondecreasing functions of γ with $c_1(0) = c_2(0) = 1$.

Proof: Proof can be found in Appendix E. ■

Theorem V.1 highlights three key points. First, as can be seen from (62), the variance of the QEM-based gradient estimator is at least $c_1(\gamma)$ times the variance of the unmitigated gradient estimator. For depolarizing noise, the function $c_1(\gamma) \geq 1$ is nondecreasing with γ , implying that QEM-based gradient estimator has larger variance than the biased noisy gradient estimator.

Second, the upper bound V^{QEM} in (63) on the variance of the QEM-based gradient estimator comprises the contribution of shot noise, as captured by the first term, as well as of the circuit sampling noise, as captured by the second term. Accordingly, the bound (63) suggests that, even when an infinite number of circuits is sampled, i.e., when $N_c \rightarrow \infty$, the finite number of per-circuit measurements results in a nonzero variance. Finally, inequality (64) relates the bound V^{QEM} of the variance with QEM to the bound $V^\mathcal{E}$ obtained without QEM, as defined in (43). Noting that functions $c_1(\gamma)$ and $c_2(\gamma)$ are nondecreasing with γ for depolarizing noise, the inequality in (64) suggests that the variance of the QEM-based estimator generally increases with gate noise.

D. CONVERGENCE OF QEM-BASED SGD

Based on the analysis of the stochastic gradient in the previous section, we have the following convergence result for SGD using the QEM-based gradient estimator \hat{g}_t^{QEM} .

Theorem V.2: Under Assumption III.1, for any given initial point θ^0 , the following bound on the optimality gap holds for the SGD employing QEM-based gradient estimator in (58), given any fixed learning rate $\eta_t = \eta \leq 1/\mathcal{L}$

$$\begin{aligned} \mathbb{E}[L(\theta^T)] - L^* &\leq (1 - \eta\mu)^T (\mathbb{E}[L(\theta^0)] - L^*) \\ &\quad + \frac{1}{2} \left[\frac{\eta\mathcal{L}V^{\text{QEM}}}{\mu} \right] \end{aligned} \quad (65)$$

where the expectations are taken over quantum measurements and noisy circuit samples, and V^{QEM} is defined as in (62). Furthermore, given some target error level $\delta > 0$, for learning rate $\eta = \eta^{\text{QEM}} \leq \min\{\frac{1}{\mathcal{L}}, \frac{\delta\mu}{\mathcal{L}V^{\text{QEM}}}\}$, a number of iterations

$$T^{\text{QEM}} = \tilde{\mathcal{O}} \left(\log \frac{1}{\delta} + \frac{V^{\text{QEM}}}{\delta\mu} \right) \frac{\mathcal{L}}{\mu} \quad (66)$$

is sufficient to ensure the error $\mathbb{E}[L(\theta^T)] - L^* = \mathcal{O}(\delta)$.

Theorem V.2 can be used to compare the convergence of QEM-based SGD with that of the SGD under shot and gate noise in Theorem IV.1. While the presence of quantum gate noise forces SGD to settle at an error floor of $\mathcal{O}(\delta)$ with $\delta > B^\mathcal{E}/\mu$, QEM can achieve any error floor, $\mathcal{O}(\delta)$ for $\delta > 0$, in a finite number T^{QEM} of iterations. We refer to Section I for further discussion on this result.

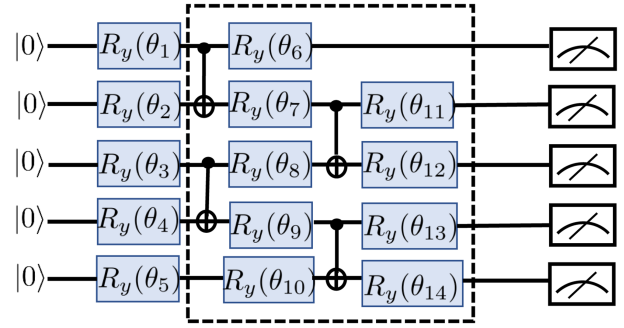


FIGURE 4. Hardware-efficient ansatz used in the experiments. The gates in the dashed box are repeated L times.

VI. EXPERIMENTS

In this section, we present numerical results concerning the VQE to solve a weighted max-cut combinatorial optimization problem.

A. WEIGHTED MAX-CUT HAMILTONIAN

Consider an undirected graph $\mathcal{G} = (V, E)$ with $V = \{1, \dots, n\}$ denoting the set of vertices of the graph and $E \subseteq V \times V$ denoting the set of edges. Each edge $(i, j) \in E$ has an associated weight $w_{i,j} > 0$ such that $w_{i,j} = w_{j,i}$. A cut of the graph defines a partition of the vertices into two distinct subsets. Specifically, a cut assigns a binary variable $x_v \in \{0, 1\}$ to each vertex $v \in V$ depending on whether it belongs to one subset or to the other. In the weighted max-cut optimization problem, the goal is to find the cut that maximizes the sum of the weights of the edges that connect the vertices belonging to the two distinct subsets, i.e., of the edges crossing the cut. This corresponds to maximizing the cost function

$$C(\mathbf{x}) = \sum_{(i,j) \in E} w_{i,j} x_i (1 - x_j) + \sum_i w_{i,i} x_i \quad (67)$$

over the binary vector $\mathbf{x} = (x_1, \dots, x_n)$ with $x_i \in \{0, 1\}$ for all $i = 1, \dots, n$. Note that (67) also imposes an additional penalty on the self-weights assigned to each vertices via the second summation.

The objective function in (67) can be converted to an Ising Hamiltonian via the mapping $x_i \mapsto (1 - Z_i)/2$ where $Z_i = (I \otimes \dots \otimes I \otimes Z \otimes I \otimes \dots \otimes I)$ denotes an n -qubit operator that applies a Pauli-Z gate on the i th qubit. The weighted max-cut problem of maximizing the cost in (67) over \mathbf{x} can be then equivalently expressed as minimizing the expected value $\langle H \rangle_{|\Psi(\theta)\rangle}$ of the Ising Hamiltonian

$$H = \sum_{i=1}^n w_{i,i} z_i + \sum_{i < j} w_{i,j} z_i z_j \quad (68)$$

over the quantum state $|\Psi(\theta)\rangle$. To produce the state $|\Psi(\theta)\rangle = U(\theta)|0\rangle$, we consider the hardware-efficient ansatz adopted in [8], as shown in Fig. 4. Accordingly, the PQC comprises of single-qubit Pauli Y -rotation gates $R_y(\theta)$, as well as two-qubit CNOT gates. The dashed box in Fig. 4 represents a layer




FIGURE 5. Convergence analysis of the SGD—when exact gradients can be computed, when only shot-noise is present, when shot and gate noise are present, and when QEM is employed—as a function of the number of iterations for $n = 3$ qubit system subjected to depolarizing noise on the CNOT gates. The learning rate is $\eta_t = 0.4/t$; $N_c = 8$; $N_m = 400$; $L = 2$ and (left) $\epsilon = 0.02$ and (right) $\epsilon = 0.1$.

of the PQC, which can be repeated to obtain deep variational quantum circuits.

B. RESULTS

We first consider a simple complete graph \mathcal{G} with $n = 3$ vertices and a random 3×3 weight matrix $w = [0.41, 0.44, 0.55; 0.44, 0.97, 0.22; 0.55, 0.22, 0.89]$ whose (i, j) th element $w_{i,j}$ denotes the weight of the edge $(i, j) \in E$. We consider the PQC defined by the hardware-efficient ansatz in Fig. 4 with two layers, i.e., with $L = 2$. The PQC is initialized to the state $|+\rangle^{\otimes n}$ by setting the last rotation angle on each qubit to $\pi/2$ and setting all other rotation angles to zero [8]. We assume that depolarizing noise, with error probability ϵ , acts only on the CNOT gates, resulting in a PQC with eight noisy gates.

We use the open-source quantum computing framework Qiskit [39] to simulate our experiments. In particular, the noise model in our numerical simulations is constructed by making use of the *NoiseModel* class available in the *Aer* library of Qiskit. We use the function *add_quantum_error* to add depolarizing two-qubit error on each of the CNOT gates with desired error strength. The assumption that the depolarizing noise acts only on CNOT gates is motivated by the experimental observations that noisy CNOT gates yield larger error in measurement outcomes as compared to noisy single-qubit gates [28].

