

Iteration Complexity of Variational Quantum Algorithms

Vyacheslav Kungurtsev¹, Georgios Korpas^{1,2}, Jakub Marecek¹, and Elton Yechao Zhu³

¹Department of Computer Science, Czech Technical University in Prague, Karlovo nám. 13, Prague 2, Czech Republic

²Archimedes Research Unit on AI, Data Science and Algorithms, Athena Research and Innovation Center, 15125 Marousi, Greece

³Fidelity Center for Applied Technology, FMR LLC, Boston, MA 02210, USA

There has been much recent interest in near-term applications of quantum computers, i.e., using quantum circuits that have short decoherence times due to hardware limitations. Variational quantum algorithms (VQA), wherein an optimization algorithm implemented on a classical computer evaluates a parametrized quantum circuit as an objective function, are a leading framework in this space. An enormous breadth of algorithms in this framework have been proposed for solving a range of problems in machine learning, forecasting, applied physics, and combinatorial optimization, among others.

In this paper, we analyze the iteration complexity of VQA, that is, the number of steps that VQA requires until its iterates satisfy a surrogate measure of optimality. We argue that although VQA procedures incorporate algorithms that can, in the idealized case, be modeled as classic procedures in the optimization literature, the particular nature of noise in near-term devices invalidates the claim of applicability of off-the-shelf analyses of these algorithms. Specifically, noise makes the evaluations of the objective function via quantum circuits *biased*. Commonly used optimization procedures, such as SPSA and the parameter shift rule, can thus be seen as derivative-free optimization algorithms with biased function evaluations, for which there are currently no iteration complexity guarantees in the literature. We derive the missing guarantees and find that the rate of convergence is unaffected. However, the level of bias contributes unfavorably to both the constant therein, and the asymptotic distance to stationarity, i.e., the more bias, the farther one is guaranteed, at best, to reach a stationary point of the VQA objective.

1 Introduction

There has been a great deal of recent interest in near-term applications of quantum computers. Variational quantum algorithms [1, 2, 3, 4] are the most popular approach by a large margin, as evidenced by the quantity of papers and citations thereof studying their performance on a wide range of practical and theoretical problems. Recent studies show that they can potentially outperform existing approaches in a number of important tasks, such as combinatorial optimization [5], quantum chemistry [6, 7], particle physics [8, 9] and many others.

In variational quantum algorithms (VQAs), a parametric quantum circuit is run and “evaluated”, or endowed with measurement that collapses the state to some noisy eigenvalue of an observable operator. This measurement is in turn used to classically guide the estimate of the parameters, which are adjusted according to the evaluation, and then used to prepare the (updated) parametric quantum circuit again. Thus, it is a scheme to simultaneously use classical and quantum computing in an alternating manner, developed as a means of utilizing the potential speedup properties of quantum computing despite the inability to obtain useful output from circuits beyond a relatively shallow depth in the so-called near-term devices [3].

The classical computation component can be viewed as an unconstrained optimization problem, wherein a function of the so-called variational parameters is minimized with respect to these variables. This objective function has certain specific properties, which we shall discuss in detail in the following. However, one important and immediate point is that the evaluations of the function (and its gradients and any higher-order information) are ultimately noisy, meaning that one evaluation with a certain set of parameters may yield a different value than another evaluation with the same set of variational parameters.

We can distinguish two classes of algorithms used to compute the steps of the classical optimization procedure in VQA. One popular option [10], inspired by automatic differentiation, evaluates closed-form expressions for the gradient [11, 12]. Alternatively, one can use “zero order” or “derivative free optimization” methods, which only make use of function evaluations rather than seeking to compute gradients. A popular choice is the *Simultaneous Perturbation Stochastic Approximation* (SPSA) algorithm, a classic method introduced by Spall et al. [13] to solve stochastic root-finding problems using noisy function evaluations to approximate a gradient. There it was found that a careful evaluation of the problem function at distinct points yielded a more reliable estimate of the gradient than finite difference methods, relative to the number of evaluations required, and it became a standard tool of zeroth-order stochastic approximation. Despite being a comparatively unused method in classical settings, known to perform substandard compared to other methods, in the case of VQA, practical numerical experience has yielded SPSA to be the most popular zero order optimization method, see, e.g. [14]. In studying the convergence properties under biased noise, we hope to elucidate on the potential unique value of SPSA to this class of problems.

More commonly in the general optimization literature, an approach described (ambiguously, since SPSA also used two evaluations) as the “two-point function approximation” rule is the most extensively analyzed. As such, as a point of both theoretical and numerical comparison, we consider the application of both SPSA vis-a-vis the two point function approximation rule for optimizing VQA objectives.

The third popular procedure, which is unique to the functional form of parametrized quantum circuits, is called the “parameter shift rule”. It can be understood as an intermediate step between zero’th- and first-order methods. In particular, although by strict

definition the procedure is zero-order in that it estimates a gradient using function evaluations, it is based on the original observation that for certain classes of gates and choices of perturbation, the approximation is exact. Based on select choices of the direction and magnitude of the perturbation around the base point, the parameter shift rule can be calculated to produce the exact gradient for certain circuits [15, 16]. The class of circuits for this applicability, however, is limited, especially when one considers the entire pipeline through an ultimate evaluation of an observable. Recent work has expanded the scope of procedures that appear equivalent to analytical expressions as well as the estimation accuracy of those that are not [17]. As a consequence, given that the parameter shift rule computes analytically exact in certain contexts, and approximations in many others, it is best considered as a method in between zero- and first-order, and is best analyzed in a manner attuned to the problem at hand.

A recent numerical comparison of these solvers is given [14]. Extensive testing and hyperparameter tuning was used to obtain the best settings for computational performance on a set of representative problems. SPSA is compared with a custom solver and a few classic second-order optimization methods. There was no clear winner, with different methods being competitive for different problem types and parameters.

In regards to claiming analytical results on these methods, while there are plenty of studies on the iteration complexity of classical optimization algorithms, their applicability to VQA is suspect due to the presence and nature of the errors arising in near-term circuits. That is: beyond the potential stochastic nature of a quantum circuit evaluation (i.e., for a number of problems, the theoretical output can be one of several eigenstates of an observable operator), the gates in near-term devices are far from implemented with perfect fidelity. The errors or unintended operations in the quantum state [18] are difficult to model. Common models consider rotation (*dephasing*) and dampening of an amplitude (*dissipation*). While there is active research in developing hardware and error-correcting codes to mitigate the scale and impact of these errors, for the time being, it is clear that such systematic errors in the evaluation of quantum circuits shall remain endemic for some time.

Whereas there is thorough theoretical understanding of the performance of many variants of SGD and zero-order optimization methods, the analysis is typically inspired by statistical and machine learning applications. In this case, the function being minimized is a loss function in the form of a large sum of discrepancies from actual to predicted estimates across a large sum of data samples. The stochasticity comes from the necessity of sampling one or multiple (minibatch) samples in evaluation of the function or gradient during the course of the optimization procedure, due to the frequently large sample size and unrealizability of storing an entire function or gradient in memory. A key feature of this class of problems is that a function or gradient evaluation can be performed in such a way that it is unbiased, i.e., while it is noisy, the statistical expectation of the evaluation is equal to a desired deterministic quantity.

By contrast, as we shall see below, the errors that appear in the evaluation of quantum circuits in VQA tend to exhibit "bias". To formalize this notion, consider an optimization problem of the form,

$$\min_{\boldsymbol{\theta}} f(\boldsymbol{\theta}) := \mathbb{E}_{\xi \sim \Xi(\boldsymbol{\theta})}[F(\boldsymbol{\theta}, \xi)], \quad (1)$$

where $\boldsymbol{\theta} \in \mathbb{R}^p$ are the deterministic VQA parameters and $\xi \in \mathbb{R}$ is the stochastic sample. We consider that Ξ is the underlying distribution on which we want to minimize this first moment quantity. However, in practice, we may not have access to unbiased samples of the quantity in (1), but instead each physical evaluation of $F(\boldsymbol{\theta}, \cdot)$ satisfies a difference

distribution, formally,

$$F(\boldsymbol{\theta}, \cdot) \sim F(\boldsymbol{\theta}, \xi), \xi \sim \chi(\boldsymbol{\theta}), \quad (2)$$

where in general the state-dependent distributions over ξ are not the same nor do they have the same moments, i.e.,

$$\mathbb{E}_{\xi \sim \Xi(\boldsymbol{\theta})}[F(\boldsymbol{\theta}, \xi)] \neq \mathbb{E}_{\xi \sim \chi(\boldsymbol{\theta})}[F(\boldsymbol{\theta}, \xi)].$$

Thus, in this work the desired minimization is with respect to the distribution denoted by $\Xi(\boldsymbol{\theta})$, but only estimates associated with the distribution $\chi(\boldsymbol{\theta})$ are available. We assume that there is a reasonable amount of resemblance in the sense that $\text{supp}(\cup_{\boldsymbol{\theta}} \Xi(\boldsymbol{\theta})) \subseteq \text{supp}(\cup_{\boldsymbol{\theta}} \chi(\boldsymbol{\theta}))$. In the sequel, for parsimony we will drop the state-dependence of the distributions, i.e., write $\Xi(\boldsymbol{\theta})$ as just Ξ , observing that the notation limitation does not hinder the clarity of the presented results.

In standard Stochastic Approximation (SA) Theory, one iteratively computes,

$$\mathbf{g}_k \approx \nabla F(\boldsymbol{\theta}_k, \xi_k) = \nabla f(\boldsymbol{\theta}_k) + \mathbf{n}(\boldsymbol{\theta}_k, \xi_k) \quad (3)$$

i.e., forming some possibly approximate estimate of a sample of ∇F from a distribution, written formally as the exact gradient of the (hidden) function f plus an approximation error denoted as $\mathbf{n}(\boldsymbol{\theta}_k, \xi_k)$. This error represents the deviation from the exact gradient and can be due to systemic stochastic noise and algorithmic inexact approximation of the gradient. We consider the general case of state-dependent noise, where $\boldsymbol{\theta}_k$ is the state.

SA then involves iterating,

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \alpha_k \mathbf{g}_k \quad (4)$$

where α_k is a *gain* sequence satisfying

$$\sum_{k=0}^{\infty} \alpha_k = \infty, \quad \sum_{k=0}^{\infty} (\alpha_k)^2 < \infty. \quad (5)$$

Standard SGD methods assume that such an evaluation satisfies $\mathbb{E}_{\xi}[\mathbf{n}(\boldsymbol{\theta}_k, \xi)] = 0$, that is, they assume the noise is unbiased. The convergence properties of SA under asymptotically vanishing biased noise are considered throughout many classic texts [19]. Several recent works do consider systemic bias that remains throughout the iterative process, i.e., $\mathbb{E}_k[\mathbf{n}(\boldsymbol{\theta}_k, \xi_k)] \neq 0, \forall k$, and, even stronger, a stationary $\limsup_k \|\mathbb{E}_k[\mathbf{n}(\boldsymbol{\theta}_k, \xi_k)]\| = b \neq 0, \forall k$, has been considered in works such as [20, 21] and most recently [22]. Of course, with systematic bias, one cannot hope to reach, even asymptotically, even a local minimizer. Instead, guarantees an iteration is in some neighborhood of a stationary point.

Analyses of biased zeroth-order optimization methods in the literature are scant, to the best of the authors' knowledge. In [23] this problem is claimed to be uninteresting since one can simply minimize $f(\boldsymbol{\theta}_k) - \mathbb{E}_{\xi}[\mathbf{n}(\boldsymbol{\theta}_k, \xi_k)]$. However, we believe this is misguided, at least in the case of VQA, as one does not have a closed-form expression for the noise, and its estimation may add significant computational expense. That said, we are also unaware of any other applications, other than VQA, wherein there is a known (noisy but with known distribution) function to minimize, with the actual observations subject to biased noise with unknown distribution and even moments thereof. As such, the novelty of the work is appropriate for the highlight, and to the knowledge of the authors, sole relevant application of VQA.

The presence of biased noise indicates that one cannot hope to achieve, even asymptotically, a stationary point of $\min_{\boldsymbol{\theta}} \mathbb{E}_{\xi}[F(\boldsymbol{\theta}, \xi)]$. Instead, there are notions of approximation

and robustness that become relevant: given a bound for the noise, for a zero-order algorithm, 1) how close can we guarantee that iterates get to stationary points and 2) how can we guarantee an estimate from the algorithm that is the best choice under consideration of the worst case outcome of the noise? We should also mention that there exists one work studying the asymptotic guarantees of SPSA under bounded noise [24] with subsequent application to Gaussian mixture models [25]. In this work the authors established a quantitative bound on the asymptotic bias associated with the mean squared error of the gradient for SPSA and other zero-order optimization procedures.

Our contribution In this paper, we are interested in studying the properties of classical optimization algorithms as used for variational quantum algorithms (VQA) under non-trivial noise in the evaluation of a quantum circuit. The bias in zeroth-order methods would consist of *function evaluations* yielding an error, i.e.,

$$f \sim F(\boldsymbol{\theta}, \xi) = f(\boldsymbol{\theta}) + n(\boldsymbol{\theta}, \xi), \quad \xi \in \chi \quad (6)$$

with $\mathbb{E}_{\xi \sim \chi}[n(\boldsymbol{\theta}, \xi)] \neq 0$.

The rest of this paper proceeds as follows. Section 2 presents the background on VQA and applicable optimization methods. We also present our argument in regards to the necessity of considering systematic bias in the analysis of the convergence of optimization algorithms for VQA and review the literature on convergence guarantees in the presence of bias. Section 3 presents our novel contribution in providing nonasymptotic guarantees of neighborhood convergence. Section 4 presents the results of numerical performance on a set of quantum optimization problems. Finally, we discuss the implications of our work and insights for future research in Section 5. Appendix A provides the proofs of our main results. In summary, this work represents a more faithful-to-reality convergence analysis for quantum variational solvers, providing first-principles insight into the intuitive practice of eschewing attempts to evaluate a gradient and consider a zero-order method instead.

2 Background

2.1 Variational Quantum Algorithms

The advent of *variational quantum algorithms* (VQAs) [26, 27, 28, 1, 2, 3] came as a practical response to the challenges of *noisy intermediate scale quantum* (NISQ) devices [29]. NISQ devices suffer from coherent and incoherent noise, short coherence times, and limitations on qubit connectivity. VQA attempt to achieve quantum advantage under those practical limitations. Of course, VQAs also have many disadvantages. By incorporating a hybrid classical-quantum framework, a significant overhead appears in preparing the initial quantum state to feed as input to the circuit, and significant computing potential is lost by frequent wave collapse by observation. On the other hand, one of their attractive features is that they provide a general framework that fits a range of problems, from quantum simulation [26] to solving differential equations [30] and to designing quantum autoencoders [31], to name a few. No matter what class of problem one is interested in solving using VQAs, the basic structure of these algorithms is invariant, despite the fact that certain applications might benefit from different algorithmic subroutines.

