

Stochastic optimization algorithms for quantum applications

J. Gidi^{1,*}, B. Candia¹, A. D. Muñoz-Moller¹, A. Rojas¹, L. Pereira², M. Muñoz³, L. Zambrano⁴, and A. Delgado¹

¹*Instituto Milenio de Investigación en Óptica y Departamento de Física, Facultad de Ciencias Físicas y Matemáticas, Universidad de Concepción, Casilla 160-C, Concepción, Chile*

²*Instituto de Física Fundamental IFF-CSIC, Calle Serrano 113b, Madrid 28006, Spain*

³*Departamento de Ingeniería Matemática y Centro de Investigación en Ingeniería Matemática (CIPMA), Facultad de Ciencias Físicas y Matemáticas, Universidad de Concepción, Casilla 160-C, Concepción, Chile*

⁴*ICFO-Institut de Ciències Fotoniques, The Barcelona Institute of Science and Technology, 08860 Castelldefels, Barcelona, Spain*



(Received 18 July 2023; accepted 14 August 2023; published 8 September 2023)

Hybrid classical quantum optimization methods have become an important tool for efficiently solving problems in the current generation of noisy intermediate-scale quantum computers. These methods use an optimization algorithm executed in a classical computer, fed with values of the objective function obtained in a quantum processor. A proper choice of optimization algorithm is essential to achieve good performance. Here, we review the use of first-order, second-order, and quantum natural gradient stochastic optimization methods, which are defined in the field of real numbers, and propose stochastic algorithms defined in the field of complex numbers. The performance of all methods is evaluated by means of their application to variational quantum eigensolver, quantum control of quantum states, and quantum state estimation. In general, complex number optimization algorithms perform best, with first-order complex algorithms consistently achieving the best performance, closely followed by complex quantum natural algorithms, which do not require expensive hyperparameter calibration. In particular, the scalar formulation of the complex quantum natural algorithm allows to achieve good performance with low classical computational cost.

DOI: [10.1103/PhysRevA.108.032409](https://doi.org/10.1103/PhysRevA.108.032409)

I. INTRODUCTION

The current generation of quantum hardware has been described as noisy intermediate-scale quantum (NISQ) devices [1], characterized by noisy entangling gates, short coherence times, and large sampling errors. A promising approach to achieve quantum advantage in NISQ devices are hybrid quantum-classical optimization algorithms [2–5]. These evaluate an objective function through a parametrized quantum circuit in a quantum computer and feed the values of the objective function to a classical optimization algorithm running on a classical computer. Thus, hybrid optimization algorithms are used whenever the objective function can be evaluated more efficiently on a quantum computer than on a classical one. This is the case for applications to quantum chemistry [6–8], quantum control [9–11], quantum simulation [12,13], entanglement detection [14–16], state estimation [17–21], quantum machine learning [22–26], error correction [27], graph theory [28–30], differential equations [31–33], and finances [34].

The performance of hybrid quantum-classical algorithms is affected by the optimization landscape associated with the objective function and the choice of the optimization algorithm. For instance, it has been recently shown [35,36] that a very general class of objective functions exhibits a barren plateau, that is, a region in the optimization landscape where the objective function gradient vanishes and its standard deviation

decreases exponentially with the number of qubits. In particular, this affects applications in which random quantum circuits are used, such as, for instance, quantum machine learning [23]. Also recently, several studies [37–40] have been carried out to establish general guidelines to choose the optimization algorithm with the best performance, according to a predefined metric, for a certain class of problems. These consider methods such as stochastic gradient descent [41], adaptive gradient algorithms [42], root-mean-square propagation [43], Adam and variations [44–46], the Nelder-Mead method [47], the Powell method [48], and the Newton conjugate gradient [49], among many others [50–56].