In Fig. 5, we study the convergence of SGD when exact gradients can be computed; when only shot noise is present, with number of measurements $N_m = 400$; when both shot and gate noise are present, with the latter modeled by the mentioned depolarizing noise with error probability ϵ ; and when QEM is used to combat gate noise with number of sampled noisy circuits $N_c = 8$. We compare the figure on the left, corresponding to a smaller noise of $\epsilon = 0.02$, with that on the right with a larger noise level of $\epsilon = 0.1$. The error levels are chosen to be comparable with the experiments in [28] and [40].

For each of the three scenarios, namely shot noise, shot+gate noise, and QEM, we conduct ten experiments with

the SGD iterations starting from the same fixed initial point θ^0 to obtain a set of parameter iterates. In each experiment, the cost $\langle H \rangle_{|\Psi(\theta^t)\rangle}$ corresponding to the obtained parameter iterate θ^t , for $t = 1, \dots, T$, is evaluated exactly (i.e., $N_m \rightarrow \infty$). The mean of the costs over the experiments is indicated by the bold curve, while the spread is indicated by the lighter shadows.

Overall, the experimental findings of Fig. 5 corroborate our theoretical analysis. In particular, by comparing the two figures in Fig. 5, we observe that a larger strength of the depolarizing noise causes a larger error floor for SGD in the presence of shot and gate noise. In contrast, QEM achieves a lower error floor under both low and high noise strengths. When the noise strength is smaller and a sufficient number of circuits are sampled as in Fig. 5(left), the variance of the QEM-based estimator is reduced (see Table 2), thereby ensuring convergence at a rate comparable to the shot-noise only case. In contrast, when the noise strength is larger, by Theorem V.1, QEM requires more circuits to be sampled per iteration.

In Fig. 6, we study the exact loss $L(\theta^T)$ (evaluated with $N_m \rightarrow \infty$) of the SGD iterate after $T = 10$ iterations when exact gradient can be computed; when shot-noise is present, with $N_m = 10240$ measurement shots; when shot and gate noise are present; and when QEM is employed to mitigate the bias, as a function of the increasing noise level ϵ . The figure in the left corresponds to smaller number, $N_c = 7$, of circuit samples, while the figure on the right corresponds to $N_c = 10$. The max-cut problem has $n = 5$ vertices and a random 5×5 weight matrix $w = [0.42, 0.43, 0.55, 0.96, 0.22; 0.44, 0.89, 0.07, 0.87, 0.01; 0.55, 0.07, 0.77, 0.18, 0.15; 0.96, 0.87, 0.18, 0.77, 0.51; 0.22, 0.01, 0.15, 0.51, 0.84]$. As before, we use the hardware-efficient ansatz in Fig. 4, with $L = 1$ layer and depolarizing noise assumed to act only on CNOT gates. We fix the learning rate as $\eta_t = 0.14$.

It can be seen from Fig. 6 that the presence of gate noise induces a significant bias as the strength of the noise increases. For large noise level, QEM successfully lowers the




FIGURE 6. Cost function $L(\theta^T)$ after $T = 10$ iterations of the SGD when only shot noise is present, when both shot and gate noise are present, and when QEM is employed, as a function of the noise level ϵ . (Left) Circuit samples $N_c = 7$ and (right) $N_c = 10$. Other parameters are set as $n = 5$, $L = 1$, $N_m = 10240$, and $\eta_t = 0.14$. Ground state eigenvalue for the max-cut problem is -3.24 .

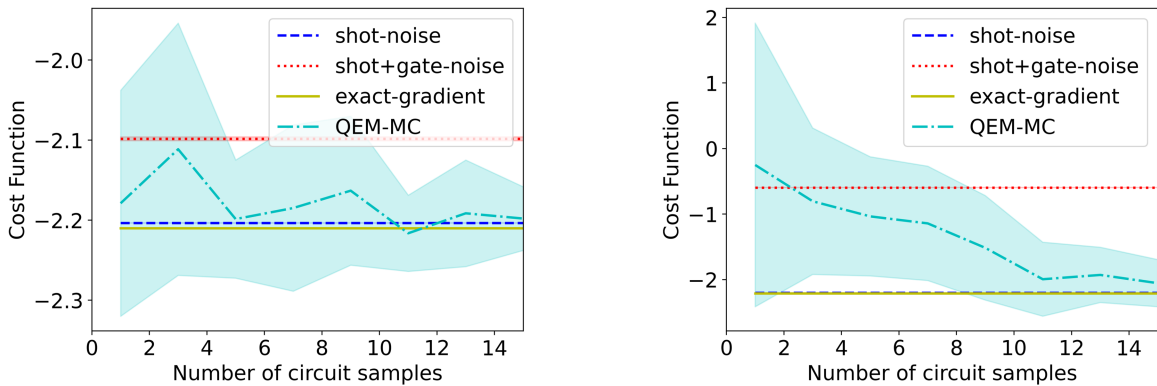


FIGURE 7. Cost function $L(\theta^T)$ after $T = 10$ iterations of the SGD when only shot noise is present, when both shot and gate noise are present, and when QEM is employed, as a function of the number of circuit samples N_c . (Left) Low noise level $\epsilon = 0.03$ and (right) high noise level ($\epsilon = 0.25$). Other parameters are set as $n = 5$, $L = 1$, $N_m = 10240$, and $\eta_t = 0.14$. Ground state eigenvalue for the max-cut problem is -3.24 .

bias induced by gate noise. However, when the noise strength is small, comparing the figures in the left and right shows that the higher variance of the QEM, resulting due to small N_c , may offset the benefit of QEM. Finally, Fig. 7 demonstrates the impact of increasing number of circuit samples on the loss $L(\theta^T)$ of the SGD iterate after $T = 10$ iterations. The experimental setting is same as in Fig. 6. As can be seen from the figure, increasing the number of circuit samples significantly reduces the variance of the QEM and ensure performance close to the SGD with perfect error mitigation (blue curve).

VII. CONCLUSION

This article seeks an answer to the following question: Can quasi-probabilistic error mitigation (QEM) be beneficial in improving the convergence of SGD for the implementation of VQEs in NISQ devices? By analyzing the convergence properties of SGD for VQEs, we have shown that the quantum gate noise inherent in the operation of NISQ devices results in a nonzero error floor. Achieving lower error levels requires the use of QEM. However, for larger error levels, the increase in the variance of the stochastic gradient estimate caused by QEM may entail the need for a larger number of SGD iterations when QEM is deployed. In particular, when

the noise strength is high and insufficient time is available at each SGD iteration to sample circuits for QEM, QEM may require a larger number of iterations to converge to the desired error level. Conversely, when the number of circuits sampled per iteration is sufficiently large, QEM can yield a significant reduction in the number of SGD iterations as compared to a conventional system without error mitigation.

APPENDIX A PROOF OF LEMMA 3.1

To prove Lemma 3.1, we fix the iteration index t , which is dropped from the notation. Let $X_{d,\pm} = \widehat{\langle H \rangle}_{|\Psi(\theta \pm \frac{\pi}{2} e_d)\rangle} - \langle H \rangle_{|\Psi(\theta \pm \frac{\pi}{2} e_d)\rangle}$ denote the difference between the estimated and true expectation of the observable H under the quantum state $|\Psi(\theta \pm \frac{\pi}{2} e_d)\rangle$ whose d th parameter is phase-shifted by $\pm\pi/2$. For brevity, we also henceforth use the notation $|\Psi_{d,\pm}\rangle = |\Psi(\theta \pm e_d \frac{\pi}{2})\rangle$. The variance of the gradient estimate (26) can be written as

$$\text{var}(\xi) = \mathbb{E} \left[\sum_{d=1}^D \left(\frac{1}{2} \left(\widehat{\langle H \rangle}_{|\Psi_{d,+}\rangle} - \widehat{\langle H \rangle}_{|\Psi_{d,-}\rangle} \right) - \frac{1}{2} \left(\langle H \rangle_{|\Psi_{d,+}\rangle} - \langle H \rangle_{|\Psi_{d,-}\rangle} \right) \right)^2 \right]$$

$$\begin{aligned}
 &= \sum_{d=1}^D \frac{1}{4} \mathbb{E} \left[(X_{d,+} - X_{d,-})^2 \right] \\
 &= \sum_{d=1}^D \frac{1}{4} \left(\mathbb{E} \left[X_{d,+}^2 \right] + \mathbb{E} \left[X_{d,-}^2 \right] \right) \quad (69)
 \end{aligned}$$

where the expectation is with respect to the N_m measurements of the quantum state $|\Psi(\theta + e_d \frac{\pi}{2})\rangle$ and the N_m measurements of the quantum state $|\Psi(\theta - e_d \frac{\pi}{2})\rangle$ for $d = 1, \dots, D$. The random variables $X_{d,+}$ and $X_{d,-}$ are, thus, independent for $d = 1, \dots, D$, which results in the equality in (69).