Specifically, VQAs are built around a quantum-classical loop. First, assuming that one contains data that can be encoded in a parametrized quantum circuit, one defines a cost function $C(\boldsymbol{\theta})$, $\boldsymbol{\theta} \in \mathbb{R}^p$, whose true minimum, obtained by the optimizer $\boldsymbol{\theta}^* :=$

Quantity	Description	Reference
F, C	cost functions	Eq. (1) and Eq. (7)
$\boldsymbol{\theta}$	variational parameter in $\mathbb{R}^p$	Eq. (1) and Sec. 2.1
$\boldsymbol{\theta}^*$	optimal variational parameter	Sec. 2.1
ξ	stochastic sample	Eq. (1)
$\Xi(\boldsymbol{\theta})$	ideal (noise) distribution	Eq. (1)
$\chi(\boldsymbol{\theta})$	physical (noise) distribution	Eq. (2)
$n(\boldsymbol{\theta}, \xi)$	random variable describing mean zero and noisy bias	Eq. (6)
$\mathbf{g}_k$	(stochastic) gradient at step k	Eq. (3)
α_k	gain in the negative gradient direction k	Eq. (4) and Eq. (5)
$ \psi_0\rangle, \rho_0$	initial quantum state	Eq. (7) and Sec. 2.1
$ \psi(\boldsymbol{\theta})\rangle, \rho(\boldsymbol{\theta})$	parametrized quantum state	Eq. (7) and Sec. 2.1
$U(\boldsymbol{\theta})$	parametrized variational ansatz	Eq. (8) and Eq. (8)
O, B	observable that defines the cost function	Eq. (12), Sec. 2.2 and Fig. 1 resp.
$\mathcal{E}(\rho)$	quantum noise channel	Eq. (47) and Sec. B
$\text{Herm}(\mathbb{C}^n)$	space of Hermitian operators over $\mathbb{C}^n$	Before Eq. (7)
$\text{Hilb}(\mathbb{C}^n)$	Hilbert space over $\mathbb{C}^n$	Eq. (47) and Sec. B
Δ_k	k -th simultaneous perturbation vector	Sec. 2.3 and Eq. (18)
μ	smoothing parameter > 0	Eq. (20)
f_μ	Gaussian smoothing function	Eq. (20)
N	number of SA iterations	Sec. 2.5
L	modulus of Lipschitz continuity	Eqs. (21)
γ	step size parameter	Eq. (23)
C_n	constant > 0	Assumption 2.1
c, c_k	set of scalars (errors) in the gradient approximations	Eq. (18) and Eq. (19)
b, b_k	upper bound of bias	Eq. (24)
p, L_V	number of layers in the VQA	Eq. (8) and Eq. (33) resp.

Table 1: Summary of frequent notation used throughout the paper.

$\text{argmin}_{\boldsymbol{\theta}} C(\boldsymbol{\theta})$, corresponds to the solution of the problem. Usually, $C(\boldsymbol{\theta})$ is given as the expectation value of some operator $O \in \text{Herm}(\mathbb{C}^d)$,

$$\begin{aligned}
C(\boldsymbol{\theta}) &= \langle \psi(\boldsymbol{\theta}) | O \psi(\boldsymbol{\theta}) \rangle \\
&= \langle \psi_0 | U^\dagger(\boldsymbol{\theta}) O U(\boldsymbol{\theta}) | \psi_0 \rangle,
\end{aligned} \tag{7}$$

where $|\psi_0\rangle \in \mathbb{C}^d$ denotes the initially prepared state and $U \in U(d)$. Furthermore, it is assumed that the minimum of $C(\boldsymbol{\theta})$ is not efficiently computable using only a classical computer. (This requires, among others, the use of gates outside of the Clifford group and limits the applicability of common quantum error correction mechanisms.) Therefore, the quantum part of the information loop mentioned above amounts to:

- (1) Encoding (loading) the data into a parametrized quantum circuit.
- (2) Initializing the parameters $\boldsymbol{\theta} \in \mathbb{R}^p$ of the quantum circuit.
- (3) Measuring the parametrized output of the circuit $C(\boldsymbol{\theta})$ or its gradient.

Once an iteration of the quantum part of the loop has been performed, a classical optimizer is used to minimize the measured cost function with the aim of updating the parameters $\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}'$, $\boldsymbol{\theta}' \in \mathbb{R}^p$ such that $C(\boldsymbol{\theta}') \leq C(\boldsymbol{\theta})$. In Eq. (7), the unitary $U = U(\boldsymbol{\theta})$ is referred to as the *variational ansatz*, which explicitly contains precise information about the variational parameters $\boldsymbol{\theta} \in \mathbb{R}^p$ and corresponds to the actual quantity that is *classically trainable*. In the literature, one encounters problem-inspired or problem-specific ansätze, suited for certain problems, as well as generic ansatz architectures that are problem-agnostic [3].

Generically, in the simplest form, $U(\boldsymbol{\theta})$ is decomposed as the sequential application of unitaries parameterized by the parameters:

$$U(\boldsymbol{\theta}) = U_p(\theta_p) \dots U_1(\theta_1), \quad (8)$$

where each parameterized unitary $U_i \in U(d)$, where d is the level of the quantum system, takes the form

$$U_i(\theta_i) := W_i \exp(-i\theta_i H_i), \quad (9)$$

for $H_i \in \text{Herm}(\mathbb{C}^d)$, $W_i \in U(d)$ and $\theta_i \in \mathbb{R}$, for any $i \in [p]$. The unitaries W_i are not subject to variation by the classical optimizer, while Hermitian operators H_i , the *Hamiltonians*, are optimally chosen according to the problem at hand. For example, for a certain type of VQAs, the so-called *quantum approximate optimization algorithm* (QAOA) [27], as well as generalization thereof such as the quantum alternating operator ansatz [32], best suited for combinatorial optimization problems, one considers a repeating sequence of two types of Hamiltonians where H_{2i} corresponds to a sum of the Pauli σ_x matrices and H_{2i+1} corresponds to the Ising model Hamiltonian. Generally, there are many different types of ansatze, including ansatze for thermal states [33].

An important aspect of VQAs is the presence of systemic errors (bias). As a motivating example, consider the depolarizing channel of a density matrix

$$\mathcal{E}(\rho) = p \frac{\mathbb{1}}{d} + (1-p)\rho, \quad (10)$$

$\mathcal{E} \in \text{End}(\text{Herm}(\mathbb{C}^d))$, which replaces each qubit (or the whole system) in state ρ with a totally mixed state with some probability $p > 0$. Consider that subsequently, an observable O is applied to compute the energy, entropy or entanglement, for instance. It is clear that if in all circuit runs the state ρ is input into the operator $\mathcal{E}$ any evaluation will not have an expectation $\langle O \rangle := \text{Tr}(\rho O)$, rather $\langle \tilde{O} \rangle := \text{Tr}(\mathcal{E}(\rho) O)$, and no matter how many trials are performed, the sample average will not equal this desired value, that is,

$$\langle O \rangle \neq \mathbb{E}(\tilde{O}) := \frac{1}{N} \sum_{i=1}^N f_i, \quad f_i \sim \mathbb{E}_i(\tilde{O}), \quad (11)$$

where f_i corresponds to a measurement. Said differently, the estimator $\tilde{O}$ of some observable O is biased in the sense $\mathbb{E}[O - \tilde{O}] \neq 0$. It is known that, for example, on IBM machines, there exists an asymmetric read-out error [34]. See Appendix B for a numerical illustration.

Thus, it is clear that the noise arising from a classical measurement applied to a quantum circuit, as is commonly done in, e.g., VQA, is not unbiased, as far as the expectation of the (parameterized) estimator is concerned, and thus any analysis of algorithms with function evaluations of this sort must incorporate this for faithful modeling of the problem.

2.2 Parameter Shift Rule as a Gradient Estimate

Let us consider a VQA unitary in the form of Eq. (8). The cost function to be optimized takes the form

$$C(\boldsymbol{\theta}) = \langle \psi_0 | U^\dagger(\boldsymbol{\theta}) O U(\boldsymbol{\theta}) | \psi_0 \rangle, \quad (12)$$

with $U(\boldsymbol{\theta})$ as in Eq. (8). Typically, the unitaries executed in the context of VQAs are of the form of Eq. (9) where H_i are rotation-like operators such that $H_i^2 = \mathbb{1}$. This is the

case for all operators acting on single qubits; however, H_i can be any generic operator, and a viable strategy is to decompose them into Pauli strings [35], for example. In what follows we introduce the backbone of what is generically referred to as the *parameter shift rule* (PSR).

The gradient of Eq. (12) is of particular interest in this article. Precisely for the reasons mentioned in the previous paragraph, it suffices to analyze the case where the set of Hermitian operators $\{H_i\}$ are rotation-like operators, that is, given in terms of Pauli strings. Recall that any gate can be represented as a unitary operator

$$U_i(\theta_i) = e^{-i\theta_i H_i} = \cos(\theta_i)\mathbb{1} - i\sin(\theta_i)H_i. \quad (13)$$

Any operator of the form $O(\theta_i) := U_i^\dagger(\theta_i)OU_i(\theta_i)$, as in Eq. (12), can be written as

$$A + B\cos(\theta_i) + C\sin(\theta_i), \quad (14)$$

with A, B, C operators independent of the variational parameter θ_i and only dependent on the unitary H_i . Additionally, using standard trigonometric rules,

$$\begin{aligned} \frac{d\cos(x)}{dx} &= \frac{\cos(x+s) - \cos(x-s)}{2\sin(s)} \\ \frac{d\sin(x)}{dx} &= \frac{\sin(x+s) - \sin(x-s)}{2\sin(s)}, \end{aligned} \quad (15)$$

for $s \neq n\pi$ where $n \in \mathbb{Z}$, we can easily compute the derivative of $O(\theta_i)$ as

$$\frac{dO(\theta_i)}{d\theta_i} = \frac{O(\theta_i + s) - O(\theta_i - s)}{2\sin(s)}. \quad (16)$$

Therefore, computing the derivative with respect to one of the variational parameters amounts to computing the actual observable twice in symmetrically shifted points. The parameter-shift rule, as defined in [15, 36, 16] corresponds to $s = \pi/2$. Note that for any rotation-like operator H_i , the derivative (16) is exact (except for integer multiples of π). The gradient can be easily evaluated using the same exact rule by redefining the expectation value. That is, to compute the derivative with respect to θ_1 , we redefine $\tilde{O}_1 := U_2^\dagger(\theta_2) \dots U_L^\dagger(\theta_L)OU_L(\theta_L) \dots U_2(\theta_2)$ and $\tilde{O}_1(\theta_1) := U_1^\dagger(\theta_1)\tilde{O}_1U_1(\theta_1)$, while for $1 < j < L$ we define $\tilde{O}_j := U_{j+1}^\dagger(\theta_{j+1}) \dots U_L^\dagger(\theta_L)OU_L(\theta_L) \dots U_{j+1}(\theta_{j+1})$ and $|\tilde{\psi}_j\rangle = U_{j-1}(\theta_{j-1})|\psi_0\rangle$ such that the derivative with respect to θ_j is taken for $\tilde{O}_j(\theta_j) = U_j^\dagger(\theta_j)\tilde{O}_jU_j(\theta_j)$. The gradient is then given as

$$\nabla_{\theta}O(\theta) = \sum_{i=1}^L \frac{d\tilde{O}_i(\theta_i)}{d\theta_i}. \quad (17)$$

Higher-order derivatives can also be derived by iterations of the same rule [37].

For quantum variational algorithms, the parameter shift rule has emerged as a standard means by which to analytically compute the gradients of a parametrized circuit (see, e.g., [17]). The rule became popular because for noisy near-term quantum computers, it is quite difficult to implement accurate finite difference estimates since the two required shifted evaluations, of difference of order $O(10^{-7})$, cannot be effectively distinguished from the noise. However, the parameter-shift rules are not only exact, unlike the finite-difference methods, but also involve a relatively large shift away from the noise threshold. However, the expectation values required to compute Eq. (17) are still subject to the limitation

that only a finite number of shots will be available to construct empirical estimates, and since the ultimate objective evaluation is an expectation, a sampling error will appear. In addition, though, as we shall see in the discussion of bias in Sec. B, the evaluations will not be of the ideal desired sampled distribution but, as a result of systemic noise, biased. We shall study the effect of this on the convergence properties in the following section.

Description	Year	Name	Reference	Context	Exactness
A Small Sample of Classical Precursors					
Random Optimization method	1977	Dorea	[38]	OPT	No, stochastic in nature
Automatic Differentiation	1983	Wengert	[39]	OPT	Yes
SPSA	1992	Spall	[13]	OPT	No, stochastic in nature
Nesterov Smoothing	2005	Nesterov	[40]	OPT	Yes
Two-point Approximation	2015	Duchi <i>et. al.</i>	[41]	OPT	No, stochastic in nature
Gaussian Smoothing	2017	Nesterov & Spokoyny	[42]	OPT	No, stochastic in nature
Parameter Shift Rule (PSR)					
Gradient of expectation values of observables	2005	Khaneja <i>et. al.</i>	[43]	QOC	First-order approximation
Gradient evaluation with two parameters	2017	Li <i>et. al.</i>	[36]	QOC	First-order approximation
First paper on PSR with two parameters	2018	Mitarai <i>et. al.</i>	[15]	QML	Yes
Extension of [15] for more generic circuits	2018	Vidal & Theis	[44]	QML	Yes
Follow up of [15] and slight generalization	2019	Schuld <i>et. al.</i>	[16]	QML	Yes for Pauli generators and no for non-Pauli generators
Follow up of [16] allowing more generic gates	2019	Crooks	[35]	QML	Yes due to the Pauli decomposition of non-Pauli generators
Extension of PSR for generators with general eigenspectrum	2021	Mari <i>et. al.</i>	[37]	QML	Yes, iterative in nature
Generalization of the PSR for higher-order derivatives	2021	Izmaylov <i>et. al.</i>	[45]	QML	Yes, iterative in nature
Stochastic generalization of the original PSR	2021	Banchi & Crooks	[17]	QML	No, stochastic in nature
Stochastic generalization for gates with more than two eigenvalues	2021	Wierichs <i>et. al.</i>	[46]	QML	No, stochastic in nature
Fourier analysis for feasibility and characterization of PSRs	2021	Theis	[47]	OPT	N/A
Four-term PSR for fermionic Hamiltonians	2021	Kottmann <i>et. al.</i>	[48]	OPT	Yes
Quantum gradient computation with classical resources	2022	Koczor & Benjamin	[49]	QML	No, classical computation

Table 2: The evolution of parameter shift rule (PSR) approaches from the ancient (top) to the most recent (bottom). Acronyms used in context include QOC: Quantum Optimal Control, QML: Quantum Machine Learning, OPT: Optimization.

Several extensions to the “original” parameter shift rules have been discovered, including stochastic versions, all of which we summarize in Table 2. Let us review some of the landmark literature regarding the parameter shift rule. Parameter shift rules originate from the work on quantum optimal control [43, 36]. The “original” parameter shift rules, as they are currently understood in the context of VQAs, were first derived in [15, 44, 16]. In Ref. [35], generalized the “original” work for operators that are not given as Pauli strings by using gate decomposition to simpler gates. Similar generalizations were studied in Ref. [37] while in Ref. [45] parameter shift rules were derived for higher-order derivatives (Hessian, etc.). Ref. [17] produced an stochastic generalization of the parameter shift rules that




Figure 1: The evolution of parameter shift rule (PSR) approaches loosely based on [46, Fig. 1], from the ancient (top) to the most recent (bottom). Our methods apply to the PSR approaches highlighted in green. Following [46], for an observable B , we denote by $\langle B(x) \rangle$, that is, the expectation value at $x \in \mathbb{R}^n$. In the formula for the equidistant PSR, $x_\mu = \frac{2\mu-1}{2R}\pi$ and $f(\mu) = \frac{(-1)^{\mu-1}}{4R \sin(x_\mu/2)}$ (see Ref. [46, App. A]). In the equidistant General PSR, the General PSR and the General Stochastic PSR $\mathbb{E}_\mu$ is a function of $\langle B(x_\mu) \rangle$. Specifically, $\mathbb{E}(x) := \langle \psi | U^\dagger(x) B U(x) | \psi \rangle$ while for the General Stochastic PSR, the actual definition is provided below the corresponding PSR (for the precise definitions see Ref. [46]). Finally, in the above, U_F denotes the stochastic “unitary” $\exp i(H + F)$, with H the Hermitian generator with eigenvalues in the set $\{+1, -1\}$, F any other Hermitian (that does not have to commute with G), and $R \in \mathbb{N}$.

was claimed to be unbiased due to the inherent randomness of the quantum measurements. Using the quantum Fourier transform, Wierichs *et. al.* proposed a generalized parameter shift rule (abbreviated as PSR in Fig. 1) that is capable of significantly reducing the number of circuit evaluations required for the computation of the gradients with respect to gate parameters [46]. Ref. [47] discussed optimization/feasibility aspects of parameter shift rules using Fourier analysis. Finally, Ref. [49] proposed a classical algorithm for computing quantum gradients based on a conservative number of circuit evaluations at nearby points to a reference point.