In the growing list of optimization methods used in hybrid optimization, stochastic optimization algorithms [57–61] play an important role. State initialization, quantum gates, and measurements are noisy processes leading to noisy evaluation of the objective function. This intrinsically stochastic behavior of the objective function negates mathematical guarantees on the convergence of commonly used classical optimization methods [39]. However, certain stochastic optimization methods have convergence proofs that admit the presence of noise. In this scenario, a method that achieves good performance in various applications of hybrid optimization is the simultaneous perturbation stochastic approximation (SPSA) method [62]. The main advantages of SPSA are its robustness to noise, ubiquitous in quantum mechanics, and that it can approximate the gradient of an objective function with only two measurements. In particular, this approximation does not require knowing the operational form of the objective function. SPSA has been successfully implemented in several experimental

*jorgegidi@udec.cl

platforms and is one of the standard methods for training variational quantum eigensolvers (VQEs) [63–67], quantum neural networks (QNNs) [68–70], and quantum tomography [17,18,71].

Given the success of stochastic optimization algorithms within quantum computing, efforts have been made to improve their performance in solving certain tasks. One proposal is second-order SPSA (2SPSA), which improves the convergence rate of SPSA by preconditioning the gradient with the inverse of a simultaneous perturbation estimate of the Hessian of the objective function [72,73]. This method is inspired by the deterministic Newton-Raphson algorithm and requires four evaluations of the objective function per iteration to estimate both gradient and Hessian. It has been shown that this method achieves a nearly optimal asymptotic error for well-conditioned problems. However, for a poorly conditioned Hessian, the error is several orders of magnitude larger [74]. Another proposal focused on quantum computing is quantum natural gradient optimization [75]. The SPSA algorithm explores the parameter space within a flat geometry, which can lead to an unfavorable update of parameters. In contrast, quantum natural gradient optimization uses information about the geometry of the parametric quantum state to update the parameters appropriately. The Fubini-Study metric tensor represents this information. Natural gradient optimization provides several advantages over vanilla (or standard) methods, that is, methods in their unmodified form. This is because the natural gradient is invariant under reparametrization [76] and approximately invariant under overparametrization [77]. The version of SPSA based on the quantum natural gradient (QN-SPSA) uses a simultaneous perturbation estimate of the Fubini-Study metric tensor [78]. This estimation requires four fidelity evaluations per iteration and the two function evaluations required to estimate the gradient. The fidelity evaluation can be performed efficiently using the swap test [79], among other alternatives [80]. This method is appropriate in contexts where the evaluation of the objective function is too expensive, for example, in estimating the fundamental energy of molecules [65,66,81,82]. However, similarly to 2SPSA, ill-conditioned metrics can reduce the performance of QN-SPSA [83–85].

Optimization methods can also be extended to work in the field of complex numbers by means of Wirtinger calculus [86]. Some examples are the complex Newton-Raphson algorithm [87] and the complex quantum natural gradient [88]. These methods optimize the objective function without resorting to the real and imaginary parts of complex variables. It has been argued in the literature that optimization methods formulated within the complex numbers could achieve better performance, which has been observed in a small set of examples [89–91]. This seems to be a more natural approach to optimization in quantum mechanics, where most functions have complex arguments. For example, continuous-variable quantum computing employs displacement and squeezing operators, which depend on complex parameters [92,93]. Recently, the complex simultaneous perturbation stochastic approximation (CSPSA) method [19] has been introduced. This is a generalization of SPSA that optimizes within the field of complex numbers. It has been shown that CSPSA can deliver better results in the estimation of pure states [19] and is robust against noise [21]. It has been applied

to entanglement estimation [16], quantum state discrimination [26], and violation of the Clauser-Horne-Shimony-Holt inequality [94].

Here, we present a comparative analysis of several stochastic optimization methods applied to real-valued functions of complex variables. We first review the basic principles of the SPSA algorithm. Subsequently, we review the 2SPSA and QN-SPSA algorithms using SPSA as a guideline. We also reviewed the CSPSA algorithm and developed two optimization algorithms based on the CSPSA algorithm: 2CSPSA and QN-CSPSA. These are the complex field formulations of their real counterparts 2SPSA and QN-SPSA, respectively. We study the performance of the introduced methods by comparing their convergence rate as a function of the number of iterations with respect to SPSA, 2SPSA, and QN-SPSA. This comparison is carried out in three contemporary applications: variational quantum eigensolver, quantum control, and quantum state estimation. We use a variational quantum eigensolver to obtain the ground-state energy of the Heisenberg Hamiltonian for a ten-qubit ring configuration, which is a ubiquitous and relatively simple model that describes the interactions within a chain of spins [63]. We implement the gradient ascent pulse engineering (GRAPE) method [95], which is used to engineer quantum gates and states. This method approximates a control pulse by a sequence of constant-intensity pulses. The control parameters of this pulse are optimized to find the best implementation of a given gate or state, even in the presence of noise [9]. In particular, we apply GRAPE to the generation of five-qubit pure states. Finally, in quantum state estimation, we implement self-guided quantum tomography (SGQT) [17], which is based on the minimization of the infidelity between an unknown state and a known parametrized state, to characterize six-qubit pure states. Since the studied optimization methods are stochastic, we use numerical simulations and sampling to estimate the mean, variance (or standard deviation), median, and interquartile range of the relevant figures of merits. Measurements are simulated using a finite sample of various sizes.