The expectation $\mathbb{E}[X_{d,+}^2]$ is equal to the variance $\text{var}(\widehat{H}_{|\Psi_{d,+}\rangle})$ of the random variable $\widehat{H}_{|\Psi_{d,+}\rangle}$. Let Y be the random variable that defines the index of the measurement of the observable H , and let us write as $H = h_Y$ for the corresponding measurement output. We denote as $W_y = \mathbb{I}\{Y = y\}$, for $y = 1, \dots, N_h$, the Bernoulli random variable determining whether $Y = y$ ($W_y = 1$) or not ($W_y = 0$). Noting that the quantum measurements are i.i.d., it follows from the definition of expectation $\widehat{H}_{|\Psi_{d,+}\rangle}$ in (27) that

$$\begin{aligned}
 \mathbb{E} \left[X_{d,+}^2 \right] &= \frac{1}{N_m} \text{var} \left(\sum_{y=1}^{N_h} h_y W_y \right) \quad (70) \\
 &= \frac{1}{N_m} \mathbb{E} \left[\left(\sum_{y=1}^{N_h} h_y (W_y - p(y|\theta + e_d \frac{\pi}{2})) \right)^2 \right] \quad (71)
 \end{aligned}$$

$$\stackrel{(a)}{\leq} \frac{1}{N_m} \left(\sum_{y=1}^{N_h} h_y^2 \right) \sum_{y=1}^{N_h} \text{var}(W_y) \quad (72)$$

$$\begin{aligned}
 &\stackrel{(b)}{=} \frac{1}{N_m} \left(\sum_{y=1}^{N_h} h_y^2 \right) \sum_{y=1}^{N_h} v \left(p \left(y \left| \theta + e_d \frac{\pi}{2} \right. \right) \right) \\
 &\leq \frac{N_h}{N_m} \left(\sum_{y=1}^{N_h} h_y^2 \right) v = \frac{N_h \text{Tr}(H^2)}{N_m} v \quad (73)
 \end{aligned}$$

where (a) follows from the Cauchy–Schwarz inequality; (b) follows since the variance of the Bernoulli random variable W_y can be computed as

$$\text{var}(W_y) = \mathbb{E}[W_y^2] - (\mathbb{E}[W_y])^2 = v \left(p \left(y \left| \theta + e_d \frac{\pi}{2} \right. \right) \right)$$

with $v(x) = x(1 - x)$ for $x \in (0, 1)$. The last inequality follows from the definition of the quantity v in (30).

In a similar way, it can be shown that the following inequality holds:

$$\mathbb{E}[X_{d,-}^2] \leq \frac{N_h \text{Tr}(H^2)}{N_m} v. \quad (74)$$

Using these in (69) then yields the inequality

$$\text{var}(\xi) \leq \frac{DN_h \text{Tr}(H^2)v}{2N_m} \quad (75)$$

concluding the proof.

APPENDIX B SMOOTHNESS OF LOSS FUNCTION $L(\theta)$

In the following lemma, we show that the loss function $L(\theta)$ in (6) is indeed always \mathcal{L} -smooth for a constant \mathcal{L} that depends on the number of PQC parameters D , and on the observable H . In contrast, the loss function $L(\theta)$ need not necessarily satisfy the PL-condition.

Lemma B.1: The loss function $L(\theta)$ as defined in (6) is \mathcal{L} -smooth with $\mathcal{L} = D^{3/2} \sum_{y=1}^{N_h} |h_y|$.

Proof: To prove this result, we note that the smoothness condition can be equivalently written as the inequality $\|\nabla^2 L(\theta)\|_2 \leq \mathcal{L}$ on the Hessian $\nabla^2 L(\theta)$. We then note the following steps:

$$\left[\nabla^2 L(\theta) \right]_{i,j} = \frac{\partial^2 L(\theta)}{\partial \theta_i \partial \theta_j} \quad (76)$$

$$\begin{aligned}
 &\stackrel{(a)}{=} \frac{1}{4} \text{Tr} \left[H \left(\Psi \left(\theta + e_i \frac{\pi}{2} + e_j \frac{\pi}{2} \right) \right. \right. \\
 &\quad \left. \left. - \Psi \left(\theta + e_i \frac{\pi}{2} - e_j \frac{\pi}{2} \right) \right. \right. \\
 &\quad \left. \left. - \Psi \left(\theta - e_i \frac{\pi}{2} + e_j \frac{\pi}{2} \right) \right. \right. \\
 &\quad \left. \left. + \Psi \left(\theta - e_i \frac{\pi}{2} - e_j \frac{\pi}{2} \right) \right) \right]
 \end{aligned}$$

$$\begin{aligned}
 &\stackrel{(b)}{\leq} \frac{1}{4} \left| \text{Tr} \left[H \left(\Psi \left(\theta + e_i \frac{\pi}{2} + e_j \frac{\pi}{2} \right) \right. \right. \right. \\
 &\quad \left. \left. - \Psi \left(\theta + e_i \frac{\pi}{2} - e_j \frac{\pi}{2} \right) \right. \right. \\
 &\quad \left. \left. - \Psi \left(\theta - e_i \frac{\pi}{2} + e_j \frac{\pi}{2} \right) \right. \right. \\
 &\quad \left. \left. + \Psi \left(\theta - e_i \frac{\pi}{2} - e_j \frac{\pi}{2} \right) \right) \right] \right|
 \end{aligned}$$

$$\begin{aligned}
 &\stackrel{(c)}{\leq} \frac{\|H\|_\infty}{4} \left\| \Psi \left(\theta + e_i \frac{\pi}{2} + e_j \frac{\pi}{2} \right) \right. \\
 &\quad \left. - \Psi \left(\theta + e_i \frac{\pi}{2} - e_j \frac{\pi}{2} \right) \right. \\
 &\quad \left. - \Psi \left(\theta - e_i \frac{\pi}{2} + e_j \frac{\pi}{2} \right) \right. \\
 &\quad \left. + \Psi \left(\theta - e_i \frac{\pi}{2} - e_j \frac{\pi}{2} \right) \right\|_1 \quad (77)
 \end{aligned}$$

$$\begin{aligned}
 &\stackrel{(d)}{\leq} \frac{\|H\|_\infty}{4} \left(\left\| \Psi \left(\theta + e_i \frac{\pi}{2} + e_j \frac{\pi}{2} \right) \right. \right. \\
 &\quad \left. \left. - \Psi \left(\theta + e_i \frac{\pi}{2} - e_j \frac{\pi}{2} \right) \right\|_1 \right. \\
 &\quad \left. + \left\| -\Psi \left(\theta - e_i \frac{\pi}{2} + e_j \frac{\pi}{2} \right) \right. \right.
 \end{aligned}$$

$$\begin{aligned}
 & + \Psi \left(\theta - e_i \frac{\pi}{2} - e_j \frac{\pi}{2} \right) \Big\|_1 \\
 & \stackrel{(e)}{\leq} \|H\|_\infty \leq \sum_{y=1}^{N_h} |h_y|. \quad (78)
 \end{aligned}$$

The equality in (a) follows from a double application of parameter shift rule [8]; the inequality in (b) follows since the inequality $x \leq |x|$ holds for $x \in \mathbb{R}$; (c) follows from tracial matrix Hölder's inequality, i.e., $|\text{Tr}(A^\dagger B)| \leq \|A\|_\infty \|B\|_1$ [41]; (d) follows from the triangle inequality of the trace norm, $\|A\|_1 = \sqrt{\text{Tr}(AA^\dagger)}$ [42]; inequality (e) follows since the trace norm of a density matrix is bounded as $0 \leq \|\rho\|_1 \leq 2$ [42]; and the last inequality follows from triangle inequality and the inequality $\|\Pi_y\|_\infty \leq 1$, i.e., $\|H\|_\infty = \|\sum_{y=1}^{N_h} h_y \Pi_y\|_\infty \leq \sum_{y=1}^{N_h} |h_y| \|\Pi_y\|_\infty \leq \sum_{y=1}^{N_h} |h_y|$. This in turn implies the inequalities $\|\nabla^2 L(\theta)\|_2 \leq \sqrt{D} \|\nabla^2 L(\theta)\|_\infty \leq D^{3/2} \sum_{y=1}^{N_h} |h_y|$. ■

In contrast to the smoothness assumption, the convexity of the loss function $L(\theta)$ in Assumption III.1 is not always satisfied. It provides a useful working condition for the analysis, which was also adopted in [10] and [34]. Convexity may be recovered by applying the optimization approach for PQC detailed in [43]. We leave an investigation of this idea to future works.