2.3 Derivative-Free (Zero-Order) Optimization

The parameter-shift approach specific to VQA, as well as other popular algorithms for minimizing VQA objectives such as SPSA, fall under the class of derivative-free, or zero-order, optimization algorithms, meaning that they only use function evaluations, rather

than analytical gradient and higher derivative computations, in the iterative procedure of solving the problem.

In the classical optimization literature [50], this consists of two classes of strategies. The first class is mesh-based methods that have a set of predefined directions to consider at each iteration, with the function evaluated at some number of such directions and compared to each other and the original point. Alternatively, the so-called model-based methods use function evaluation information to approximate derivatives and higher-order information, and subsequently use some algorithm similar to classical derivative-based optimization procedures, such as trust region and gradient descent. There are a number of possible techniques to approximate a gradient using function evaluations, many of which are compared in [51]. The two we consider in this paper, the two-point function approximation and SPSA, are classic in their widespread use, in particular for the problems of interest in this paper, despite, however, their overall weaker precision and performance compared to alternatives as found in this work.

The SPSA method, introduced in [13] is defined to compute a stochastic gradient approximation at x by evaluating the noisy objective function at two random points. In particular, letting $\{\Delta_k\}$ be an independent set of random variables with $\mathbb{E}[\|\Delta_k\|^{-1}] \leq D$ and c_k a set of scalars with $c_k \rightarrow 0$, compute the gradient estimate with each component defined as,

$$[g]_i \sim \frac{F(\theta_k + c_k \Delta_k, \xi_k^1) - F(\theta_k - c_k \Delta_k, \xi_k^2)}{2c_k [\Delta_k]_i} \quad (18)$$

This was compared to schemes based on simple finite difference estimations of the gradient and it was found both theoretically and numerically that SPSA could achieve similar performance in convergence speed by SA iteration, while requiring fewer function evaluations for the approximation at each iteration [52].

We note the resemblance of this procedure to the notion of two-point function approximation popularized in [41] inspired by "Nesterov smoothing" [40]. See also [42] for a convergence analysis. In this case we evaluate,

$$g_k \sim \frac{F(\theta_k + c_k \Delta_k, \xi_k^1) - F(\theta_k, \xi_k^2)}{c_k} \Delta_k \quad (19)$$

We shall analyze the performance of this method alongside SPSA because of its strong resemblance to relying on two function evaluations to estimate the gradient. We note that ultimately this latter method has received more attention in the literature in the past decade, and thus is more known. However, SPSA tends to be the algorithm of choice for quantum applications. With these conceptual trade-offs, we shall cover both approaches in the subsequent analysis and experiments. A detailed overview of the available SPSA analyses, culminating in [24], is presented in the Appendix.

The two-point approximation [41] suggests that the gradient can be approximated by (19). The technique is based on the Gaussian smoothing technique [38, 42], which defines a function for a smoothing parameter corresponding to the perturbation for the function evaluation, here c_k . Considering a generic $c > 0$, we write:

$$f_c(\theta) = \frac{1}{(2\pi)^{\frac{n}{2}}} \int f(\theta + c\Delta) e^{-\frac{1}{2}\|\Delta\|^2} d\Delta = \mathbb{E}_{\Delta \sim \mathcal{N}(0, I)}[f(\theta + c\Delta)] \quad (20)$$

and satisfies some nice properties that shall be used in the sequel (see, e.g., [40, 53]) for

L -Lipschitz smooth $f(\boldsymbol{\theta})$,

$$\nabla f_c(\boldsymbol{\theta}) = \frac{1}{(2\pi)^{\frac{n}{2}}} \int \frac{f(\boldsymbol{\theta} + c\boldsymbol{\Delta}) - f(\boldsymbol{\theta})}{c} \boldsymbol{\Delta} e^{-\frac{1}{2}\|\boldsymbol{\Delta}\|^2} d\boldsymbol{\Delta} \quad (21a)$$

$$|f_c(\boldsymbol{\theta}) - f(\boldsymbol{\theta})| \leq \frac{c^2}{2} Lp \quad (21b)$$

$$\|\nabla f_c(\boldsymbol{\theta}) - \nabla f(\boldsymbol{\theta})\| \leq \frac{c}{2} L(p+3)^{\frac{3}{2}} \quad (21c)$$

Subsequently, [54] contributed to supplying the convergence theory for the case that $f(\boldsymbol{\theta})$ is generally smooth but nonconvex, and 1) the problem is constrained— $\boldsymbol{\theta}$ is required to be in some compact set $\mathcal{X}$, or 2) the gradient $\nabla F(\boldsymbol{\theta}, \xi)$ is almost surely sparse. Furthermore, it is assumed that there is a global bound on the gradient norm. In this setting, function evaluations are assumed to be unbiased, that is, for $\hat{F} \sim F(\boldsymbol{\theta}, \xi)$, we have $\mathbb{E}[\hat{F}] = f(\boldsymbol{\theta})$.

2.4 A Review of Asymptotic Results for SPSA

We present the single work in the literature that we were able to find with guarantees for SPSA under conditions of biased but bounded noise on the function evaluations, namely [24]. In this work, several assumptions are made regarding the noise and the algorithm, and an *asymptotic* result is given, indicating the eventual convergence to an approximately optimal point. We note that this work considers a time-varying instance where the objective function to minimize is changing over time, and as such, we simplify the original results for the purposes of our work.

Now we review the assumptions in the paper, considering that an evaluation of the noise function f generates a value in the form of (6), which for reference we repeat below:

$$f' \sim F(\boldsymbol{\theta}, \xi) = f(\boldsymbol{\theta}) + n(\boldsymbol{\theta}, \xi), \quad \xi \in \chi$$

Assumption 2.1 [24] *It holds that*

(1) *It holds that $\mathbb{E}[\|(n(\boldsymbol{\theta}_k, \xi_k) - n(\boldsymbol{\theta}_{k-1}, \xi_{k-1}))\|^2] \leq C_n$ for some $C_n > 0$.*

(2) *The function $f(\boldsymbol{\theta})$ is μ -strongly convex, i.e.,*

$$\langle \boldsymbol{\theta} - \boldsymbol{\theta}^*, \mathbb{E}[\nabla F(\boldsymbol{\theta}, \xi)] \rangle \geq \mu \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|^2$$

for some unique minimum $\boldsymbol{\theta}^*$, for all $\boldsymbol{\theta}$.

(3) *It holds that for the unique minimum,*

$$\mathbb{E}[\|\nabla F(\boldsymbol{\theta}^*, \xi)\|^2] \leq G$$

for some $G > 0$

(4) *The gradient is L -Lipschitz for all ξ , that is,*

$$\|\nabla F(\boldsymbol{\theta}, \xi) - \nabla F(\boldsymbol{\theta}', \xi)\| \leq L \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|$$

Among these assumptions, given that the objective function corresponds to the realization of a quantum circuit, and the parameters $\boldsymbol{\theta}$ rotate an element of a Lie group (the quantum gates), and each source of potential noise is bounded (see Section B above), assumptions related to boundedness of the variance clearly holds. In addition, from Section 2.1 it is

clear that all the objective function is continuously differentiable with respect to θ , being matrix exponentials, and thus possibly expressed as trigonometric functions, which are all globally Lipschitz continuous.

However, the second assumption, requiring strong convexity, is one that is rarely satisfied in our running example of minimizing functions that evaluate runs of quantum circuits. The imaginary exponentials, equivalently trigonometric functions, that are characteristic of VQA objectives are oscillatory in nature and thus exhibit both positive and negative curvature, thus violating this condition.

The main Theorem in that work can be summarized as,

Theorem 2.1 [24, Theorem 1] *It holds that for all $\epsilon > 0$ there exists some K such that for $k \geq K$,*

$$\sqrt{\mathbb{E} \|\theta_k - \theta^*\|^2} \leq M + \epsilon$$

where M depends on problem and algorithmic constants.

Thus for this case there is a mean squared error result asymptotically.

In this paper we seek to analyze the method to derive a non-asymptotic result, i.e., one that describes the rate of convergence, and do so applicably for nonconvex functions. For insight we turn to other related work.

2.5 A Review of Guarantees for Two Point Function Approximation

We consider now the SA iteration with the gradient approximation g_k given by (19). The work [41] presents convergence guarantees for the smooth convex case as well as the convex but possibly nonsmooth set of functions.

Subsequently, [54] contributed to supplying convergence theory for the case that $f(\theta)$ is generally smooth but nonconvex, and 1) the problem is constrained— θ is required to be in some compact set $\mathcal{X}$, or 2) the gradient $\nabla F(\theta, \xi)$ is almost surely sparse. Furthermore, it is assumed that there is a global bound on the gradient norm. In our context, the presence of $\mathcal{X}$ could correspond to some a priori knowledge in regards to bounds on the permissible θ , which would not be standard in practice, i.e., typically the θ are taken to be unconstrained and arbitrary in $\mathbb{R}^p$.

With the Frank-Wolfe gap defined as $g_G^k := \langle \nabla f(\theta_k), \theta_k - \hat{\theta} \rangle$, $\hat{\theta} := \operatorname{argmin}_{u \in \mathcal{X}} \langle \nabla f(\theta_k), u \rangle$, the work demonstrates in [54, Theorem 2.1] that $\mathbb{E}[g_G^R] \leq \frac{C}{\sqrt{N}}$, with N the number of iterations and the constant C depending on the noise, the gradient bound, and the size of the feasible set. The superscript R indicates a randomly chosen combination of iterations.

In this setting, function evaluations are assumed to be unbiased, i.e., for $\hat{F} \sim F(\theta, \xi)$, it holds that $\mathbb{E}[\hat{F}] = f(\theta)$.

Finally, closest to the analysis framework we use here, the seminal work [53] considers both gradient and zero-order schemes for stochastic optimization (still unbiased, in both cases). In this case, one considers a fixed budget of N iterations, and with $c_k \equiv c = O(1/\sqrt{N})$ and $\alpha_k \equiv \alpha = O(1/\sqrt{N})$ one obtains that the gradient at a particular randomly chosen combination of iterates is bounded by a term that is itself $O(1/\sqrt{N})$ for generally nonconvex functions plus an additive term indicating the bias due to the approximation accuracy. This is the same asymptotic rate as in the case of unbiased SGD, however, now with only an approximate stationarity guarantee, and larger constant terms appearing in the expressions for the convergence rate.

2.6 A Review of Guarantees for Biased SGD

Ref. [20] explores the asymptotic guarantees of biased stochastic approximation under the model that in (3) we have $\mathbf{n}(\boldsymbol{\theta}_k, \xi_k) = \boldsymbol{\zeta}_k + \boldsymbol{\eta}_k$ with $\mathbb{E}[\boldsymbol{\zeta}_k] = 0$ and $\limsup_{k \rightarrow \infty} \|\boldsymbol{\eta}_k\| < \infty$ and prove, depending on the regularity of the function, various bounds with regard to $\limsup_{k \rightarrow \infty} \|\nabla f(\boldsymbol{\theta}_k)\|$ and $\text{dist}(\boldsymbol{\theta}_k, \boldsymbol{\Theta}^*)$ with $\boldsymbol{\Theta}^*$ the set of stationary points.

More recently, with quantitative bounds, the work in [22] considers the model for biased SGD (3) by:

$$\mathbf{n}(\boldsymbol{\theta}, \xi) = \mathbf{b}(\boldsymbol{\theta}) + \mathbf{r}(\boldsymbol{\theta}, \xi), \quad \|\mathbf{b}(\boldsymbol{\theta})\| \leq m \|\nabla f(\boldsymbol{\theta})\|^2 + \zeta^2, \quad \mathbb{E}_\xi \|\mathbf{r}(\boldsymbol{\theta}, \xi)\|^2 \leq M \|\nabla f(\boldsymbol{\theta}) + \mathbf{b}(\boldsymbol{\theta})\|^2 + \sigma^2 \quad (22)$$

they obtain:

Lemma 2.1 [22, Lemma 3] *If $\alpha_k \equiv \alpha \leq \frac{1}{(1+M)L}$ then it holds that after K iterations,*

$$\frac{1}{K} \sum_{k=1}^K \mathbb{E} \|\nabla f(\boldsymbol{\theta}_k)\|^2 \leq \frac{2F}{K\alpha(1-m)} + \frac{\alpha L \sigma^2}{1-m} + \frac{\zeta^2}{1-m}, \quad (23)$$

where F is the distance between $f(\boldsymbol{\theta}_0)$ and the local minimum attained.

Theorem 2.2 [22, Theorem 4] *For $\epsilon > 0$, if $\alpha = \min \left\{ \frac{1}{(1+M)L}, \frac{\epsilon(1-m)+\zeta^2}{2L\sigma^2} \right\}$, then with $K = \mathcal{O} \left(\frac{M+1}{\epsilon(1-m)+\zeta^2} + \frac{\sigma^2}{\epsilon^2(1-m)^2+\zeta^4} \right)$ LF iterations, it holds that*

$$\frac{1}{K} \sum_{k=1}^K \mathbb{E} \|\nabla f(\boldsymbol{\theta}_k)\|^2 = \mathcal{O} \left(\epsilon + \frac{\zeta^2}{1-m} \right).$$

For completeness, we shall also mention [21] which presents a (complementary) non-asymptotic analysis of the convergence guarantees associated with biased SGD under a similar model.

Thus for circuits in which the gradient is computed analytically these results indicate the ultimate performance and guarantees in regards to the algorithm's output. As an insight into the sequel, we can expect that using zero-order methods to approximately estimate the gradient shall result in at least as similar perturbations, as in relative inaccuracy to stationarity.

3 New Iteration Complexity Results

Our main result is the analysis of iteration complexity of the two-point approximation and the SPSA algorithm under the setting of biased function evaluations. We remind the reader that we consider the estimation as based on the expansion given in (6). That is: the biased function evaluation f' is the true function $f(\boldsymbol{\theta})$ corrupted by noise $\mathbf{n}(\boldsymbol{\theta}, \xi)$:

$$f' \sim F(\boldsymbol{\theta}, \xi) = f(\boldsymbol{\theta}) + \mathbf{n}(\boldsymbol{\theta}, \xi), \quad \xi \in \chi(\boldsymbol{\theta}),$$

wherein $\boldsymbol{\theta} \in \mathbb{R}^p$ are the deterministic VQA parameters and $\xi \in \mathbb{R}$ is the sample generated from the random field, as in Eqs. (1) and (6).

We make two assumptions. First, we assume the noise $\mathbf{n}(\boldsymbol{\theta}_k, \xi_k)$ associated with a function evaluation is a mixture of a mean-zero component $\mathbf{r}_k$ and a biased component $\mathbf{b}_k$, wherein the bias is bounded from above $\|\mathbf{b}_k\| \leq b$, almost surely (a.s.).