Our comparative analysis shows that the first-order CSPSA algorithm consistently performs best in all three applications. In the case of the variational quantum eigensolver, the performance of CSPSA is achieved at the expense of calibrating gain coefficients. Without the calibration, CSPSA performed poorly. To avoid the calibration, a good alternative is to use quantum natural algorithms, which achieve a performance close to the calibrated CSPSA at the expense of increasing the number of measurements and the classical computational cost. The latter can be avoided with the scalar version of quantum natural algorithms. On the other hand, for quantum state estimation and quantum control, the CSPSA algorithm performs the best without hyperparameter calibration.

While second-order algorithms do not provide an advantage, the quantum natural algorithms are competitive against first-order algorithms. As the number of qubits increases, we expect quantum natural algorithms to become more relevant. In this scenario, however, the cost of quantum natural algorithms increases. This can be partially mitigated using their scalar versions, which render the classical computational cost of quantum natural algorithms feasible even for a very large number of qubits.

In general, complex-based optimization methods tend to outperform real-based optimization methods, although the difference in performance may be slight.

This article is organized as follows: In Sec. II, we review the stochastic optimization methods SPSA, 2SPSA, QN-SPSA, and CSPA and formulate the methods 2CSPSA and QN-CSPSA. Also, we review and introduce modifications that may improve the performance of the methods. In Sec. III, we apply the previously developed optimization methods to variational quantum eigensolver, quantum control, and quantum state estimation. In Sec. IV, we summarize our main results and conclusions.

II. STOCHASTIC OPTIMIZATION ALGORITHMS

Let us consider the problem of optimizing a real function f of p complex variables, $f: \mathbb{C}^p \rightarrow \mathbb{R}$, that is, finding an argument $\mathbf{z}^* \in \mathbb{C}^p$ such that $f(\mathbf{z}^*)$ is a local minimum of the function f . This problem can be solved by mapping the complex variables to the field of the real numbers through the relation $\mathbf{z} = \mathbf{x} + i\mathbf{y}$, in which case f becomes $f(\boldsymbol{\theta})$ with $\boldsymbol{\theta} = (\mathbf{x}, \mathbf{y})^T \in \mathbb{R}^{2p}$. Then, one can use real-variable optimization algorithms to find $\boldsymbol{\theta}^* = (\mathbf{x}^*, \mathbf{y}^*)^T$ such that $f(\boldsymbol{\theta}^*)$ is a minimum of f , and retrieve the solution for the original complex-variable problem as $\mathbf{z}^* = \mathbf{x}^* + i\mathbf{y}^*$. It is possible, nevertheless, to solve the optimization problem using Wirtinger calculus [86,96], which does not resort to mapping complex variables to real ones.

While both approaches are equivalent, the process of solving one or the other is not. It has been conjectured that a complex-variable reformulation of real-variable optimization algorithms may lead to increased performance [89–91], which has been observed when working on pure-state quantum tomography [19]. Furthermore, for applications in quantum theory, which are natively stated in terms of complex variables, the transformation to real variables adds an extra step in the optimization process. For this reason, here we review some real-variable optimization methods relevant to quantum applications and present their complex-variable analogs.