APPENDIX C PROOF OF LEMMA 4.1

We first obtain an upper bound on the variance $\text{var}(\xi_t^\mathcal{E}) = \mathbb{E}[\|\hat{g}_t^\mathcal{E} - g_t^\mathcal{E}\|^2]$ of the gradient estimator. As shown in Appendix A, we drop the index t and define $X_{d,\pm} = \langle \widehat{H} \rangle_{\rho^\mathcal{E}(\theta \pm e_d \frac{\pi}{2})} - \langle H \rangle_{\rho^\mathcal{E}(\theta \pm e_d \frac{\pi}{2})}$ as the difference between the estimated and true expectation of the observable H under the noisy quantum state $\rho^\mathcal{E}(\theta \pm \frac{\pi}{2} e_d)$. By (69), we have the equality

$$\text{var}(\xi^\mathcal{E}) = \sum_{d=1}^D \frac{1}{4} \left(\mathbb{E}[X_{d,+}^2] + \mathbb{E}[X_{d,-}^2] \right) \quad (79)$$

where $\mathbb{E}[X_{d,\pm}^2] = \text{var}(\langle \widehat{H} \rangle_{\rho^\mathcal{E}(\theta \pm e_d \frac{\pi}{2})})$. Recall from (27) the following definition:

$$\langle \widehat{H} \rangle_{\rho^\mathcal{E}(\theta + e_d \frac{\pi}{2})} = \sum_{y=1}^{N_h} \frac{h_y}{N_m} \sum_{j=1}^{N_m} \mathbb{I}\{Y_j^\mathcal{E} = y\} \quad (80)$$

of estimated expectation of the observable under noisy state $\rho^\mathcal{E}(\theta + e_d \frac{\pi}{2})$ holds, where $Y_j^\mathcal{E}$ is the random variable corresponding to the index of the j th measurement of the observable H . In a manner similar to (73), it can be shown that the following upper bound holds:

$$\mathbb{E}[X_{d,\pm}^2] \leq \frac{N_h \text{Tr}(H^2)}{N_m} \sup_{\theta \in \mathbb{R}^D, y \in \{1, \dots, N_h\}} v(p^\mathcal{E}(y|\theta)) \quad (81)$$

where $p^\mathcal{E}(y|\theta)$ is defined in (35). Using (81) in (79) yields the upper bound (43).

To obtain the relation in (44), we use the decomposition (40) of the noisy quantum state $\rho^\mathcal{E}(\theta)$ as a convex combination of the noiseless ideal state $\Psi(\theta)$ and of the error density matrix $\tilde{\rho}(\theta)$, defined in (41). With the definition (41), the following set of relation holds for any $y \in \{1, \dots, N_h\}$ and $\theta \in \mathbb{R}^D$:

$$\begin{aligned}
 v(p^\mathcal{E}(y|\theta)) &= ((1-\gamma)p(y|\theta) + \gamma\tilde{p}(y|\theta)) \\
 &\quad \times (1 - (1-\gamma)p(y|\theta) - \gamma\tilde{p}(y|\theta)) \\
 &= (1-\gamma)p(y|\theta)(1 - (1-\gamma)p(y|\theta)) \\
 &\quad + \gamma\tilde{p}(y|\theta)(1 - (1-\gamma)p(y|\theta)) \\
 &\quad - \gamma(1-\gamma)p(y|\theta)\tilde{p}(y|\theta) - \gamma^2\tilde{p}(y|\theta)^2 \\
 &= (1-\gamma)v(p(y|\theta)) + \gamma \left[(1-\gamma)p^2(y|\theta) \right. \\
 &\quad \left. + \tilde{p}(y|\theta)(1 - (1-\gamma)p(y|\theta)) \right. \\
 &\quad \left. - (1-\gamma)p(y|\theta)\tilde{p}(y|\theta) - \gamma\tilde{p}(y|\theta)^2 \right] \\
 &= (1-\gamma)v(p(y|\theta)) + \\
 &\quad \gamma \left[(1-\gamma)(p(y|\theta) - \tilde{p}(y|\theta))^2 + \tilde{p}(y|\theta) - \tilde{p}(y|\theta)^2 \right] \\
 &= (1-\gamma)v(p(y|\theta)) \\
 &\quad + v((1-\gamma)(p(y|\theta) - \tilde{p}(y|\theta))^2 + \gamma v(\tilde{p}(y|\theta))).
 \end{aligned}$$

We now analyze the bias term. By (38), we have the equality

$$\begin{aligned}
 \|\text{bias}\|^2 &= \sum_{d=1}^D \frac{1}{4} \left(\text{Tr} \left(H \left(\rho^\mathcal{E} \left(\theta + e_d \frac{\pi}{2} \right) \right. \right. \right. \\
 &\quad \left. \left. - \rho^\mathcal{E} \left(\theta - e_d \frac{\pi}{2} \right) - \Psi \left(\theta + e_d \frac{\pi}{2} \right) \right. \right. \\
 &\quad \left. \left. + \Psi \left(\theta - e_d \frac{\pi}{2} \right) \right) \right)^2 \quad (82)
 \end{aligned}$$

$$\stackrel{(a)}{\leq} \sum_{d=1}^D \frac{1}{4} \|H\|_\infty^2 \left(\left\| \rho^\mathcal{E} \left(\theta + e_d \frac{\pi}{2} \right) - \Psi \left(\theta + e_d \frac{\pi}{2} \right) \right\|_1 \right. \\
 \left. + \left\| \rho^\mathcal{E} \left(\theta - e_d \frac{\pi}{2} \right) - \Psi \left(\theta - e_d \frac{\pi}{2} \right) \right\|_1 \right)^2 \quad (83)$$

$$\stackrel{(b)}{\leq} 4D \|H\|_\infty^2 \gamma \quad (84)$$

where the inequality in (a) follows from tracial matrix Hölder's inequality and the triangle inequality for trace norm [42]. To obtain the inequality in (b), we use the following relationship between trace distance and quantum fidelity $F(\rho, \sigma) = \left(\text{Tr}(\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}) \right)^2$ [42]:

$$\begin{aligned}
 & \left\| \rho^\mathcal{E} \left(\theta \pm e_j \frac{\pi}{2} \right) - \Psi \left(\theta \pm e_j \frac{\pi}{2} \right) \right\|_1 \\
 & \leq 2 \sqrt{1 - F \left(\rho^\mathcal{E} \left(\theta \pm e_j \frac{\pi}{2} \right), \Psi \left(\theta \pm e_j \frac{\pi}{2} \right) \right)}. \quad (85)
 \end{aligned}$$

Furthermore, the quantum fidelity in (85) can be upper bounded using (40) as

$$\begin{aligned} F(\rho^\mathcal{E}(\theta), \Psi(\theta)) &= \langle \Psi(\theta) | \rho^\mathcal{E}(\theta) | \Psi(\theta) \rangle \\ &= (1 - \gamma) \langle \Psi(\theta) | \Psi(\theta) | \Psi(\theta) \rangle \\ &\quad + \gamma \langle \Psi(\theta) | \tilde{\rho}(\theta) | \Psi(\theta) \rangle \\ &= (1 - \gamma) + \gamma \langle \Psi(\theta) | \tilde{\rho}(\theta) | \Psi(\theta) \rangle \quad (86) \\ &\geq 1 - \gamma \quad (87) \end{aligned}$$

for all $\theta \in \mathbb{R}^D$. Here, the last inequality follows since the density operator is positive semidefinite, making the second term in (86) nonnegative. Using (87) in (85) yields the inequality in (84).