Assumption 3.1 The function evaluation (6) satisfies,

$$n(\boldsymbol{\theta}_k, \xi_k) = r_k + b_k, \text{ with } \mathbb{E}[r_k] = 0, \mathbb{E}\|r_k\|^2 \leq \sigma^2, \text{ and } \|b_k\| \leq b \text{ a.s.} \quad (24)$$

Indeed, recalling the form for the objective function for VQA, starting from the discussion from (8),

$$C(\boldsymbol{\theta}) = \langle \psi_0 | U^\dagger(\boldsymbol{\theta}) O U(\boldsymbol{\theta}) | \psi_0 \rangle, \text{ with } U(\boldsymbol{\theta}) = U_p(\theta_p) \dots U_1(\theta_1) \text{ and } U_i(\theta_i) = W_i \exp(-i\theta_i H_i)$$

In considering noise as an error that takes place at one of the gates, this would correspond to a change of H_i . Note that regardless, the modulus of the exponential term is one, $|\exp(-i\theta_i \tilde{H}_i)| = 1$ for any perturbed $\tilde{H}_i$. For illustration, let us compute the error corresponding to an evaluation of a three layer circuit, with a gate error in the middle layer:

$$\begin{aligned} & \langle \psi_0 | W_3 \exp(i\theta_3 H_3) W_2 \exp(i\theta_2 \tilde{H}_2) W_1 \exp(i\theta_1 H_1) O W_3 \exp(-i\theta_3 H_3) W_2 \exp(-i\theta_2 \tilde{H}_2) W_1 \exp(-i\theta_1 H_1) | \psi_0 \rangle \\ & - \langle \psi_0 | W_3 \exp(i\theta_3 H_3) W_2 \exp(i\theta_2 H_2) W_1 \exp(i\theta_1 H_1) O W_3 \exp(-i\theta_3 H_3) W_2 \exp(-i\theta_2 H_2) W_1 \exp(-i\theta_1 H_1) | \psi_0 \rangle \\ & = \langle \psi_0 | W_3 \exp(i\theta_3 H_3) W_2 \exp(i\theta_2 \tilde{H}_2) W_1 \exp(i\theta_1 H_1) O W_3 \exp(-i\theta_3 H_3) W_2 \exp(-i\theta_2 \tilde{H}_2) W_1 \exp(-i\theta_1 H_1) | \psi_0 \rangle \\ & - \langle \psi_0 | W_3 \exp(i\theta_3 H_3) W_2 \exp(i\theta_2 H_2) W_1 \exp(i\theta_1 H_1) O W_3 \exp(-i\theta_3 H_3) W_2 \exp(-i\theta_2 \tilde{H}_2) W_1 \exp(-i\theta_1 H_1) | \psi_0 \rangle \\ & + \langle \psi_0 | W_3 \exp(i\theta_3 H_3) W_2 \exp(i\theta_2 H_2) W_1 \exp(i\theta_1 H_1) O W_3 \exp(-i\theta_3 H_3) W_2 \exp(-i\theta_2 \tilde{H}_2) W_1 \exp(-i\theta_1 H_1) | \psi_0 \rangle \\ & - \langle \psi_0 | W_3 \exp(i\theta_3 H_3) W_2 \exp(i\theta_2 H_2) W_1 \exp(i\theta_1 H_1) O W_3 \exp(-i\theta_3 H_3) W_2 \exp(-i\theta_2 H_2) W_1 \exp(-i\theta_1 H_1) | \psi_0 \rangle \\ & \leq \langle \psi_0 | W_3 \exp(i\theta_3 H_3) W_2 (\exp(i\theta_2 \tilde{H}_2) - \exp(i\theta_2 H_2)) W_1 \exp(i\theta_1 H_1) O W_3 \exp(-i\theta_3 H_3) W_2 \exp(-i\theta_2 \tilde{H}_2) W_1 \exp(-i\theta_1 H_1) | \psi_0 \rangle \\ & + \langle \psi_0 | W_3 \exp(i\theta_3 H_3) W_2 \exp(i\theta_2 H_2) W_1 \exp(i\theta_1 H_1) O W_3 \exp(-i\theta_3 H_3) W_2 (\exp(-i\theta_2 \tilde{H}_2) - \exp(-i\theta_2 H_2)) W_1 \exp(-i\theta_1 H_1) | \psi_0 \rangle \end{aligned}$$

Recalling that $W_i \in U(d)$, and operations $\exp(i\theta H)$ are norm preserving, we can obtain a bound as based on the operation error. Namely, the observation operator O is the one component of the function that can conscript or enlarge the norm of the input. Thus, an explicit bound can be characterized as,

$$b = 2 \|\exp(i\theta H) - \exp(i\theta \tilde{H})\| \sigma^M(O)$$

where $\sigma^M(O)$ is the largest eigenvalue of O , which is of course constant for a given problem, producing the upper bound. Any statistically symmetric additional noise can be considered as incorporated into r_k .

Second, we have the standard assumption on the Lipschitz continuity of the gradient. From the equations above explicitly detailing the function evaluations, it is clear that the expressions are entirely analytic in the parameters θ :

Assumption 3.2 The gradient is L -Lipschitz for all ξ , that is,

$$\|\nabla F(\boldsymbol{\theta}, \xi) - \nabla F(\boldsymbol{\theta}', \xi)\| \leq L \|\boldsymbol{\theta} - \boldsymbol{\theta}'\| \quad (25)$$

To comment on the applicability to VQA: a Lipschitz constant on the gradient of a function corresponds to the upper bound on the largest Hessian singular value. Note that the use of ξ in the definition is ideal noise. In this case, recalling that each gate is norm preserving, in between a map from and onto the Bloch sphere, the Lipschitz bound would be $L = \sigma^M(O)$.

Non-Asymptotic Convergence Guarantees for Biased Two-Point Approximation We now state our first main result on the convergence rate of the two-point approximation with bias. In particular, we consider two regimes:

- (1) one in which we analyze the performance after a given number K of iterations,
- (2) one in which we analyze the performance of a mixture $\boldsymbol{\theta}_R$ of the iterates with a diminishing step size. That is, $\boldsymbol{\theta}_R$ is a random variable where each iterate (out of K iterates, with K countable) is sampled with a probability that depends on the diminishing step size.

Across both regimes, we bound the expectation of a norm of the gradient, which could be seen as a measure of the distance to the nearest critical point. Across both regimes, the bound involves three terms:

- one that is independent of the noise and depends only on the modulus L of Lipschitz continuity of the gradient (25) and a certain measure of the objective range $4(f_c(\boldsymbol{\theta}_k) - f_c^*)$, where f_c^* is a lower bound on the objective
- one that depends on the variance σ^2 (24) of the expectation-zero noise (24) and on the magnitude b of the bias (24),
- one that depends on the magnitude b of the bias (24).

The first two are seen to asymptotically decrease to zero with the iterations, while the last presents a lower bound on the measure of optimality. Let us recall the smoothed function f_c from (21).

Theorem 3.1 *Consider the two-point function approximation algorithm with iteration (4) for a problem that meets the Assumptions 3.1 and 3.2.*

- (1) *Consider a fixed budget of K iterations. Let $\alpha_k \equiv \alpha/\sqrt{K}$ with $\alpha \leq L/4$. It holds that the iterates satisfy:*

$$\min_{k \in [K]} \mathbb{E} \|\nabla f(\boldsymbol{\theta}_k)\|^2 \leq \frac{4(f_c(\boldsymbol{\theta}_k) - f_c^*)}{\alpha\sqrt{K}} + \frac{2L\alpha(\sigma^2 + \frac{8b^2}{c^2})}{\alpha\sqrt{K}} + \left[\frac{2b^2}{c^2} + \frac{c^2}{2}L(p+3)^3 \right] \quad (26)$$

- (2) *Consider a diminishing stepsize of $\alpha_k = \alpha/k^\gamma$ with $\gamma \in (1/2, 1]$ and $\alpha \leq L/4$. Then it holds that, taking,*

$$\boldsymbol{\theta}_R \sim \boldsymbol{\theta}_k \text{ w.p. } \alpha_k / \sum_{l=1}^K \alpha_l$$

we have,

$$\mathbb{E} \|\nabla f(\boldsymbol{\theta}_R)\|^2 \leq \frac{4(f_c(\boldsymbol{\theta}_k) - f_c^*)}{\alpha K^\gamma} + \frac{2L\alpha(\sigma^2 + \frac{8b^2}{c^2})}{\alpha K^\gamma} + \left[\frac{2b^2}{c^2} + \frac{c^2}{2}L(p+3)^3 \right] \quad (27)$$

The proof is given in Appendix A on page 35. Note that this implies that regardless of the number of iterations performed, the measure of stationarity will always have a lower bound given by the right-hand side (“effect of bias”), which is:

$$\frac{2b^2}{c^2} + \frac{c^2}{2}L(p+3)^3. \quad (28)$$

When we minimize the effect of bias (28), as a function of c , we obtain the following:

$$c^* = \left[\frac{4b^2}{L(p+3)^3} \right]^{1/4}. \quad (29)$$




Figure 2: The (logarithm of the) bound surface of the optimal perturbation c as a function of b, p .

Non-Asymptotic Guarantees for SPSA Now, we consider applying a similar line of analysis to SPSA. Recall that the noisy gradient estimate is of the form (18)

$$[\mathbf{g}_k]_i \sim \frac{F(\boldsymbol{\theta}_k + c_k \boldsymbol{\Delta}_k, \xi_k^1) - F(\boldsymbol{\theta}_k - c_k \boldsymbol{\Delta}_k, \xi_k^2)}{2c_k [\boldsymbol{\Delta}_k]_i}.$$

We shall make use of a Lemma in the literature characterizing the bias associated with this estimate in the case of being able to compute unbiased samples of the noisy objective function itself.

Lemma 3.1 [13, Lemma 1] Suppose that $\{\boldsymbol{\Delta}_k\}$, and its components, are i.i.d. and symmetrically distributed around zero with $[\boldsymbol{\Delta}_k]_i \leq B_0$ almost surely and $\mathbb{E}[[\boldsymbol{\Delta}_k]_i^{-1}] \leq B_1$, with $B_0, B_1 > 0$. Furthermore, assume that for some K , for all $k \geq K$, there exists $\delta > 0, B_2 > 0$ such that $\left\| \frac{\partial^3}{\partial \boldsymbol{\theta}^3} f(\boldsymbol{\theta}_k) \right\| \leq B_2$ for all $\boldsymbol{\theta}$ in a δ -neighborhood around $\boldsymbol{\theta}_k$. Then it holds that, for such $\boldsymbol{\theta}_k$, the SPSA estimate (18) without bias in function evaluations yields

$$\mathbf{g}_k - \nabla f(\boldsymbol{\theta}_k) = \mathbf{r}(\xi_k^1, \xi_k^2, \boldsymbol{\Delta}_k) + \mathbf{b}_0(\boldsymbol{\theta}_k), \mathbb{E}[\mathbf{r}(\xi_k^1, \xi_k^2, \boldsymbol{\Delta}_k)] = 0, \|\mathbf{b}_0(\boldsymbol{\theta}_k)\| \leq b_0 c_k^2 \quad (30)$$

for some b_0 depending on B_0, B_1, B_2 , and p .

We will also need an additional assumption:

Assumption 3.3 The perturbation is bounded away from zero, i.e., there exists $B_3 > 0$ such that $[\boldsymbol{\Delta}_k]_i^{-1} \leq B_3$ almost surely for all k .

With this, we are able to prove the convergence rate of the SPSA procedure with bias. Again, we consider two regimes. One, where we analyze the performance after a given number K of iterations. Second, we analyze the performance of mixture of the outcomes

of iterations with a diminishing step size. Across both regimes, we bound the expectation of a norm of the gradient as a measure of the distance to the nearest stationary point. Again, the bounds involve three terms, the first one independent of the noise and approaching zero with the iterations, the second one depending on the noise and bias but also approaching zero, and the third the asymptotic bias.

We incorporate the previous Lemma in the following sense in the proof of the Theorem, which is also in the Appendix: we consider the SPSA procedure in the case wherein there is no bias in function evaluations, then consider a perturbation of this. Finally:

Theorem 3.2 *Consider the SPSA algorithm with the iteration (4) and gradient estimate (18) for a problem satisfying Assumptions 3.1 and 3.2. Also, let the SPSA algorithm satisfy the conditions of Lemma 3.1 and Assumption 3.3. Then it holds that*

- (1) *Consider a fixed budget of K iterations. Let $\alpha_k \equiv \alpha/\sqrt{K}$ with $\alpha \leq L/2$. It holds that the iterates satisfy:*

$$\min_{k \in [K]} \mathbb{E} \|\nabla f(\boldsymbol{\theta}_k)\|^2 \leq \frac{4(f_c(\boldsymbol{\theta}_k) - f_c^*)}{\alpha\sqrt{K}} + \frac{2L\alpha \left[\frac{\sigma^2 B_3^2}{c_k^2} + 2b_0 c_k^2 + \frac{2b^2 B_3^2}{c_k^2} \right]}{\sqrt{K}} + \left[b_0 c_k^2 + \frac{b^2 B_3^2}{c_k^2} \right] \quad (31)$$

- (2) *Consider a diminishing stepsize of $\alpha_k = \alpha/k^\gamma$ with $\gamma \in (1/2, 1]$ and $\alpha \leq L/2$. Then it holds that, taking,*

$$\boldsymbol{\theta}_R \sim \boldsymbol{\theta}_k \text{ w.p. } \alpha_k / \sum_{l=1}^K \alpha_l$$

we have,

$$\mathbb{E} \|\nabla f(\boldsymbol{\theta}_R)\|^2 \leq \frac{4(f_c(\boldsymbol{\theta}_k) - f_c^*)}{\alpha K^\gamma} + \frac{2L\alpha \left[\frac{\sigma^2 B_3^2}{c_k^2} + 2b_0 c_k^2 + \frac{2b^2 B_3^2}{c_k^2} \right]}{K^\gamma} + \left[b_0 c_k^2 + \frac{b^2 B_3^2}{c_k^2} \right] \quad (32)$$

The proof is in Appendix A on page 37.

Guarantees for the Parameter Shift Rule Now, we consider the original Parameter Shift rule as given by (16), rewritten for convenience below with $s = \frac{\pi}{2}$,

$$\frac{dO_i(\theta_i)}{d\theta_i} = \frac{O_i(\theta_i + \frac{\pi}{2}) - O_i(\theta_i - \frac{\pi}{2})}{2}, \quad \nabla O(\boldsymbol{\theta}) = \sum_{i=1}^L \frac{d\tilde{O}_i(\theta_i)}{d\theta_i}$$

Now, we have that, without bias in the function evaluations, this is an *exact* gradient computation. For noisy operations under the ideal probability distribution Ξ , this would be an unbiased stochastic gradient. For the real noise distribution χ , however, biased noise will appear from the evaluations of $O(\theta_i + \frac{\pi}{2})$ and $O(\theta_i - \frac{\pi}{2})$. By our biased noise model given in Assumption 3.1 we have that,

$$n(\boldsymbol{\theta}, \xi) = O(\boldsymbol{\theta}, \xi \sim \Xi) + (O(\boldsymbol{\theta}, \xi \sim \chi) - O(\boldsymbol{\theta}, \xi \sim \Xi))$$

and

$$b = L_V \max_i \sup_{\xi \in \chi} \left| O_i\left(\theta_i \pm \frac{\pi}{2}, \xi \sim \chi\right) - O_i\left(\theta_i \pm \frac{\pi}{2}\right) \right| \quad (33)$$




Figure 3: The bound surface of Eq. (31) with $L = 8$, $\alpha = 4$, $\sigma = 1$, $K = 10^5$ for $0 \leq b, c \leq 1$. The z -axis is in logarithmic scale.

where L_V are the number of layers in the VQA, to distinguish it from the Lipschitz constant. Now, given that the ideal PS computation is the exact stochastic gradient, we are in the biased SGD regime, with the noise setting (22) with $\zeta^2 = b$, and it can be seen that we can take $m = M = 0$. From (23) we have,

$$\frac{1}{K} \sum_{k=1}^K \|\nabla f(\theta_k)\|^2 \leq \frac{2F}{K\alpha} + \alpha L\sigma^2 + \zeta^2 = \frac{2F}{K\alpha} + \alpha L\sigma^2 + b \quad (34)$$

Discussion and Comparison Let us compare the convergence guarantees for the two point function approximation, SPSA and Parameter Shift rule by (37), (31) and (34), respectively.