A particularly suitable class of methods for optimizing multivariate functions in the presence of noisy measurements are the stochastic approximation methods [57–61]. This family of methods originates from the Robbins-Monro algorithm [97] designed to find a root θ of a function $M(x)$ given by the expectation of a random variable $Y(x)$. Here M is unknown, just like the probability function of Y , and the Robbins-Monro algorithm gives an estimate of θ by making successive observations on Y . From the Robbins-Monro algorithm, it is possible to consider M as a regression function [98] and propose a scheme to estimate the maximum of M . Therefore, the use of stochastic approximations arises to deliver an algorithm that converges to an optimal value of a function f using the Kiefer-Wolfowitz procedure when $M = \nabla f$.

A widely used family of stochastic approximation (SA) methods is based on the iterative rule

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

where the descent step series $a_k = a/(k+A)^s$ is fixed by the externally selected gain parameters a , A , and s . The quantity \mathbf{g}_k is a stochastic approximation of the gradient of the

objective function at $\boldsymbol{\theta}_k$, which depends on the gain coefficient $b_k = b/k^t$, where b and t are externally fixed gain parameters.

In the following sections, we review the SPSA algorithm and its extension to the second-order and quantum natural gradient algorithms, 2SPSA and QN-SPSA, respectively. Subsequently, we review the CSPA algorithm for complex variables and develop two extensions to it: the second-order algorithm 2CSPSA and the quantum natural gradient algorithm QN-CSPSA. This work is conducted, in a similar way to the SPSA algorithm, by considering an iterative rule as Eq. (1) for the case of complex variables. Last, we present typical modifications to improve the performance of the optimization algorithms, namely, blocking and resampling, and introduce two further variations: an alternative Hessian postprocessing procedure and a scalar approximation to second-order and quantum natural algorithms that reduce their classical computational cost.

A. Real-variable methods

1. SPSA

The simultaneous perturbation stochastic approximation (SPSA) is a multivariate optimization method for real functions of real variables. While the SPSA denomination came later, the method was first presented by Spall [62] and corresponded to an improvement over the finite difference stochastic approximation (FDSA) from Kiefer and Wolfowitz [98]. Both the FDSA and SPSA algorithms optimize the function $f(\boldsymbol{\theta})$ with $\boldsymbol{\theta} \in \mathbb{R}^p$ by following the recursive stochastic approximation rule [Eq. (1)]. However, the main feature of SPSA is that, instead of estimating each of the p components of the gradient as a stochastic finite difference approximation, it defines the estimator \mathbf{g}_k as

$$\mathbf{g}_k(\boldsymbol{\theta}) = \frac{f(\boldsymbol{\theta} + b_k \boldsymbol{\Delta}_k) - f(\boldsymbol{\theta} - b_k \boldsymbol{\Delta}_k)}{2b_k} \begin{pmatrix} 1/\Delta_{k,1} \\ \vdots \\ 1/\Delta_{k,p} \end{pmatrix}, \quad (2)$$

where $\boldsymbol{\Delta}_k$ is a random perturbation vector with p components typically chosen from the set $\{\pm 1\}$ with uniform probability, and the finite difference approximation step $b_k = b/k^t$ is controlled by the externally selected gain parameters b and t . It is worth noting that while $\mathbf{g}_k(\boldsymbol{\theta}_k)$ does not necessarily have the direction of the gradient at each iteration, it is an asymptotically unbiased estimator of the gradient, meaning that it converges at the statistical limit to the same solution as following the gradient. Furthermore, Eq. (2) makes the SPSA algorithm especially suitable for high-dimensional problems since it always requires two function evaluations per iteration, regardless of the number p of variables, in contrast to the FDSA algorithm that requires $2p$ function evaluations per iteration.

An iteration of the SPSA algorithm is given by Eqs. (1) and (2) and requires a total of two objective function evaluations.

2. 2SPSA

Since the iterative rule used in the SPSA algorithm is derived from a first-order gradient descent approximation, the rate of convergence of the algorithm could be accelerated using a second-order iterative rule coming from the