APPENDIX D ADDITIONAL PROPERTIES

In this section, we provide the following additional properties of the variance term $v(p^\mathcal{E}(y|\theta))$ defined in (44). All these properties can be easily verified using the definition (44), and hence, we omit the detailed derivations here.

- 1) The inequality $v(p^\mathcal{E}(y|\theta)) \geq (1 - \gamma)v(p(y|\theta)) + \gamma v(\tilde{p}(y|\theta))$ holds for all $y \in \{1, \dots, N_h\}$ and $\theta \in \Theta$.
- 2) $v(p^\mathcal{E}(y|\theta))$ is a concave function of γ . It is increasing in the range $\gamma \in [0, \gamma^*(y, \theta))$, and decreasing in the range $\gamma \in [\gamma^*(y, \theta), 1)$, where

$$\gamma^*(y, \theta) = \min \left\{ 1, 0.5 \left(1 - \frac{v(p(y|\theta)) - v(\tilde{p}(y|\theta))}{(\tilde{p}(y|\theta) - p(y|\theta))^2} \right) \right\}. \quad (88)$$

APPENDIX E PROOF OF THEOREM 5.1

In this section, we first compute the variance of the QEM-based stochastic gradient estimator (58). With the notation $\theta_{d,\pm}^t = \theta^t \pm e_d \frac{\pi}{2}$, we have the equalities

$$\text{var} \left(\xi_t^{\text{QEM}} \right) = \mathbb{E}[\|\hat{g}_t^{\text{QEM}} - g_t\|^2] \quad (89)$$

$$= \mathbb{E}[\|\hat{g}_t^{\text{QEM}} - g_t^{\text{circ}} + g_t^{\text{circ}} - g_t\|^2] \quad (90)$$

$$\stackrel{(a)}{=} \underbrace{\mathbb{E}[\|\hat{g}_t^{\text{QEM}} - g_t^{\text{circ}}\|^2]}_A + \underbrace{\mathbb{E}[\|g_t^{\text{circ}} - g_t\|^2]}_B \quad (91)$$

where the expectation is over measurements as well as over sampled noisy circuits, and we used the definition

$$\begin{aligned} [g_t^{\text{circ}}]_d &= \frac{Z}{2N_c} \sum_{l=1}^{N_c} \text{sgn}(q_{s_l^D}) \\ &\quad \times \left(\langle H \rangle_{\rho_{s_l^D}(\theta_{d,+}^t)} - \langle H \rangle_{\rho_{s_l^D}(\theta_{d,-}^t)} \right). \quad (92) \end{aligned}$$

The equality in (a) follows from the equality $\mathbb{E}[(\hat{g}_t^{\text{QEM}} - g_t^{\text{circ}})^T (g_t^{\text{circ}} - g_t) | s_{1:N_c}^D] = 0$ with the expectation taken over the quantum measurements conditioned on the sampled noisy circuit indices $s_{1:N_c}^D$. We now analyze each of the terms

in (91) separately. The first term A in (91) evaluates as shown in (93) at the bottom of this page

$$\begin{aligned} A &= \sum_{d=1}^D \frac{Z^2}{4N_c^2} \\ &\quad \times \mathbb{E} \left[\left(\sum_{l=1}^{N_c} \text{sgn}(q_{s_l^D}) \left(\underbrace{\langle \widehat{H} \rangle_{\rho_{s_l^D}(\theta_{d,+}^t)} - \langle H \rangle_{\rho_{s_l^D}(\theta_{d,+}^t)}}_{A_{1,l}} \right. \right. \right. \\ &\quad \left. \left. \left. + \underbrace{\langle H \rangle_{\rho_{s_l^D}(\theta_{d,-}^t)} - \langle \widehat{H} \rangle_{\rho_{s_l^D}(\theta_{d,-}^t)}}_{A_{2,l}} \right) \right)^2 \right] \\ &\stackrel{(a)}{=} \frac{Z^2}{4N_c^2} \sum_{d=1}^D \text{var} \left(\sum_{l=1}^{N_c} \text{sgn}(q_{s_l^D}) (A_{1,l} + A_{2,l}) \right) \\ &\stackrel{(b)}{=} \frac{Z^2}{4N_c^2} \sum_{d=1}^D \sum_{l=1}^{N_c} \mathbb{E}[A_{1,l}^2 + A_{2,l}^2] \quad (93) \end{aligned}$$

where the equality in (a) follows by noting that for the random variable $C = \sum_{l=1}^{N_c} C_l$, with $C_l = \text{sgn}(q_{s_l^D})(A_{1,l} + A_{2,l})$, we have the equality $\mathbb{E}[C | s_{1:N_c}^D] = 0$, where the expectation is over quantum measurements conditioned on the indices $s_{1:N_c}^D$ of the sampled circuits. The equality in (b) follows since random variables C_l for $l = 1, \dots, N_c$, are independent, enabling the equality $\text{var}(C) = \sum_{l=1}^{N_c} \text{var}(C_l)$. Furthermore, we have $\text{var}(C_l) = \mathbb{E}[(A_{1,l} + A_{2,l})^2]$, which is in turn equal to $\mathbb{E}[A_{1,l}^2 + A_{2,l}^2]$, since, conditioned on a sampled noisy circuit, the expectation $\mathbb{E}[A_{1,l}A_{2,l} | s_l^D] = 0$ holds.

Now, following an analysis similar to (71), defining the Bernoulli random variable $W_{y,s_l^D} = \{Y_{s_l^D}^{\text{QEM}} = y\}$, with $Y_{s_l^D}^{\text{QEM}}$ being the random variable that denotes the index of the measurement of the observable H , we obtain the inequality

$$\begin{aligned} \mathbb{E}[A_{1,l}^2 | s_l^D] &= \frac{1}{N_m/N_c} \mathbb{E} \left[\left(\sum_{y=1}^{N_h} h_y (W_{y,s_l^D} - p_{s_l^D}(y|\theta_{d,+}^t)) \right)^2 \middle| s_l^D \right] \quad (94) \end{aligned}$$

$$\leq \frac{\text{Tr}(H^2)}{N_m/N_c} \sum_{y=1}^{N_h} v(p_{s_l^D}(y|\theta_{d,+}^t)). \quad (95)$$

A similar bound can be obtained for the term $\mathbb{E}[A_{2,l}^2 | s_{1:N_c}^D]$.

Using (95), we first obtain an upper bound on the term A in (93) as

$$A \leq \frac{Z^2}{4N_c^2} \frac{\text{Tr}(H^2)}{N_m/N_c} \sum_{d=1}^D \sum_{l=1}^{N_c} \left(\sum_{y=1}^{N_h} \mathbb{E}_{s_l^D} [v(p_{s_l^D}(y|\theta_{d,+}^t))] \right)$$

$$+ \sum_{y=1}^{N_h} \mathbb{E}_{s_l^D} \left[v \left(p_{s_l^D}(y|\theta_{d,-}^t) \right) \right] \quad (96)$$

$$\leq \frac{N_h D Z^2 \text{Tr}(H^2)}{2N_m} \sup_{\theta \in \mathbb{R}^D, y \in \{1, \dots, N_h\}} \mathbb{E}_{s^D} [v(p_{s^D}(y|\theta))] \quad (97)$$

where we used the notation \mathbb{E}_{\bullet} to represent the expectation over the random variable in the subscript \bullet .