- (1) It appears that all algorithms have the same overall convergence rate.
- (2) The error grows at least with the square of the maximum eigenvalue of the observable $\sigma_M(O)$ through the influence of this factor on L and b .
- (3) The Parameter Shift Rule (PSR) exhibits the smallest effect of bias, in general, as well as the smallest constants for the convergence, which is reasonable given it is an exact gradient computation
- (4) SPSA, however, seems to be the most tuneable, i.e., it is clear that B_3 and c_k are freely chosen by the user, and as such can be adjusted so as to significantly mitigate the bias. Thus, in cases wherein significant bias mitigation is called for in order to obtain sensible result, aggressive parameter choices can be made in SPSA to achieve this. Note that the influence of the circuit depth, corresponding to p , does not influence the SPSA convergence except via its influence of b_c , in contrast to the two point function approximation.




Figure 4: The QAOA circuit for 4 qubits. Within the red crossed lines we denote one layer of the QAOA and with dots we denote the repetition of the layers to the desired depth.

- (5) The two-point function approximation would perform better than SPSA for generic derivative-free optimization using classical algorithmic parameter choices, but is clearly the option of last resort in variational quantum algorithms, given that there is only one parameter c to adjust to mitigate the noise (compared to SPSA's two: c and B_3), while at the same time underperforming PSR, if the circuit is amenable to gradient computations using PSR. Furthermore, the convergence error scales with the cube of the circuit depth, or number of variational parameters, q , due to the additional bias incurred with greater dimensional problems in the two point function approximation to the gradient.
- (6) Taking more samples in each function evaluation would reduce the symmetric noise, and thus σ^2 , which would also have a favorable effect on both the rate of convergence and quality of the final iterate.
- (7) The final bias scales with the error profile as defined by the worst norm difference that occurs in the case of a modified operation, $\|\exp(i\theta H) - \exp(i\theta \tilde{H})\|$.

4 Numerical Experiments

In this section, we illustrate the iteration complexity of VQAs in the context of a standard QAOA test cast, the MAX-CUT problem on a 4-regular graph.

The Quantum Approximate Optimization Algorithm [27] is a VQA that can tackle a variety of quadratic unconstrained binary optimization problems (QUBOs) by finding approximate solutions. The way it works, as a VQA, is to maximize the cost function of output bit strings. It is a (quantum) random search algorithm where, e.g., for the MAX-CUT problem, it can be viewed as a sequence of quantum random walks on the hypercube defined by the problem. At each iteration, the algorithm provides an improved cost function by optimizing it and feeding it back to the quantum computer for the next iteration. In the standard QAOA one starts with the uniform quantum superposition state and then two different Hamiltonians are encoded, the mixer Hamiltonian H_M , and the problem Hamiltonian H_C which is equivalent to the $d = 1$ Ising model Hamiltonian. The corresponding circuit is shown in Figure 4.

The standard quantum circuit is shown in Fig. 4, where the gates are defined as usual (cf. Appendix C).

In order to perform tests with varying amounts of noise and varying bounds on the bias therein, we utilize the Qiskit **Aer** simulator [55] with the native gateset comprised of identity, R_z , $\sqrt{X}$, and CNOT gates, to which we transpile the U_2 and $R(\theta, \phi)$ gates of Fig. 4, and a custom noise model. The custom noise model captures three sources of error:

- imperfect fidelity of 1-qubit gates: we associate each application of a 1-qubit gate (after transpiling to the native gateset) with a depolarizing channel (cf. p. 38) acting upon ρ as:

$$\mathcal{E}_1(\rho) := 0.999\rho + 0.001\text{Tr}(\rho)\frac{\mathbb{1}}{2} \quad (35)$$

- imperfect fidelity of 2-qubit gates: we associate each application of a CNOT gate (i.e., after transpiling $R(\theta, \phi)$ gates of Fig. 4 to the native gateset with CNOT only) with a depolarizing channel acting on ρ as:

$$\mathcal{E}_2(\rho) := 0.99\rho + 0.01\text{Tr}(\rho)\frac{\mathbb{1}}{4} \quad (36)$$

- bias in the read-out: we consider a matrix P defining the so-called assignment probabilities, i.e., conditional probability of recording a noisy measurement outcome as α given that the measurement outcome should have been β , for all representations α, β of bit-strings on 4 qubits. We detail two such matrices P below.

In what follows, we first test the optimization performance for a random choice of assignment probabilities within the read-out, while afterwards we define a custom read-out noise model based on the average 1-qubit local read-out errors of some standard freely available low-qubit IBMQ devices; see Table 3. The need of such a custom noise model stems from the fact that it becomes practically impossible to distinguish between the gate errors and the read-out error for anything more than one qubit.

ibmq_	armonk	lima	santiago	bogota	belem	quito	manila
qubits	1	5	5	5	5	5	5
$p(1 0)$	0.0322	0.0078	0.0274	0.0496	0.0052	0.0186	0.0114
$p(0 1)$	0.0358	0.0260	0.0448	0.1916	0.0352	0.0548	0.034
bias $\tilde{b}$	0.00359	0.0182	0.0174	0.1402	0.0300	0.0362	0.0226

Table 3: Assignment probabilities for qubit-flipped measurements across various IBM devices. Observe that all entries of the last row are larger than those of the row before for all devices, clearly showcasing the read-out asymmetry.

A random noise model

First, we revisit the convergence of VQA utilizing SPSA and a randomly (and highly non-physical) probability transition matrix P_{rand} :

$$P_{\text{rand}} = \begin{pmatrix} 0.55 & 0.27 & 0.17 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.58 & 0.14 & 0.16 & 0.04 & 0.00 & 0.02 & 0.04 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.10 & 0.10 & 0.60 & 0.13 & 0.01 & 0.01 & 0.02 & 0.00 & 0.00 & 0.01 & 0.01 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.32 & 0.64 & 0.01 & 0.01 & 0.02 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.30 & 0.68 & 0.01 & 0.00 & 0.01 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.02 & 0.52 & 0.44 & 0.00 & 0.00 & 0.01 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.90 & 0.06 & 0.01 & 0.01 & 0.00 & 0.02 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.94 & 0.02 & 0.02 & 0.01 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.25 & 0.29 & 0.13 & 0.28 & 0.00 & 0.00 & 0.00 & 0.02 & 0.02 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.35 & 0.05 & 0.38 & 0.19 & 0.01 & 0.01 & 0.01 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.32 & 0.27 & 0.09 & 0.07 & 0.19 & 0.03 & 0.03 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.81 & 0.10 & 0.06 & 0.01 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.41 & 0.02 & 0.24 & 0.01 & 0.30 & 0.01 & 0.00 & 0.01 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.42 & 0.46 & 0.07 & 0.02 & 0.01 & 0.00 & 0.02 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.15 & 0.18 & 0.64 & 0.03 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.10 & 0.75 & 0.08 & 0.03 & 0.01 & 0.00 & 0.01 & 0.01 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \end{pmatrix}$$

This matrix was randomly generated subject to the conditions of sparsity and row-stochasticity. Using this random noise model P_{rand} , we performed several runs of the SPSA to record its iteration complexity. We observe that the mean-valued run converges after approximately 60 iterations, see Fig. 8.




Figure 5: The (noisy) objective values of SPSA for the random noise model P_{rand} for a 4-regular graph. Although this noise model is not physical, we see that convergence is already achieved after approximately 50 iterations.

A realistic noise model

Next, we consider a realistic biased noise model given by the matrix P :

$$P = \begin{pmatrix} x & \frac{1-x}{4} & \frac{1-x}{4} & \frac{1-x}{4} & \frac{1-x}{4} & & & & & & \\ & y & & & \frac{1-y}{4} & & \frac{1-y}{4} & & \frac{1-y}{4} & & \\ & & y & & \frac{1-y}{4} & & \frac{1-y}{4} & & \frac{1-y}{4} & & \\ & & & y & & \frac{1-y}{4} & \frac{1-y}{4} & \frac{1-y}{4} & & & \\ & & & & y & \frac{1-y}{4} & \frac{1-y}{4} & \frac{1-y}{4} & & & \\ & \alpha \frac{1-x}{2} & & & \alpha \frac{1-x}{2} & y & & & & & \\ & & \alpha \frac{1-x}{2} & & \alpha \frac{1-x}{2} & & y & & \alpha \frac{1-y}{2} & \frac{1-y}{2} & \frac{1-y}{2} \\ & & & \alpha \frac{1-x}{2} & \alpha \frac{1-x}{2} & & & y & \alpha \frac{1-y}{2} & \alpha \frac{1-y}{2} & \\ & \alpha \frac{1-x}{2} & & \alpha \frac{1-x}{2} & & & & & \alpha \frac{1-y}{2} & \alpha \frac{1-y}{2} & \\ & & \alpha \frac{1-x}{2} & \alpha \frac{1-x}{2} & & & & y & & & \\ & \alpha \frac{1-x}{2} & \alpha \frac{1-x}{2} & & & & & & \alpha \frac{1-y}{2} & & \alpha \frac{1-y}{2} \\ & & & & & \alpha \frac{1-x}{3} & \alpha \frac{1-x}{3} & \alpha \frac{1-x}{3} & & y & \alpha(1-y) \\ & & & & \alpha \frac{1-x}{3} & \alpha \frac{1-x}{3} & \alpha \frac{1-x}{3} & & & y & \alpha(1-y) \\ & & & \alpha \frac{1-x}{3} & \alpha \frac{1-x}{3} & & & \alpha \frac{1-x}{3} & & & \alpha(1-y) \\ & & & & & \alpha \frac{1-x}{3} & \alpha \frac{1-x}{3} & \alpha \frac{1-x}{3} & & y & \alpha(1-y) \\ & & & & & & \beta \frac{1-x}{4} & \beta \frac{1-x}{4} & \beta \frac{1-x}{4} & \beta \frac{1-x}{4} & y \end{pmatrix}$$

In this model, $x = P(1|0)$, $y = P(0|1)$, and α, β are normalization factors such that the sum of each row is one, preserving the probability. Our model only takes into account single bit flips in the classical measurement. For example, it assigns probabilities as described for states $|\alpha_1\alpha_2\alpha_3\alpha_4\rangle$ where any (but only one) of the α_i flips its bit, but assigns zero probability for two or more bit flips.

Using this model and taking into account the conditional probabilities of Table 3, we performed several runs of the SPSA using the probability matrix induced by 1-qubit probability matrix of three IBMQ devices, **Armonk**, **Lima**, and **Bogota**. In Figure 8, we present the multiple noisy trajectories of the objective function values for five runs, together with the mean (golden crosses) and one standard deviation added to and subtracted from the mean (shaded region) for the five runs.

In this setting, consider four different manifestations of the bias:

- bias in the read-out model. In particular, a fixed difference $\tilde{b}(c) := |c - x|$ input into our choice P of the model of the read-out error of a 4-qubit device, which defines dividing the magnitude of the SPSA perturbation parameter c across the off-diagonal elements of a row of P above.
- the empirical distribution of objective function values obtained by running the circuit of Fig. 4 for QAOA on the 4-vertex MAXCUT instance described above with a particular read-out model with a particular choice of $\tilde{b}$. In Fig. 7, we display the empirical distributions obtained by running 1 iteration of QAOA for four different choices of $\tilde{b}$ in the read-out model: 0.01, 0.02, 0.03, 0.04, with 1024 runs per each value of $\tilde{b}$.
- an estimate $\hat{b}$ of the bias in the empirical distribution of the objective function values, as obtained by performing parameter estimation of a mixture model (24) composed of a normal and a uniform distribution on a set of samples of the QAOA run for




Figure 6: The iteration complexity of SPSA for different IBMQ devices.




Figure 7: Four empirical distributions of the objective function values for four values of $\tilde{b}$. Clock-wise from the top-left: $\tilde{b}$: 0.01, 0.02, 0.03, 0.04, respectively, that yield $\hat{b} = 0.31224, 0.32779, 0.33484, 0.36894$. The modes, indicated by the vertical black lines, are located at $-2.224, -2.221, -2.2635, -2.219$. The colours match those of Fig. 8, except for 30% transparency. For comparison, see Fig. 9, which suggests that the mean of the normally-distributed component of the mixture is -1.95 .

a representative value of θ . In Fig. 7, clock-wise from the top-left, the estimates $\hat{b}$ are 0.31224, 0.32779, 0.33484, 0.36894 based on $\tilde{b}(c)$ of 0.01, 0.02, 0.03, and 0.04, respectively. For comparison, see Fig. 9 for the case of no read-out error, but realistic fidelity of the gates.

- the impact of the bias on the rate of convergence. In Fig. 8, for the same values of $c, \tilde{b}$ and estimated bias $\hat{b}$ as in Fig. 7, we illustrate the trajectory of objective values by SPSA iteration while considering the variable noise model given by P .

The empirical results in Fig. 8 seem to corroborate the analytical results of Thm. 3.2. As expected from Thm. 3.2 and Fig. 3, the iteration complexity grows with the bias.

4.1 Comparison Between SPSA, 2PFA, and PSR

Now we shall compare SPSA, the two point function approximation and the PSR as zero order gradient approximation schemes. For the experiment we chose 10 random starting points, and then started 20 trials for all three algorithms across parameter values $c = 1e-1, 1e-2, 1e-3$ and the stepsize $\alpha = 1e-1, 1e-2, 1e-3$. We are interested in studying the average final output value, the number of iterations until reaching convergence to the stationary point or noise dominated region, and the variability of the runs with respect to these criteria. To perform this analysis, we consider four different regimes with respect to small or large c and small or large α . Below we show figures that are representative of the results obtained. Note that the perturbation is not a parameter in PSR. We used the Lima noise model for all experiments.




Figure 8: The iteration complexity of SPSA for four different values of bias $\hat{b}$ with error $c = 0.01$.




Figure 9: The empirical distribution of objective function values from 4096 evaluations in the ideal read-out scenario, i.e., where $P = \text{diag}(\mathbf{1}_{16})$. Our estimate of the bias is 2.24×10^{-14} with an empirical observation for the mean of the normal distribution at -1.95 and the mode of the mixture at -1.9737 .

4.1.1 Large Stepsize, Large Perturbation

With this regime, we would expect fast convergence, however, to a fairly high variance noise-dominated region. In addition, monotonic decrease should be consistently seen, but perhaps not as steep as otherwise, as we expect high perturbation gradient estimates to be robust but imprecise.