To get a lower bound on the term A , we start by observing that

$$\mathbb{E}[A_{1,l}^2] = \mathbb{E}_{s_l^D} \left[\mathbb{E} \left[A_{1,l}^2 | s_l^D \right] \right] \quad (98)$$

$$\stackrel{(a)}{\geq} p_{s^D=\{0\}^D} \mathbb{E} \left[A_{1,l}^2 | s^D = \{0\}^D \right] \quad (99)$$

$$= p_{s^D=\{0\}^D} N_c \text{var} \left(\widehat{\langle H \rangle}_{\rho^\mathcal{E}(\theta_{d,+}^t)} \right) \quad (100)$$

with the inequality in (a) obtained by lower bounding the outer expectation in (98) with the choice of the $s^D = \{0\}^D$ circuit. This determines the noisy operation $\mathcal{O}_{s^D}^\theta(\cdot)$ with Pauli operator \mathcal{P}_0 acting on all gates times the probability $p_{s^D=\{0\}^D}$. With this choice of sampled circuit, we then have that $p_{s_l^D}(y|\theta_{d,\pm}^t) = p^\mathcal{E}(y|\theta_{d,\pm}^t)$ and $W_{y,s_l^D} = \mathbb{I}\{Y^\mathcal{E} = y\}$. This results in the last equality with $\text{var} \left(\widehat{\langle H \rangle}_{\rho^\mathcal{E}(\theta_{d,+}^t)} \right)$ defined as in (80). Similar bound holds for $\mathbb{E}[A_{2,l}^2]$. Subsequently, one obtains the lower bound

$$A \geq \frac{Z^2}{4N_c} p_{s^D=\{0\}^D} \sum_{d=1}^D \left[N_c \text{var} \left(\widehat{\langle H \rangle}_{\rho^\mathcal{E}(\theta_{d,+}^t)} \right) + N_c \text{var} \left(\widehat{\langle H \rangle}_{\rho^\mathcal{E}(\theta_{d,-}^t)} \right) \right] \quad (101)$$

$$= Z^2 p_{s^D=\{0\}^D} \text{var}(\xi_t^\mathcal{E}) = c_1(\gamma) \text{var}(\xi_t^\mathcal{E}) \quad (102)$$

with $\text{var}(\xi_t^\mathcal{E})$ defined as in (79).

We now evaluate the second term B in (91). Toward this, we note the inequality

$$\begin{aligned} B &= \sum_{d=1}^D \mathbb{E} \left[\left(\left[g_t^{\text{circ}} \right]_d - [g_t]_d \right)^2 \right] \\ &= \sum_{d=1}^D \sum_{l=1}^{N_c} \frac{Z^2}{4N_c^2} \\ &\quad \times \text{var} \left(\text{sgn}(q_{s_l^D}) \left[\langle H \rangle_{\rho_{s_l^D}(\theta_{d,+}^t)} - \langle H \rangle_{\rho_{s_l^D}(\theta_{d,-}^t)} \right] \right) \\ &\stackrel{(a)}{\leq} \sum_{d=1}^D \sum_{l=1}^{N_c} \frac{Z^2}{4N_c^2} \end{aligned}$$

$$\times \mathbb{E} \left[\left(\text{sgn}(q_{s_l^D}) \left[\langle H \rangle_{\rho_{s_l^D}(\theta_{d,+}^t)} - \langle H \rangle_{\rho_{s_l^D}(\theta_{d,-}^t)} \right] \right)^2 \right] \quad (103)$$

$$\begin{aligned} &= \frac{Z^2}{4N_c} \sum_{d=1}^D \mathbb{E}_{s^D} \left[\left(\langle H \rangle_{\rho_{s^D}(\theta_{d,+}^t)} - \langle H \rangle_{\rho_{s^D}(\theta_{d,-}^t)} \right)^2 \right] \\ &= \frac{Z^2}{4N_c} \sum_{d=1}^D \mathbb{E}_{s^D} \left[\left(\text{Tr} \left(H \left(\rho_{s^D}(\theta_{d,+}^t) - \rho_{s^D}(\theta_{d,-}^t) \right) \right) \right)^2 \right] \\ &\stackrel{(b)}{\leq} \frac{Z^2 \|H\|_\infty^2}{4N_c} \sum_{d=1}^D \mathbb{E}_{s^D} \left[\|\rho_{s^D}(\theta_{d,+}^t) - \rho_{s^D}(\theta_{d,-}^t)\|_1^2 \right] \\ &\stackrel{(c)}{\leq} \frac{D Z^2 \|H\|_\infty^2}{N_c} \quad (104) \end{aligned}$$

where in (a) we used $\text{var}(X) \leq \mathbb{E}[X^2]$. The inequality in (b) follows by the application of the tracial matrix Hölder's inequality [41], and the inequality (c) follows since trace norm between density matrices is upper bounded by 2 [42].

Noting the inequality $B \geq 0$ and using (102), we get the following lower bound:

$$\text{var}(\xi_t^{\text{QEM}}) \geq c_1(\gamma) \text{var}(\xi_t^\mathcal{E}).$$

Furthermore, using (97) and (104) in (91) yields the upper bound V^{QEM} of (63).

To show the relation in (64), we start with the upper bound on term A obtained in (97). One can lower bound this term as

$$\begin{aligned} &\frac{N_h D Z^2 \text{Tr}(H^2)}{2N_m} \sup_{\theta \in \mathbb{R}^D, y \in \{1, \dots, N_h\}} \mathbb{E}_{s^D} [v(p_{s^D}(y|\theta))] \\ &\geq \frac{N_h D Z^2 \text{Tr}(H^2) p_{s^D=\{0\}^D}}{2N_m} \sup_{\theta \in \mathbb{R}^D, y \in \{1, \dots, N_h\}} v(p^\mathcal{E}(y|\theta)) \\ &= Z^2 p_{s^D=\{0\}^D} V^\mathcal{E} = c_1(\gamma) V^\mathcal{E} \quad (105) \end{aligned}$$

where $V^\mathcal{E}$ is as defined in (43). Together with the bound in (104), we then get the inequality

$$V^{\text{QEM}} \geq c_1(\gamma) V^\mathcal{E} + c_2(\gamma) \frac{D \|H\|_\infty^2}{N_c} \quad (106)$$

with $c_2(\gamma) = Z^2$.

Finally, we analyze the scalar functions $c_1(\gamma)$ and $c_2(\gamma)$ when the quantum noise channel in (17) is depolarizing. In this case, [44, Th. 2] gives the equalities

$$Z_d = \frac{1 + (1 - 2^{1-2n})(1 - (1 - \gamma)^{1/D})}{(1 - \gamma)^{1/D}} \quad (107)$$

and

$$p_d(0) = \frac{2^{2n} - 1 + (1 - \gamma)^{1/D}}{2^{2n} (1 + (1 - 2^{1-2n})(1 - (1 - \gamma)^{1/D}))} \quad (108)$$

for $d = 1, \dots, D$. Note that the inequalities $Z_d \geq 1$ and $\frac{dZ_d}{d\gamma} \geq 0$ holds, whereby Z_d is a nondecreasing function of γ .

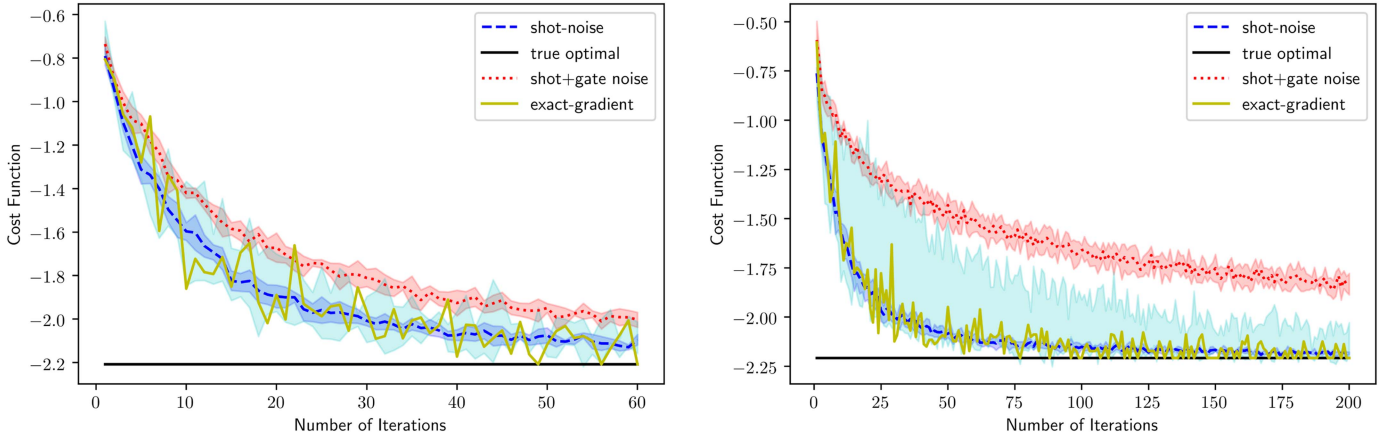


FIGURE 8. Convergence analysis of the SGD as a function of the number of iterations when exact gradients can be computed, when only shot-noise is present, when shot and gate noise are present, and when QEM is employed ($n = 3$ qubit system subjected to depolarizing noise on the CNOT gates). The learning rate is set as $\eta_t = 0.5/t$; $N_c = 8$; and $N_m = 400$. (Left) Low noise with $\epsilon = 0.09$; (Right) Large noise with $\epsilon = 0.25$. The loss function $L(\theta^t)$ is evaluated during testing by taking N_m measurements of a gate-noise free PQC.