In Figure 10 we observe:

- Impressively, even with an imprecise estimate with a high value of c , the speed of convergence of SPSA is the same as what can be considered the exact gradient of




Figure 10: Results for SPSA and 2PFA for a representative run with $c = 0.1$ and $\alpha = 0.1$ and for PSR with $\alpha = 0.1$.




Figure 11: Results for SPSA and 2PFA for a representative run with $c = 0.01, 0.001$, and $\alpha = 0.1$.

PSR. However, as expected the convergence is to a relatively large noise dominated region with a high average bias in the former case. By comparison 2PFA is both slower to converge in terms of iterations, while also being significantly less robust, with many late iterations even surpassing the objective value of the starting point.

4.1.2 Large Stepsize, Smaller Perturbation

This regime indicates aggressive optimization towards a minimum, together with a precise but noisier estimate of the gradient, with the intention of finding a stationary point with lower bias. The plots in Figure 11 are interesting and communicate some important subtleties in regards to the two zero order methods. We observe that SPSA is very noisy for even moderately small c when the stepsize is large. However, 2PFA exhibits a steady and reliable, if slow, convergence with the moderately small perturbation to a fairly good objective value. Meanwhile it is completely useless with a higher variance estimate using $c = 0.001$.

4.1.3 Small Stepsize, Large Perturbation

This regime is the “safest” in terms of obtaining a decent objective value reliably. A small stepsize means a smaller noise dominated region near optimality and more locally steps. A large perturbation results in a relatively low variance and but imprecise and high bias gradient. From Figure 12 we observe, remarkably, that SPSA converges faster than even the PSR in the safe regime. Otherwise, all three algorithms exhibit steady decrease with the iterations, with 2PFA being the noisiest of the algorithms.

4.1.4 Medium Stepsize, Small Perturbation

This regime indicates a steady decrease, but with a precise gradient estimate, thus obtaining a final objective value lower in mean and higher in variance. Finally, Figure 13 indicates




Figure 12: Results for SPSA and 2PFA for a representative run with $c = 0.1$ and $\alpha = 0.001$ and PSR with $\alpha = 0.001$.




Figure 13: Results for SPSA and 2PFA for a representative run with $c = 0.01, 0.001$ and $\alpha = 0.1$.

that with a small perturbation, the variance of the two point function approximation estimates practically eliminates any potential advantage of the steady reliable decrease of a smaller stepsize. For a more moderate perturbation, the convergence is both slow and noisy, albeit at least relatively steady. By contrast SPSA clearly has an optimal threshold, but with a small value of c , some iterates still obtain a good quality objective value. The top Figure indicates particularly strong performance in speed and reliability, together with the lowest bias seen among the comparisons.

5 A Discussion and Conclusion

We have analyzed the convergence guarantees of variational quantum algorithms in the presence of bias in the quantum part. Let us interpret these results.

In regards to the asymptotic bias in the approximate measure of stationarity, we can notice that, predictably, the greater the bias in the function evaluation, the greater the error lower bound in the ultimate optimality criterion. A coarser discretization for estimating the gradient appears to stabilize the effect of the function evaluation bias but carries the trade-off in introducing additional asymptotic optimality bias in the overall rougher accuracy in gradient estimation. In addition, as typical for zero-order methods, the asymptotic bias scales poorly as a function of the problem dimension.

With respect to the iteration complexity, we can see that the overall rate is not affected by the presence of bias. The rate of convergence is the same as in a standard stochastic gradient analysis. That said, for the two terms in the convergence guarantee that do drop to zero with more iterations, exhibit a larger constant with greater bias.

More broadly, this is a small step towards a better understanding of the performance of variational quantum algorithms. Until recently, these have been seen as heuristics with no guarantees whatsoever. This and recent related work [5, e.g.] suggests that their performance can be characterized in some detail.

Acknowledgement This work is a collaboration between the Fidelity Center for Applied Technology, Fidelity Labs, LLC., and the Czech Technical University in Prague.

References

- [1] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. “The theory of variational hybrid quantum-classical algorithms”. *New Journal of Physics* **18**, 023023 (2016).
- [2] Xiao Yuan, Suguru Endo, Qi Zhao, Ying Li, and Simon C Benjamin. “Theory of variational quantum simulation”. *Quantum* **3**, 191 (2019).
- [3] Marco Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, et al. “Variational quantum algorithms”. *Nature Reviews Physics* **3**, 625–644 (2021).
- [4] Amira Abbas, Andris Ambainis, Brandon Augustino, Andreas Bärttschi, Harry Buhrman, Carleton Coffrin, Giorgio Cortiana, Vedran Dunjko, Daniel J. Egger, Bruce G. Elmegreen, Nicola Franco, Filippo Fratini, Bryce Fuller, Julien Gacon, Constantin Gonceiulea, Sander Gribling, Swati Gupta, Stuart Hadfield, Raoul Heese, Gerhard Kircher, Thomas Kleinert, Thorsten Koch, Georgios Korpas, Steve Lenk, Jakub Marecek, Vanio Markov, Guglielmo Mazzola, Stefano Mensa, Naeimeh Mohseni, Giacomo Nannicini, Corey O’Meara, Elena Peña Tapia, Sebastian Pokutta, Manuel

- Proissl, Patrick Rebentrost, Emre Sahin, Benjamin C. B. Symons, Sabine Törnøw, Victor Valls, Stefan Woerner, Mira L. Wolf-Bauwens, Jon Yard, Sheir Yarkoni, Dirk Zechiel, Sergiy Zhuk, and Christa Zoufal. “Quantum Optimization: Potential, Challenges, and the Path Forward” (2023). [arXiv:2312.02279](#).
- [5] Daniel J. Egger, Jakub Mareček, and Stefan Woerner. “Warm-starting quantum optimization”. [Quantum](#) **5**, 479 (2021).
 - [6] Bela Bauer, Sergey Bravyi, Mario Motta, and Garnet Kin-Lic Chan. “Quantum algorithms for quantum chemistry and quantum materials science”. [Chemical Reviews](#) **120**, 12685–12717 (2020).
 - [7] Lin Lin and Yu Tong. “Heisenberg-limited ground-state energy estimation for early fault-tolerant quantum computers”. [PRX Quantum](#) **3**, 010318 (2022).
 - [8] Travis S Humble, Andrea Delgado, Raphael Pooser, Christopher Seck, Ryan Benink, Vicente Leyton-Ortega, C-C Joseph Wang, Eugene Dumitrescu, Titus Morris, Kathleen Hamilton, et al. “Snowmass white paper: Quantum computing systems and software for high-energy physics research” (2022).
 - [9] Andrew Blance and Michael Spannowsky. “Quantum machine learning for particle physics using a variational quantum classifier”. [Journal of High Energy Physics](#) **2021**, 1–20 (2021).
 - [10] Aram W. Harrow and John C. Napp. “Low-depth gradient measurements can improve convergence in variational hybrid quantum-classical algorithms”. [Physical Review Letters](#) **126**, 140502 (2021).
 - [11] Artur K Ekert, Carolina Moura Alves, Daniel KL Oi, Michał Horodecki, Paweł Horodecki, and Leong Chuan Kwek. “Direct estimations of linear and nonlinear functionals of a quantum state”. [Physical Review Letters](#) **88**, 217901 (2002).
 - [12] Laurin E Fischer, Daniel Miller, Francesco Tacchino, Panagiotis Kl Barkoutsos, Daniel J Egger, and Ivano Tavernelli. “Ancilla-free implementation of generalized measurements for qubits embedded in a qudit space”. [Physical Review Research](#) **4**, 033027 (2022).
 - [13] James C Spall et al. “Multivariate stochastic approximation using a simultaneous perturbation gradient approximation”. [IEEE transactions on automatic control](#) **37**, 332–341 (1992).
 - [14] Xavier Bonet-Monroig, Hao Wang, Diederick Vermetten, Bruno Senjean, Charles Moussa, Thomas Bäck, Vedran Dunjko, and Thomas E O’Brien. “Performance comparison of optimization methods on variational quantum algorithms”. [Physical Review A](#) **107**, 032407 (2023).
 - [15] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii. “Quantum circuit learning”. [Physical Review A](#) **98** (2018).
 - [16] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. “Evaluating analytic gradients on quantum hardware”. [Physical Review A](#) **99**, 032331 (2019).
 - [17] Leonardo Banchi and Gavin E. Crooks. “Measuring analytic gradients of general quantum evolution with the stochastic parameter shift rule”. [Quantum](#) **5**, 386 (2021).
 - [18] Giuliano Benenti, Giulio Casati, Davide Rossini, and Giuliano Strini. “Principles of quantum computation and information: a comprehensive textbook”. [World Scientific](#). (2019).

- [19] Vivek S Borkar. “Stochastic approximation: a dynamical systems viewpoint”. *Volume 48*. Springer. (2009).
- [20] A Doucet and V Tadic. “Asymptotic bias of stochastic gradient search”. *Annals of Applied Probability* **27**, 3255–3304 (2017).
- [21] Belhal Karimi, Blazej Miasojedow, Eric Moulines, and Hoi-To Wai. “Non-asymptotic analysis of biased stochastic approximation scheme”. In Alina Beygelzimer and Daniel Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*. Volume 99 of *Proceedings of Machine Learning Research*, pages 1944–1974. PMLR (2019). url: <https://proceedings.mlr.press/v99/karimi19a.html>.
- [22] Ahmad Ajalloeian and Sebastian U. Stich. “Analysis of SGD with biased gradient estimators” (2020). [arXiv:2008.00051](https://arxiv.org/abs/2008.00051).
- [23] Ruobing Chen. “Stochastic derivative-free optimization of noisy functions”. PhD thesis. Lehigh University. (2015). url: <https://preserve.lehigh.edu/lehigh-scholarship/graduate-publications-theses-dissertations/theses-dissertations/stochastic-7>.
- [24] Oleg Granichin and Natalia Amelina. “Simultaneous perturbation stochastic approximation for tracking under unknown but bounded disturbances”. *IEEE Transactions on Automatic Control* **60**, 1653–1658 (2014).
- [25] Andrei Boiarov, Oleg Granichin, and Hou Wenguang. “Simultaneous perturbation stochastic approximation for clustering of a gaussian mixture model under unknown but bounded disturbances”. In *2017 IEEE Conference on Control Technology and Applications (CCTA)*. Pages 1740–1745. IEEE (2017).
- [26] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O’Brien. “A variational eigenvalue solver on a photonic quantum processor”. *Nature Communications* **5** (2014).
- [27] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. “A quantum approximate optimization algorithm” (2014). [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
- [28] Dave Wecker, Matthew B. Hastings, and Matthias Troyer. “Progress towards practical quantum variational algorithms”. *Physical Review A* **92** (2015).
- [29] John Preskill. “Quantum Computing in the NISQ era and beyond”. *Quantum* **2**, 79 (2018).
- [30] Oleksandr Kyriienko, Annie E. Paine, and Vincent E. Elfving. “Solving nonlinear differential equations with differentiable quantum circuits”. *Physical Review A* **103**, 052416 (2021).
- [31] Razin A. Shaikh, Sara Sabrina Zemljic, Sean Tull, and Stephen Clark. “The Conceptual VAE” (2022). [arXiv:2203.11216](https://arxiv.org/abs/2203.11216).
- [32] Stuart Hadfield, Zhihui Wang, Bryan O’Gorman, Eleanor Rieffel, Davide Venturelli, and Rupak Biswas. “From the quantum approximate optimization algorithm to a quantum alternating operator ansatz”. *Algorithms* **12**, 34 (2019).
- [33] Guillaume Verdon, Jacob Marks, Sasha Nanda, Stefan Leichenauer, and Jack Hidary. “Quantum Hamiltonian-Based Models and the Variational Quantum Thermalizer Algorithm” (2019). [arXiv:1910.02071](https://arxiv.org/abs/1910.02071).
- [34] Benjamin Nachman, Miroslav Urbanek, Wibe A. de Jong, and Christian W. Bauer. “Unfolding quantum computer readout noise”. *npj Quantum Information* **6** (2020).

- [35] Gavin E. Crooks. “Gradients of parameterized quantum gates using the parameter-shift rule and gate decomposition” (2019). [arXiv:1905.13311](#).
- [36] Jun Li, Xiaodong Yang, Xinhua Peng, and Chang-Pu Sun. “Hybrid quantum-classical approach to quantum optimal control”. *Physical Review Letters* **118** (2017).
- [37] Andrea Mari, Thomas R. Bromley, and Nathan Killoran. “Estimating the gradient and higher-order derivatives on quantum hardware”. *Physical Review A* **103** (2021).
- [38] C. C. Y. Dorea. “Expected number of steps of a random optimization method”. *Journal of Optimization Theory and Applications* **39**, 165–171 (1983).
- [39] R. Wengert. “A simple automatic derivative evaluation program”. *Communications of the ACM* **7**, 463–464 (1964).
- [40] Yurii Nesterov. “Smooth minimization of non-smooth functions”. *Mathematical programming* **103**, 127–152 (2005).
- [41] John C Duchi, Michael I Jordan, Martin J Wainwright, and Andre Wibisono. “Optimal rates for zero-order convex optimization: The power of two function evaluations”. *IEEE Transactions on Information Theory* **61**, 2788–2806 (2015).
- [42] Yurii Nesterov and Vladimir Spokoiny. “Random gradient-free minimization of convex functions”. *Foundations of Computational Mathematics* **17**, 527–566 (2015).
- [43] Navin Khaneja, Timo Reiss, Cindie Kehlet, Thomas Schulte-Herbrüggen, and Steffen J. Glaser. “Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms”. *Journal of Magnetic Resonance* **172**, 296–305 (2005).
- [44] Javier Gil Vidal and Dirk Oliver Theis. “Calculus on parameterized quantum circuits” (2018). [arXiv:1812.06323](#).
- [45] Artur F Izmaylov, Robert A Lang, and Tzu-Ching Yen. “Analytic gradients in variational quantum algorithms: Algebraic extensions of the parameter-shift rule to general unitary transformations”. *Physical Review A* **104**, 062443 (2021).
- [46] David Wierichs, Josh Izaac, Cody Wang, and Cedric Yen-Yu Lin. “General parameter-shift rules for quantum gradients”. *Quantum* **6**, 677 (2022).
- [47] Dirk Oliver Theis. “Optimality of finite-support parameter shift rules for derivatives of variational quantum circuits” (2021).
- [48] Jakob S. Kottmann, Abhinav Anand, and Alán Aspuru-Guzik. “A feasible approach for automatically differentiable unitary coupled-cluster on quantum computers”. *Chemical Science* **12**, 3497–3508 (2021).
- [49] Bálint Koczor and Simon C Benjamin. “Quantum analytic descent”. *Physical Review Research* **4**, 023017 (2022).
- [50] Andrew R Conn, Katya Scheinberg, and Luis N Vicente. “Introduction to derivative-free optimization”. *SIAM*. (2009).
- [51] Albert S Berahas, Liyuan Cao, Krzysztof Choromanski, and Katya Scheinberg. “A theoretical and empirical comparison of gradient approximations in derivative-free optimization”. *Foundations of Computational Mathematics* **22**, 507–560 (2022).
- [52] Daniel C Chin. “Comparative study of stochastic algorithms for system optimization based on gradient approximations”. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* **27**, 244–249 (1997).