Subsequently, we have the equality $c_2(\gamma) = Z^2 = (Z_d)^{2D} \geq 1$ and $c_2(\gamma)$ is a nondecreasing function of γ .

We now analyze the term $c_1(\gamma) = Z^2 p_{s^D=\{0\}^D} = (Z_d^2 p_{d,s=0})^D$, where we have $Z_d^2 p_{d,s=0} = \frac{\text{num}}{\text{den}}$ with

$$\text{num} = (2^{2n+1} - 4 + 2^{1-2n}) + (1 - \gamma)^{1/D} (5 - 4/2^{2n} - 2^{2n}) - (1 - \gamma)^{2/D} (1 - 2^{1-2n})$$

and $\text{den} = 2^{2n} (1 - \gamma)^{2/D}$. The following two properties can be easily verified. First, we show that

$$\begin{aligned} \text{num} - \text{den} &= (2^{2n+1} - 4 + 2^{1-2n}) + (1 - \gamma)^{1/D} (5 - 4/2^{2n} - 2^{2n}) \\ &\quad - (1 - \gamma)^{2/D} (1 - 2^{1-2n} + 2^{2n}) \\ &\stackrel{(a)}{\geq} (2^{2n+1} - 4 + 2^{1-2n}) + (1 - \gamma)^{1/D} (5 - 4/2^{2n} - 2^{2n} \\ &\quad - 1 + 2^{1-2n} - 2^{2n}) \\ &= (2^{2n+1} - 4 + 2^{1-2n}) - (1 - \gamma)^{1/D} (-4 + 2/2^{2n} + 2^{2n+1}) \\ &\geq 0 \end{aligned}$$

where inequality (a) follows since $(1 - \gamma)^2 \leq (1 - \gamma)$ and the last inequality follows since $(1 - \gamma)^{1/D} \leq 1$. As a consequence, we have the equality $Z_d^2 p_{d,s=0} = \text{num}/\text{den} \geq 1$, and thus, the inequality $c_1(\gamma) \geq 1$.

Lastly, it can be verified that

$$\begin{aligned} \frac{d(Z_d^2 p_{d,s=0})}{d\gamma} &= \frac{d}{d\gamma} \frac{\text{num}}{\text{den}} = \frac{1}{D(1 - \gamma)^{2/D+1}} \\ &\quad \times \left(4 - \frac{8}{2^{2n}} + \frac{4}{2^{4n}} \right) \\ &\quad + \frac{1}{D(1 - \gamma)^{1/D+1}} \left(\frac{5}{2^{2n}} - \frac{4}{2^{4n}} - 1 \right) \\ &\geq \frac{1}{D(1 - \gamma)^{1/D+1}} \left(3 - \frac{3}{2^{2n}} \right) \geq 0 \quad (109) \end{aligned}$$

holds, where the last inequality follows by the inequality $(1 - \gamma)^2 \leq (1 - \gamma)$. Consequently, $Z_d^2 p_{d,s=0}$, and thus $c_1(\gamma)$, are nondecreasing functions of γ .

APPENDIX F ADDITIONAL EXPERIMENTS

In Fig. 5 of Section VI, we have evaluated the performance of the parameter iterate θ^t of the SGD in terms of the loss function $L(\theta)$ in (6), which can be estimated in practice given an infinite measurements on a gate-noise free PQC during testing. In this section, we assume that during testing, a gate-noise free PQC with the same measurement budget as that during training is available. As can be seen from Fig. 8, obtained under the same conditions as Fig. 5, the finite number of measurements of the PQC made during testing induces a larger variance on the evaluated loss function for all settings.

REFERENCES

- [1] D. Aharonov and M. Ben-Or, "Fault-tolerant quantum computation with constant error rate," *SIAM J. Comput.*, vol. 38, no. 4, pp. 1207–1282, 2008, doi: [10.1137/S0097539799359385](https://doi.org/10.1137/S0097539799359385).
- [2] K. Temme, S. Bravyi, and J. M. Gambetta, "Error mitigation for short-depth quantum circuits," *Phys. Rev. Lett.*, vol. 119, no. 18, 2017, Art. no. 180509, doi: [10.1103/PhysRevLett.119.180509](https://doi.org/10.1103/PhysRevLett.119.180509).
- [3] S. Endo, S. C. Benjamin, and Y. Li, "Practical quantum error mitigation for near-future applications," *Phys. Rev. X*, vol. 8, no. 3, 2018, Art. no. 031027, doi: [10.1103/PhysRevX.8.031027](https://doi.org/10.1103/PhysRevX.8.031027).
- [4] M. Schuld, I. Sinayskiy, and F. Petruccione, "An introduction to quantum machine learning," *Contemporary Phys.*, vol. 56, no. 2, pp. 172–185, 2015, doi: [10.1080/00107514.2014.964942](https://doi.org/10.1080/00107514.2014.964942).
- [5] O. Simeone, "An introduction to quantum machine learning for engineers," vol. 16, no. 1/2, 2022, pp. 1–223, doi: [10.1561/2000000118](https://doi.org/10.1561/2000000118).
- [6] A. Peruzzo et al., "A variational eigenvalue solver on a photonic quantum processor," *Nature Commun.*, vol. 5, no. 1, 2014, Art. no. 4213, doi: [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213).
- [7] A. Garcia-Saez and J. Latorre, "Addressing hard classical problems with adiabatically assisted variational quantum eigensolvers," 2018, *arXiv:1806.02287*, doi: [10.48550/arXiv.1806.02287](https://doi.org/10.48550/arXiv.1806.02287).
- [8] D. Amaro, C. Modica, M. Rosenkranz, M. Fiorentini, M. Benedetti, and M. Lubasch, "Filtering variational quantum algorithms for combinatorial optimization," *Quantum Sci. Technol.*, vol. 7, no. 1, 2022, Art. no. 015021, doi: [10.1088/2058-9565/ac3e54](https://doi.org/10.1088/2058-9565/ac3e54).

- [9] N. Moll et al., “Quantum optimization using variational algorithms on near-term quantum devices,” *Quantum Sci. Technol.*, vol. 3, no. 3, 2018, Art. no. 030503, doi: [10.1088/2058-9565/aab822](https://doi.org/10.1088/2058-9565/aab822).
- [10] L. Gentini, A. Cuccoli, S. Pirandola, P. Verrucchi, and L. Banchi, “Noise-resilient variational hybrid quantum-classical optimization,” *Phys. Rev. A*, vol. 102, no. 5, 2020, Art. no. 052414, doi: [10.1103/PhysRevA.102.052414](https://doi.org/10.1103/PhysRevA.102.052414).
- [11] P. W. Shor, “Scheme for reducing decoherence in quantum computer memory,” *Phys. Rev. A*, vol. 52, no. 4, 1995, Art. no. R2493, doi: [10.1103/PhysRevA.52.R2493](https://doi.org/10.1103/PhysRevA.52.R2493).
- [12] A. M. Steane, “Error correcting codes in quantum theory,” *Phys. Rev. Lett.*, vol. 77, no. 5, 1996, Art. no. 793, doi: [10.1103/PhysRevLett.77.793](https://doi.org/10.1103/PhysRevLett.77.793).
- [13] Y. Li and S. C. Benjamin, “Efficient variational quantum simulator incorporating active error minimization,” *Phys. Rev. X*, vol. 7, no. 2, 2017, Art. no. 021050, doi: [10.1103/PhysRevX.7.021050](https://doi.org/10.1103/PhysRevX.7.021050).
- [14] J. J. Wallman and J. Emerson, “Noise tailoring for scalable quantum computation via randomized compiling,” *Phys. Rev. A*, vol. 94, no. 5, 2016, Art. no. 052325, doi: [10.1103/PhysRevA.94.052325](https://doi.org/10.1103/PhysRevA.94.052325).
- [15] E. Knill, “Quantum computing with realistically noisy devices,” *Nature*, vol. 434, no. 7029, pp. 39–44, 2005, doi: [10.1038/nature03350](https://doi.org/10.1038/nature03350).
- [16] S. Endo, Z. Cai, S. C. Benjamin, and X. Yuan, “Hybrid quantum-classical algorithms and quantum error mitigation,” *J. Phys. Soc. Jpn.*, vol. 90, no. 3, 2021, Art. no. 032001, doi: [10.7566/JPSJ.90.032001](https://doi.org/10.7566/JPSJ.90.032001).
- [17] A. Ajallooeian and S. U. Stich, “On the convergence of SGD with biased gradients,” 2020, *arXiv:2008.00051*, doi: [10.48550/arXiv.2008.00051](https://doi.org/10.48550/arXiv.2008.00051).
- [18] P. Braccia, L. Banchi, and F. Caruso, “Quantum noise sensing by generating fake noise,” *Phys. Rev. Appl.*, vol. 17, no. 2, 2022, Art. no. 024002, doi: [10.1103/PhysRevApplied.17.024002](https://doi.org/10.1103/PhysRevApplied.17.024002).
- [19] A. Kandala et al., “Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets,” *Nature*, vol. 549, no. 7671, pp. 242–246, 2017, doi: [10.1038/nature23879](https://doi.org/10.1038/nature23879).
- [20] J. M. Kübler, A. Arrasmith, L. Cincio, and P. J. Coles, “An adaptive optimizer for measurement-frugal variational algorithms,” *Quantum*, vol. 4, 2020, Art. no. 263, doi: [10.22331/q-2020-05-11-263](https://doi.org/10.22331/q-2020-05-11-263).
- [21] S. Wang et al., “Noise-induced barren plateaus in variational quantum algorithms,” *Nature Commun.*, vol. 12, no. 1, 2021, Art. no. 6961, doi: [10.1038/s41467-021-27045-6](https://doi.org/10.1038/s41467-021-27045-6).
- [22] Y. Du, M.-H. Hsieh, T. Liu, S. You, and D. Tao, “Learnability of quantum neural networks,” *PRX Quantum*, vol. 2, no. 4, 2021, Art. no. 040337, doi: [10.1103/PRXQuantum.2.040337](https://doi.org/10.1103/PRXQuantum.2.040337).
- [23] A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, “Error mitigation extends the computational reach of a noisy quantum processor,” *Nature*, vol. 567, no. 7749, pp. 491–495, 2019, doi: [10.1038/s41586-019-1040-7](https://doi.org/10.1038/s41586-019-1040-7).
- [24] B. Zhang et al., “Hidden inverses: Coherent error cancellation at the circuit level,” *Phys. Rev. Appl.*, vol. 17, no. 3, 2022, Art. no. 034074, doi: [10.1103/PhysRevApplied.17.034074](https://doi.org/10.1103/PhysRevApplied.17.034074).
- [25] V. Leyton-Ortega, S. Majumder, and R. C. Pooser, “Quantum error mitigation by hidden inverses protocol in superconducting quantum devices,” *J. Quantum Sci. Technol.*, vol. 8, no. 1, Dec. 2022, Art. no. 014008, doi: [10.1088/2058-9565/aca92d](https://doi.org/10.1088/2058-9565/aca92d).
- [26] Z. Cai, “Multi-exponential error extrapolation and combining error mitigation techniques for NISQ applications,” *NPJ Quantum Inf.*, vol. 7, no. 1, 2021, Art. no. 80, doi: [10.1038/s41534-021-00404-3](https://doi.org/10.1038/s41534-021-00404-3).
- [27] A. Mari, N. Shammah, and W. J. Zeng, “Extending quantum probabilistic error cancellation by noise scaling,” *Phys. Rev. A*, vol. 104, no. 5, 2021, Art. no. 052607, doi: [10.1103/PhysRevA.104.052607](https://doi.org/10.1103/PhysRevA.104.052607).
- [28] A. Strikis, D. Qin, Y. Chen, S. C. Benjamin, and Y. Li, “Learning-based quantum error mitigation,” *PRX Quantum*, vol. 2, no. 4, 2021, Art. no. 040330, doi: [10.1103/PRXQuantum.2.040330](https://doi.org/10.1103/PRXQuantum.2.040330).
- [29] S. Wang, P. Czarnik, A. Arrasmith, M. Cerezo, L. Cincio, and P. J. Coles, “Can error mitigation improve trainability of noisy variational quantum algorithms?,” 2021, *arXiv:2109.01051*, doi: [10.48550/arXiv.2109.01051](https://doi.org/10.48550/arXiv.2109.01051).
- [30] R. Sweke et al., “Stochastic gradient descent for hybrid quantum-classical optimization,” *Quantum*, vol. 4, 2020, Art. no. 314, doi: [10.22331/q-2020-08-31-314](https://doi.org/10.22331/q-2020-08-31-314).
- [31] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, “Parameterized quantum circuits as machine learning models,” *Quantum Sci. Technol.*, vol. 4, no. 4, 2019, Art. no. 043001, doi: [10.1088/2058-9565/ab4eb5](https://doi.org/10.1088/2058-9565/ab4eb5).
- [32] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information*, 10th ed. Cambridge, U.K.: Cambridge Univ. Press, 2020, doi: [10.1017/CBO9780511976667](https://doi.org/10.1017/CBO9780511976667).
- [33] S. Mangini, L. Maccone, and C. Macchiavello, “Qubit noise deconvolution,” 2021, *arXiv:2112.03043*, doi: [10.1140/epjqt/s40507-022-00151-0](https://doi.org/10.1140/epjqt/s40507-022-00151-0).
- [34] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, “Evaluating analytic gradients on quantum hardware,” *Phys. Rev. A*, vol. 99, no. 3, 2019, Art. no. 032331, doi: [10.1103/PhysRevA.99.032331](https://doi.org/10.1103/PhysRevA.99.032331).
- [35] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, “Barren plateaus in quantum neural network training landscapes,” *Nature Commun.*, vol. 9, no. 1, 2018, Art. no. 4812, doi: [10.1038/s41467-018-07090-4](https://doi.org/10.1038/s41467-018-07090-4).
- [36] B. Koczor, “The dominant eigenvector of a noisy quantum state,” *New J. Phys.*, vol. 23, no. 12, 2021, Art. no. 123047, doi: [10.1088/1367-2630/ac37ae](https://doi.org/10.1088/1367-2630/ac37ae).
- [37] T. Kolda, “Multilinear operators for higher-order decompositions technical report,” Sandia Nat. Lab., Albuquerque, NM, USA and Livermore, CA, USA, Sandia Rep. sand2006-2081, 2006, doi: [10.2172/923081](https://doi.org/10.2172/923081).
- [38] Y. Xiong, S. X. Ng, and L. Hanzo, “The accuracy vs. sampling overhead trade-off in quantum error mitigation using Monte Carlo-based channel inversion,” *IEEE Trans. Commun.*, vol. 70, no. 3, pp. 1943–1956, Mar. 2022, doi: [10.1109/TCOMM.2022.3144469](https://doi.org/10.1109/TCOMM.2022.3144469).
- [39] G. Aleksandrowicz et al., “Qiskit: An open-source framework for quantum computing,” Jan. 2019, doi: [10.5281/zenodo.2562111](https://doi.org/10.5281/zenodo.2562111).
- [40] Y. Du, T. Huang, S. You, M.-H. Hsieh, and D. Tao, “Quantum circuit architecture search: Error mitigation and trainability enhancement for variational quantum solvers,” 2020, *arXiv:2010.10217*, doi: [10.1038/s41534-022-00570-y](https://doi.org/10.1038/s41534-022-00570-y).
- [41] B. Baumgartner, “An inequality for the trace of matrix products, using absolute values,” 2011, *arXiv:1106.6189*, doi: [10.48550/arXiv.1106.6189](https://doi.org/10.48550/arXiv.1106.6189).
- [42] M. M. Wilde, *Quantum Information Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2013, doi: [10.1017/CBO978113952534](https://doi.org/10.1017/CBO978113952534).
- [43] L. Banchi, J. Pereira, S. Lloyd, and S. Pirandola, “Convex optimization of programmable quantum computers,” *NPJ Quantum Inf.*, vol. 6, no. 1, 2020, Art. no. 42, doi: [10.1038/s41534-020-0268-2](https://doi.org/10.1038/s41534-020-0268-2).
- [44] R. Takagi, “Optimal resource cost for error mitigation,” *Phys. Rev. Res.*, vol. 3, no. 3, 2021, Art. no. 033178, doi: [10.1103/PhysRevResearch.3.033178](https://doi.org/10.1103/PhysRevResearch.3.033178).