- [53] Saeed Ghadimi and Guanghai Lan. “Stochastic first-and zeroth-order methods for nonconvex stochastic programming”. *SIAM Journal on Optimization* **23**, 2341–2368 (2013).
- [54] Krishnakumar Balasubramanian and Saeed Ghadimi. “Zeroth-order (non)-convex stochastic optimization via conditional gradient and gradient updates”. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*. Volume 31. Curran Associates, Inc. (2018).
- [55] Md. Sajid Anis et al. “Qiskit: An open-source framework for quantum computing” (2021).
- [56] Isaac L. Chuang Michael A. Nielsen. “Quantum computation and quantum information”. *Cambridge University Press*. (2011). 10th anniversary edition edition.
- [57] Giuliano Gadioli La Guardia. “Quantum error correction”. *Springer International Publishing*. (2020).
- [58] Swamit S Tannu and Moinuddin K Qureshi. “Mitigating measurement errors in quantum computers by exploiting state-dependent bias”. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*. Pages 279–290. (2019).
- [59] M. D. Shulman, O. E. Dial, S. P. Harvey, H. Bluhm, V. Umansky, and A. Yacoby. “Demonstration of entanglement of electrostatically coupled singlet-triplet qubits”. *Science* **336**, 202–205 (2012).
- [60] T. F. Watson, S. G. J. Philips, E. Kawakami, D. R. Ward, P. Scarlino, M. Veldhorst, D. E. Savage, M. G. Lagally, Mark Friesen, S. N. Coppersmith, M. A. Eriksson, and L. M. K. Vandersypen. “A programmable two-qubit quantum processor in silicon”. *Nature* **555**, 633–637 (2018).
- [61] Ioan M. Pop, Kurtis Geerlings, Gianluigi Catelani, Robert J. Schoelkopf, Leonid I. Glazman, and Michel H. Devoret. “Coherent suppression of electromagnetic dissipation due to superconducting quasiparticles”. *Nature* **508**, 369–372 (2014).

A Appendix: Proofs of Theoretical Results

To begin with, we prove a bound on the error associated with the two point function approximation technique:

Lemma A.1 *Under Assumption 3.1, the two point estimate (19) satisfies, for all θ*

$$\mathbf{g}(\theta, \xi^1, \xi^2) = \nabla f_c(\theta) + \mathbf{r}(\xi^1, \xi^2) + \mathbf{b}_c(\theta), \mathbb{E}[\mathbf{r}(\xi^1, \xi^2)] = \mathbf{0}, \|\mathbf{b}_c(\theta)\| \leq \frac{2b}{c}$$

Proof.

$$\begin{aligned} \mathbb{E}[\mathbf{g}] &= \mathbb{E}_{\Delta} \left[\frac{\mathbb{E}_{\xi^1}[F(\theta + c\Delta, \xi^1)] - \mathbb{E}_{\xi^2}[F(\theta, \xi^2)]}{c} \Delta \right] \\ &= \mathbb{E}_{\Delta} \left[\frac{f(\theta + c\Delta) + \mathbb{E}_{\xi^1}[n(\theta + c\Delta, \xi^1)] - f(\theta) - \mathbb{E}_{\xi^2}[n(\theta, \xi^2)]}{c} \Delta \right] \\ &= \nabla f_c(\theta) + \mathbb{E}_{\Delta} \left[\frac{\mathbb{E}_{\xi^1}[b(\theta + c\Delta, \xi^1)] - \mathbb{E}_{\xi^1}[b(\theta, \xi^2)]}{c} \Delta \right] \end{aligned}$$

where we have used the definition of ∇f_c as given in (20) and the functional form of the bias given in (24).

Now we can bound the second term,

$$\left\| \mathbb{E}_{\Delta} \left[\frac{\mathbb{E}_{\xi^1}[b(\theta + c\Delta, \xi^1)] - \mathbb{E}_{\xi^1}[b(\theta, \xi^2)]}{c} \Delta \right] \right\| \leq \frac{2}{(2\pi)^{\frac{n}{2}}} \int_0^\infty \frac{\|b(\cdot)\|_\infty + \|b(\cdot)\|_\infty}{c} |\Delta| d\Delta \leq 2b/c$$

where we used the symmetry of the Gaussian. ■

Now we can show the main result, which we restate for reference. The proof follows standard arguments, e.g., [53, Theorem 3.2].

Theorem A.1 *Consider the two point function approximation algorithm with the iteration (4) for a problem satisfying Assumptions 3.1 and 3.2.*

- (1) *Consider a fixed budget of K iterations. Let $\alpha_k \equiv \alpha/\sqrt{K}$ with $\alpha \leq L/6$. It holds that the iterates satisfy:*

$$\min_{k \in [K]} \mathbb{E} \|\nabla f(\theta_k)\|^2 \leq 4 \frac{f_c(\theta_k) - f_c^*}{\alpha \sqrt{K}} + 2 \frac{L\alpha(\sigma^2 + \frac{8b^2}{c^2})}{\alpha \sqrt{K}} + \left[\frac{2b^2}{c^2} + \frac{c^2}{2} L(p+3)^3 \right] \quad (37)$$

- (2) *Consider a diminishing stepsize of $\alpha_k = \alpha/k^\gamma$ with $\gamma \in (1/2, 1]$ and $\alpha \leq L/2$. Then it holds that, taking,*

$$\theta_R \sim \theta_k \text{ w.p. } \alpha_k / \sum_{l=1}^K \alpha_l$$

we have,

$$\mathbb{E} \|\nabla f(\theta_R)\|^2 \leq 4 \frac{f_c(\theta_k) - f_c^*}{\alpha K^\gamma} + 2 \frac{L\alpha(\sigma^2 + \frac{8b^2}{c^2})}{\alpha K^\gamma} + \left[\frac{2b^2}{c^2} + \frac{c^2}{2} L(p+3)^3 \right] \quad (38)$$

Proof. We begin with the standard Descent Lemma on $f_c(\theta)$, recalling that f_c has a Lipschitz gradient with the Lipschitz constant bounded by L and proceed with expressing the step as the composition given in Lemma A.1.

$$\begin{aligned} f_c(\theta_{k+1}) &\leq f_c(\theta_k) - \alpha_k \langle \nabla f_c(\theta_k), \mathbf{g}_k \rangle + \frac{L}{2} \alpha_k^2 \|\mathbf{g}_k\|^2 \\ &= f_c(\theta_k) - \alpha_k \|\nabla f_c(\theta_k)\|^2 - \alpha_k \langle \nabla f_c(\theta_k), \mathbf{r}(\xi_k^1, \xi_k^2) + \mathbf{b}_c(\theta_k) \rangle + \frac{L}{2} \alpha_k^2 \|\mathbf{g}_k\|^2 \end{aligned}$$

Taking conditional expectations with respect to the Filtration $\mathcal{F}_k$ corresponding to (ξ_k^1, ξ_k^2) , and summing up, we get, with f_c^* a lower bound for $f_c(\boldsymbol{\theta})$,

$$\sum_{k=1}^K \alpha_k \mathbb{E} \left[\|\nabla f_c(\boldsymbol{\theta}_k)\|^2 | \mathcal{F}_k \right] \leq f_c(\boldsymbol{\theta}_k) - f_c^* - \sum_{k=1}^K \mathbb{E} \left[\alpha_k \langle \nabla f_c(\boldsymbol{\theta}_k), \mathbf{r}(\xi_k^1, \xi_k^2) + \mathbf{b}_c(\boldsymbol{\theta}_k) \rangle | \mathcal{F}_k \right] + \frac{L}{2} \sum_{k=1}^K \alpha_k^2 \mathbb{E} \left[\|\mathbf{g}_k\|^2 | \mathcal{F}_k \right] \quad (39)$$

Now, we have, from Lemma A.1 and Assumption 3.1, and noting that cross terms involving $r(\xi_k^1, \xi_k^1)$ and deterministic quantities are zero by $\mathbb{E}[r(\xi_k^1, \xi_k^2) | \mathcal{F}_k] = 0$,

$$\mathbb{E}[\|\mathbf{g}_k\|^2 | \mathcal{F}_k] \leq 2\|\nabla f_c(\boldsymbol{\theta}_k)\|^2 + \mathbb{E}[\|r(\xi_k^1, \xi_k^2)\|^2 | \mathcal{F}_k] + 2\|\mathbf{b}_c(\boldsymbol{\theta}_k)\|^2 \leq 2\|\nabla f_c(\boldsymbol{\theta}_k)\|^2 + \sigma^2 + \frac{8b^2}{c^2}$$

Furthermore, using that $\mathbb{E}[r(\xi_k^1, \xi_k^2) | \mathcal{F}_k] = 0$ and Young's inequality,

$$-\mathbb{E}[\langle \nabla f_c(\boldsymbol{\theta}_k), r(\xi_k^1, \xi_k^2) + \mathbf{b}_c(\boldsymbol{\theta}_k) \rangle | \mathcal{F}_k] = -\langle \nabla f_c(\boldsymbol{\theta}_k), \mathbf{b}_c(\boldsymbol{\theta}_k) \rangle \leq \frac{\|\nabla f_c(\boldsymbol{\theta}_k)\|^2}{2} + \frac{\|\mathbf{b}_c\|^2}{2} \leq \frac{\|\nabla f_c(\boldsymbol{\theta}_k)\|^2}{2} + \frac{2b^2}{c^2}$$

Plugging in the previous two bounds into (39) and then taking total expectations and using the expectation tower property,

$$\sum_{k=1}^K \frac{\alpha_k}{2} (1 - 2\alpha_k L) \mathbb{E} \|\nabla f_c(\boldsymbol{\theta}_k)\|^2 \leq f_c(\boldsymbol{\theta}_k) - f_c^* + \frac{L(\sigma^2 + \frac{8b^2}{c^2})}{2} \sum_{k=1}^K \alpha_k^2 + \frac{2b^2}{c^2} \sum_{k=1}^K \alpha_k$$

Note that we have, from [53, (3.9)],

$$\|\nabla f(\boldsymbol{\theta})\|^2 - \frac{c^2}{2} L(p+3)^3 \leq 2\|\nabla f_c(\boldsymbol{\theta})\|^2$$

and so finally,

$$\sum_{k=1}^K \frac{\alpha_k}{2} (1 - 2\alpha_k L) \mathbb{E} \|\nabla f(\boldsymbol{\theta}_k)\|^2 \leq f_c(\boldsymbol{\theta}_k) - f_c^* + \frac{L(\sigma^2 + \frac{8b^2}{c^2})}{2} \sum_{k=1}^K \alpha_k^2 + \left[\frac{2b^2}{c^2} + \frac{c^2}{2} L(p+3)^3 \right] \sum_{k=1}^K \alpha_k \quad (40)$$

We can now consider budget and diminishing step-size regimes. In the first case, we have a fixed budget of K iterations and take $\alpha_k := \frac{\alpha}{\sqrt{K}}$ with $\alpha \leq L/4$. The above expression becomes

$$\frac{1}{4} \sqrt{K} \alpha \min_{k \in [K]} \mathbb{E} \|\nabla f(\boldsymbol{\theta}_k)\|^2 \leq f_c(\boldsymbol{\theta}_k) - f_c^* + \frac{L\alpha(\sigma^2 + \frac{8b^2}{c^2})}{2} + \alpha \left[\frac{2b^2}{c^2} + \frac{c^2}{2} L(p+3)^3 \right] \sqrt{K} \quad (41)$$

from which we derive the first result. ■

Now we first re-state then prove Theorem 3.2, the main convergence result for SPSA under biased noisy function evaluations.

Theorem A.2 *Consider the SPSA algorithm with the iteration (4) and gradient estimate (18) for a problem satisfying Assumptions 3.1 and 3.2. In addition, let the SPSA algorithm satisfy the conditions of Lemma 3.1 together with there existing $B_3 > 0$ such that $[\Delta_k]_i^{-1} \leq B_3$ almost surely for all k . Then it holds that,*

- (1) *Consider a fixed budget of K iterations. Let $\alpha_k \equiv \alpha/\sqrt{K}$ with $\alpha \leq L/2$. It holds that the iterates satisfy:*

$$\min_{k \in [K]} \mathbb{E} \|\nabla f(\boldsymbol{\theta}_k)\|^2 \leq \frac{4(f_c(\boldsymbol{\theta}_k) - f_c^*)}{\alpha \sqrt{K}} + \frac{2L\alpha \left[\frac{\sigma^2 B_3^2}{c_k^2} + 2b_0 c_k^2 + \frac{2b^2 B_3^2}{c_k^2} \right]}{\sqrt{K}} + \left[b_0 c_k^2 + \frac{b^2 B_3^2}{c_k^2} \right] \quad (42)$$

(2) Consider a diminishing stepsize of $\alpha_k = \alpha/k^\gamma$ with $\gamma \in (1/2, 1]$ and $\alpha \leq L/2$. Then it holds that, taking,

$$\boldsymbol{\theta}_R \sim \boldsymbol{\theta}_k \text{ w.p. } \alpha_k / \sum_{l=1}^K \alpha_l$$

we have,

$$\mathbb{E} \|\nabla f(\boldsymbol{\theta}_R)\|^2 \leq 4 \frac{f_c(\boldsymbol{\theta}_k) - f_c^*}{\alpha K^\gamma} + 2 \frac{L\alpha \left[\frac{\sigma^2 B_3^2}{c_k^2} + 4b_0 c_k^2 + \frac{4b^2 B_3^2}{c_k^2} \right]}{K^\gamma} + \left[b_0 c_k^2 + \frac{b^2 B_3^2}{c_k^2} \right] \quad (43)$$

Proof. Consider $\hat{\mathbf{g}}_k$ to be the SPSA estimate with zero bias in function evaluations and $\mathbf{g}_k$ the actual estimate. Recalling Lemma 3.1, we have that

$$\mathbf{g}_k - \nabla f(\boldsymbol{\theta}_k) = \mathbf{g}_k - \hat{\mathbf{g}}_k + \hat{\mathbf{g}}_k - \nabla f(\boldsymbol{\theta}_k) = \mathbf{g}_k - \hat{\mathbf{g}}_k + \mathbf{r}(\xi_k^1, \xi_k^2, \boldsymbol{\Delta}_k) + \mathbf{b}_0(\boldsymbol{\theta}_k)$$

Now using the characterization of the function evaluation (24) as well as the definition of $\mathbf{g}_k$, we get,

$$[\mathbf{g}_k]_i - [\hat{\mathbf{g}}_k]_i = \frac{\hat{r}_k + \hat{b}_k}{2c_k[\boldsymbol{\Delta}_k]_i}$$

with $\mathbb{E}[\hat{r}_k] = 0$ and $\|\hat{b}_k\| \leq 2b$ a.s.. Now, recalling the assumptions on $\boldsymbol{\Delta}_k$ we have that,

$$|[\mathbf{g}_k]_i - [\hat{\mathbf{g}}_k]_i| \leq |\bar{r}_k| + |b_k^1|$$

with $\mathbb{E}|\bar{r}_k| \leq \frac{\sigma B_3}{c_k}$ and $|b_k^1| \leq bB_3/c_k$ almost surely. Putting these quantities together we get that,

$$\mathbf{g}_k - \nabla f(\boldsymbol{\theta}_k) = \mathbf{r}_s(\xi_k^1, \xi_k^2, \boldsymbol{\Delta}_k) + \mathbf{b}_s(\boldsymbol{\theta}_k)$$

with $\mathbb{E}\mathbf{r}_s = 0$ and $\|\mathbf{b}_s(\boldsymbol{\theta}_k)\|_\infty \leq b_0 c_k^2 + \frac{bB_3}{c_k}$.

Now we can proceed by similar lines as in the proof of Theorem 3.1, noting that in this case, although the gradient error estimate is more complex, the accuracy is with respect to the original function as opposed to a smoothed variant, thus permitting a more simplified exposition of the main convergence result.

Indeed, starting with the standard Descent Lemma on $f(\boldsymbol{\theta})$,

$$\begin{aligned} f(\boldsymbol{\theta}_{k+1}) &\leq f(\boldsymbol{\theta}_k) - \alpha_k \langle \nabla f(\boldsymbol{\theta}_k), \mathbf{g}_k \rangle + \frac{L}{2} \alpha_k^2 \|\mathbf{g}_k\|^2 \\ &= f(\boldsymbol{\theta}_k) - \alpha_k \|\nabla f(\boldsymbol{\theta}_k)\|^2 - \alpha_k \langle \nabla f(\boldsymbol{\theta}_k), \mathbf{r}_s(\xi_k^1, \xi_k^2) + \mathbf{b}_s(\boldsymbol{\theta}_k) \rangle + \frac{L}{2} \alpha_k^2 \|\mathbf{g}_k\|^2 \end{aligned}$$

Summing up, we get, with f^* a lower bound for $f(\boldsymbol{\theta})$,

$$\sum_{k=1}^K \alpha_k \mathbb{E} [\|\nabla f(\boldsymbol{\theta}_k)\|^2 | \mathcal{F}_k] \leq f(\boldsymbol{\theta}_k) - f^* - \sum_{k=1}^K \alpha_k \mathbb{E} [\langle \nabla f(\boldsymbol{\theta}_k), \mathbf{r}_s(\xi_k^1, \xi_k^2) + \mathbf{b}_s(\boldsymbol{\theta}_k) \rangle | \mathcal{F}_k] + \frac{L}{2} \sum_{k=1}^K \alpha_k^2 \mathbb{E} [\|\mathbf{g}_k\|^2 | \mathcal{F}_k] \quad (44)$$

Continuing, from Lemma 3.1 and Assumption 3.1,

$$\mathbb{E} [\|\mathbf{g}_k\|^2 | \mathcal{F}_k] \leq 2\|\nabla f(\boldsymbol{\theta}_0)\|^2 + \mathbb{E} [\|\mathbf{r}_s(\xi_k^1, \xi_k^2)\|^2 | \mathcal{F}_k] + 2\|\mathbf{b}_s(\boldsymbol{\theta}_k)\|^2 \leq 2\|\nabla f(\boldsymbol{\theta}_k)\|^2 + \frac{\sigma^2 B_3^2}{c_k^2} + 4b_0 c_k^2 + \frac{4b^2 B_3^2}{c_k^2}$$

and, using that $\mathbb{E}[\mathbf{r}_s(\xi_k^1, \xi_k^2) | \mathcal{F}_k] = 0$,

$$\begin{aligned} -\mathbb{E} [\langle \nabla f(\boldsymbol{\theta}_k), \mathbf{r}_s(\xi_k^1, \xi_k^2) + \mathbf{b}_s(\boldsymbol{\theta}_k) \rangle | \mathcal{F}_k] &= -\langle \nabla f(\boldsymbol{\theta}_k), \mathbf{b}_s(\boldsymbol{\theta}_k) \rangle \\ &\leq \frac{\|\nabla f(\boldsymbol{\theta}_k)\|^2}{2} + \frac{\|\mathbf{b}_s\|^2}{2} \leq \frac{\|\nabla f(\boldsymbol{\theta}_k)\|^2}{2} + b_0 c_k^2 + \frac{b^2 B_3^2}{c_k^2} \end{aligned}$$

Incorporating these two bounds into (44) and taking total expectations, we obtain,

$$\sum_{k=1}^K \frac{\alpha_k}{2} (1 - 2\alpha_k L) \mathbb{E} \|\nabla f(\boldsymbol{\theta}_k)\|^2 \leq f(\boldsymbol{\theta}_0) - f^* + \frac{L \left[\frac{\sigma^2 B_3^2}{c_k^2} + 4b_0 c_k^2 + \frac{4b^2 B_3^2}{c_k^2} \right]}{2} \sum_{k=1}^K \alpha_k^2 + \left[b_0 c_k^2 + \frac{b^2 B_3^2}{c_k^2} \right] \sum_{k=1}^K \alpha_k \quad (45)$$

We can now consider budget and diminishing step-size regimes. In the first case, we have a fixed budget of K iterations and take $\alpha_k := \frac{\alpha}{\sqrt{K}}$ with $\alpha \leq L/4$. The above expression becomes

$$\frac{1}{4} \sqrt{K} \alpha \min_{k \in [K]} \mathbb{E} \|\nabla f(\boldsymbol{\theta}_k)\|^2 \leq f_c(\boldsymbol{\theta}_k) - f_c^* + \frac{L \alpha^2 \left[\frac{\sigma^2 B_3^2}{c_k^2} + 2b_0 c_k^2 + \frac{2b^2 B_3^2}{c_k^2} \right]}{2} + \alpha \left[b_0 c_k^2 + \frac{b^2 B_3^2}{c_k^2} \right] \sqrt{K} \quad (46)$$

■

B Bias in VQA and Its Effects on Convergence Properties

Noise models in quantum circuits In the circuit model of quantum computation, an ideal quantum gate performs a unitary transformation on the quantum state, which is described by density matrices $\rho \in \mathbb{S}_+^d$. The unitary transformation acts on an input state ρ_{in} as $U \rho_{\text{in}} U^\dagger = \rho_{\text{out}}$. In practice, real quantum hardware is very noisy because of its unavoidable interaction with the environment. This corresponds to an *open quantum system* whose dynamics is never purely unitary.

In the most general setting, the evolution of states in open quantum systems can be described as $\rho_{\text{out}} = \mathcal{E}(\rho_{\text{in}})$ by a linear operator $\mathcal{E} : \mathbb{S}_+^d \rightarrow \mathbb{S}_+^d$ via the (nonunique) Kraus sum representation

$$\mathcal{E}(\rho) := \sum_{k=1}^{\ell} E_k \rho_{\text{in}} E_k^\dagger, \quad (47)$$

where $\{E_k\}_{k=1}^{\ell}$ are called Kraus operators and satisfy $\sum_{k=1}^{\ell} E_k^\dagger E_k = \mathbb{1}$, $\ell \leq d^2 - 1$.

Let us consider some concrete examples of how Kraus operators appear in the noise typical noise channels of the gate model of quantum computing with the standard Pauli gates $\{X, Y, Z\}$.

- (1) The *bit flip* with probability p is described as $\mathcal{E}(\rho) = (1 - p)\rho + pX\rho X^\dagger$ with the Kraus operators $\{\mathbb{1}, X\}$.
- (2) The *dephasing channel* is described as $\mathcal{E}(\rho) = (1 - p)\rho + pZ\rho Z^\dagger$ with the Kraus operators being $\{\mathbb{1}, Z\}$.
- (3) The *depolarizing channel* is described as $\mathcal{E}(\rho) = (1 - p)\rho + \frac{p}{4}(\rho + X\rho X^\dagger + Y\rho Y^\dagger + Z\rho Z^\dagger) = \frac{p}{2}\mathbb{1} + (1 - p)\rho$.

Bias in Quantum Circuits Errors are endemic in the measurements of quantum states after a circuit run. A browse through any contemporary textbook or broad survey (e.g., [56, 18, 57]) indicates the degree to which research is being undertaken in the design and implementation of error-correcting codes as well as in hardware design for minimizing noise, attributes that are absolutely essential for the commercial implementation of fault-tolerant error corrected quantum computers as well as the NISQ computers. In fact, it

is widely accepted that the latter can handle only shallow quantum circuits due to the presence of noise, which, with increasing depth, makes the output of any measurement progressively less reliable.

Of course, the presence of noise itself does not indicate bias, and even what would be considered a lot of noise could be unbiased but with a high variance. However, at first glance, it is clear that bias is inherently present in the computation of noisy quantum circuits. There are two sources of bias in the noise:

- (1) Bias due to certain noise channels that favor certain operation errors, e.g., Pauli errors, more frequently compared to other operations.
- (2) Bias in the measurement, wherein we obtain positive counts for states that should appear with probability 0.
- (3) Bias in the estimators used for reconstructing quantum states (quantum state tomography) or quantum processes (quantum process tomography).

As far as point (1) is concerned, consider, for example, a qubit system with Hamiltonian proportional to Z where a common noise model description for this system is that of the dephasing channel, which occurs at rates much greater than the rates of other noise channels. Specifically, consider the case where the energy eigenstate basis of the state of a qubit before the execution of the circuit matches the computational basis $\{|0\rangle, |1\rangle\}$. The noise induced by the dephasing channel is much higher, for example, than the noise induced by bit flips. For example, in [58] it was reported that simply measuring a zero state had 82% fidelity (and the state of all ones 62%) on a state-of-the-art quantum device. If the circuit and the observable operator are designed so that, in the theoretically ideal scenario, there is one eigenstate that is measured with probability one, then the expectation *should be* that eigenstate. However, suppose that one is measuring one state most of the time and another state some other percentage of the time, in practice, for this circuit, that is, an evaluation $f(\theta) \sim E_1$ with probability p_1 and $f(\theta) \sim E_2 > E_1$ with probability p_2 . In such a basic scenario, if E_1 is intended, then clearly all the noise is one-directional and thus biased. Furthermore, the noise is not even consistent since we expect repeated trials of this to continue to exhibit one-sided noise. As a matter of fact, in several types of physical qubit implementations certain types of errors over-dominate others, for example in solid state spin qubits [59, 60] and superconducting qubits [61]. This is a first indication that the implemented circuits are inherently biased.

However, in the context of the practical implementation of VQAs, point (3) is what really matters. Specifically, we are interested in how the noise in the actual measurements affects the values of the observables. As an example, consider the depolarizing channel,

$$\mathcal{E}(\rho) = p\mathbb{1} + (1 - p)\rho \quad (48)$$

which replaces each qubit (or the entire system) in the state ρ with the totally mixed state with some probability $p > 0$. Subsequently, consider that an observable O is applied to compute the energy, entropy, or entanglement, for instance. It is clear that if in all circuit runs the state ρ is the input of $\mathcal{E}$, any one evaluation will not have an expectation $\langle O \rangle := \text{Tr}(\rho O)$, rather $\langle \tilde{O} \rangle := \text{Tr}(\mathcal{E}(\rho) O)$, and no matter how many trials are performed, the sample average will not equal this desired value, i.e.,

$$\langle O \rangle \neq \mathbb{E}(\tilde{O}) := \frac{1}{N} \sum_{i=1}^N f_i, \quad f_i \sim \mathbb{E}_i(\tilde{O}), \quad (49)$$




Figure 14: Bias in the measurement of a Bell state: The top two plots present the measurements of the circuit of Fig. 15, in a noise-free simulation. The bottom two plots present the corresponding measurements from IBM’s 5-qubit device `ibmq_santiago`. The fact that states $|01\rangle$ and $|10\rangle$ appear with non-zero probability suggests there is bias, which is independent of the number of shots.

State $ \psi_f\rangle$	Analytical	Hardware
$ 00\rangle$	0.5	0.43
$ 01\rangle$	0	0.02
$ 10\rangle$	0	0.02
$ 11\rangle$	0.5	0.53

Table 4: A numerical comparison of the probabilities for the Bell state of Figure 14.

where f_i corresponds to a measurement. Said differently, the estimator $\tilde{O}$ of some observable O is biased in the sense $\mathbb{E}[O - \tilde{O}] \neq 0$. It is known that, for example, in IBM machines, there exists an asymmetric read-out error [34] since $p(|0\rangle | 1\rangle) > p(|1\rangle | 0\rangle)$. Thus, it is clear that the noise arising from a classical measurement applied to a quantum circuit, as is commonly done in, e.g., VQA, is not unbiased, as far as the expectation of the (parametrized) estimator is concerned, and thus any analysis of algorithms with function evaluations of this sort must incorporate this for faithful modeling of the problem.

Next, let us consider the bias in the measurement (2). To give an intentionally simple example, consider a Bell state in a 2-qubit register obtained performing a Hadamard gate followed by a CNOT gate. Figure 14 presents the corresponding measurements obtained in hardware (`ibmq_santiago`) and in simulation (using the `qiskit` Aer simulator). Table 4 presents the same data numerically.

Errors that result in inaccurate sample averages of particular states are, of course, well-known issues that are endemic to near-term devices, and a host of error mitigation




Figure 15: An intentionally simple quantum circuit preparing a Bell state.

strategies exist to reduce the magnitude of said errors.

However, regardless of the quantitative improvement in the frequency of errors, there would remain some nonzero appearance of the mixed states $|01\rangle$ and $|10\rangle$. This in itself makes it the case that there are observables for which, when they define the objective of a VQA circuit, there will always be bias, resulting in a fundamental discrepancy between the ideal distribution Ξ and the observed distribution of objective evaluations χ .

As a simple example, consider the operation of simply projecting onto the $|01\rangle$ state. This would correspond to taking the observable $O \in \text{Herm}(\mathbb{C}^4)$ as the Hermitian matrix,

$$O = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

Define $F(\boldsymbol{\theta}, \xi)$ as the output of $\langle \psi_0 | U^\dagger(\boldsymbol{\theta}) O U(\boldsymbol{\theta}) | \psi_0 \rangle = \langle \psi_f | O | \psi_f \rangle$. Note that there is no parameter $\boldsymbol{\theta}$ that enters the expression for $|\psi_0\rangle$ or the gates U or O , and so F is constant with respect to $\boldsymbol{\theta}$, and only the noise affects the outcome of its evaluation.

We have that $F(\cdot, \cdot)$ under the ideal distribution Ξ and the actual distribution χ are,

$$F(\boldsymbol{\theta}, \cdot) \sim_{\Xi} 0 \text{ w.p. } 1, \quad F(\boldsymbol{\theta}, \cdot) \sim_{\chi} \begin{cases} 0 & \text{w.p. } \approx 0.98, \\ 1 & \text{w.p. } \approx 0.02 \end{cases},$$

thus with $f(\boldsymbol{\theta}) \equiv 0$ as the desired expectation, we have,

$$F(\boldsymbol{\theta}, \xi) = f(\boldsymbol{\theta}) + n(\boldsymbol{\theta}, \xi) = f(\boldsymbol{\theta}) + r(\boldsymbol{\theta}, \xi) + b(\boldsymbol{\theta}) \quad (50)$$

with $b(\boldsymbol{\theta}) = (\mathbb{E}_{\chi} - \mathbb{E}_{\Xi})[F(\boldsymbol{\theta}, \cdot)] = 0.02$ the bias and $r(\boldsymbol{\theta}, \xi) = 0$ the additional statistical variation. Thus, it can be clearly seen that even in the simple example, bias in the function evaluation is present, suggesting that the discrepancy in the mean and limiting sample average between the ideal and observed values of noisy VQA function evaluations is an endemic issue, and classic optimization convergence analyses that do not take this into account are not directly applicable to studying the minimization of VQA circuits. Note that increasing the sample size does not diminish the error percentage, indicating that the procedure is not statistically consistent.

C Gate definitions

For the sake of completeness, let us reiterate the definition of gates used in Figure 4:

$$\begin{aligned}
U_2(\phi, \lambda) &= \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -e^{i\lambda} \\ e^{i\phi} & e^{i(\phi+\lambda)} \end{pmatrix}, & R_z(\lambda) &= \begin{pmatrix} e^{-i\frac{\lambda}{2}} & 0 \\ 0 & e^{i\frac{\lambda}{2}} \end{pmatrix}, \\
R(\theta, \phi) &= \begin{pmatrix} \cos \frac{\theta}{2} & -ie^{-i\phi} \sin \frac{\theta}{2} \\ -ie^{i\phi} \sin \frac{\theta}{2} & \cos \frac{\theta}{2} \end{pmatrix}.
\end{aligned}$$