Newton-Raphson method, given by

$$\theta_{k+1} = \theta_k - \eta [\mathcal{H}(\theta_k)]^{-1} \left(\frac{\partial f}{\partial \theta}(\theta_k) \right)^T, \quad (3)$$

where $\eta \in \mathbb{R}^+$ is the learning rate and \mathcal{H} is the Hessian of f . A stochastic approximation based on Eq. (3) is proposed by Spall [72], deriving the so-called adaptive or second-order SPSA (2SPSA) algorithm. The iterative rule now yields

$$\theta_{k+1} = \theta_k - \bar{a}_k \bar{\mathcal{H}}_k^{-1} \mathbf{g}_k(\theta_k), \quad (4)$$

where $\bar{a}_k = 1/(k+A)^s$ no longer depends on a . The gradient estimator \mathbf{g}_k is defined by Eq. (2), as in the first-order case, and $\bar{\mathcal{H}}_k$ is a modified version of the simultaneous perturbation stochastic approximation of the Hessian matrix. In particular, we compute $\bar{\mathcal{H}}_k$ by [72]

$$\mathcal{H}'_k = \frac{\mathcal{H}_k + [\mathcal{H}_k]^T}{2}, \quad (5a)$$

$$\mathcal{H}''_k = \frac{k}{k+1} \mathcal{H}''_{k-1} + \frac{1}{k+1} \mathcal{H}'_k, \quad (5b)$$

$$\bar{\mathcal{H}}_k = \sqrt{\mathcal{H}''_k} + \varepsilon I, \quad (5c)$$

where, in execution order, Eq. (5a) ensures that the Hessian approximation is symmetric as the analytical Hessian, then Eq. (5b) stabilizes the estimator by introducing inertia from previous iterations, starting from an identity at the zeroth iteration, $\mathcal{H}''_0 = I$, and finally, Eq. (5c) with $0 < \varepsilon \ll 1$ guarantees positive definiteness.

A one-sided simultaneous perturbation stochastic approximation to the Hessian matrix is taken as

$$\mathcal{H}_k(\theta) = \frac{\mathbf{g}_k(\theta + \tilde{b}_k \tilde{\Delta}_k) - \mathbf{g}_k(\theta)}{\tilde{b}_k} \begin{pmatrix} 1/\tilde{\Delta}_{k,1} \\ \vdots \\ 1/\tilde{\Delta}_{k,p} \end{pmatrix}^T, \quad (6)$$

which allows reusing the function evaluations from the centered gradient estimator. By inserting the definition of the gradient approximation [Eq. (2)], then Eq. (6) can be rewritten by components as

$$[\mathcal{H}_k]_{ij} = \frac{\delta^2 f_k(\theta)}{2b_k \tilde{b}_k \Delta_{k,i} \tilde{\Delta}_{k,j}}, \quad (7)$$

where

$$\delta^2 f_k(\theta) = f(\theta + b_k \Delta_k + \tilde{b}_k \tilde{\Delta}_k) - f(\theta + b_k \Delta_k) - f(\theta - b_k \Delta_k + \tilde{b}_k \tilde{\Delta}_k) + f(\theta - b_k \Delta_k), \quad (8)$$

$\tilde{b}_k = \tilde{b}/k'$ is a gain series similar in nature to b_k , and $\tilde{\Delta}_k$ is a random vector formed by p components uniformly generated from the set $\{\pm 1\}$ analogous to Δ_k .

Thereby, an iteration of this method is given by Eqs. (8), (7), (5), (2), and (4) and requires a total of four objective function evaluations.

3. QN-SPSA

The gradient descent method reaches a local minimum by moving, at each iteration, along the direction of the steepest descent of the objective function in the Euclidian parameter

space, $-(\partial f / \partial \theta)^T$, limiting the magnitude of the update step, $\Delta \theta$. The steepest-descent rule can be obtained by choosing the increment as

$$\Delta \theta = \arg \min_{\Delta \theta \in \mathbb{R}^{2p}} \left\{ \left\langle \left(\frac{\partial f}{\partial \theta} \right)^T, \Delta \theta \right\rangle + \frac{1}{2\eta} \|\Delta \theta\|_2^2 \right\}, \quad (9)$$

where $\eta \in \mathbb{R}^+$ is the learning rate, $\langle \theta, \theta' \rangle = \theta^T \theta'$ is the inner product for two vectors θ and θ' , respectively, and $\|\cdot\|_2 = \sqrt{\langle \cdot, \cdot \rangle}$ is the l^2 norm. Differentiating the argument at the right-hand side of Eq. (9) with respect to $\Delta \theta$ and setting it to 0 provides the well-known gradient descent step

$$\Delta \theta = -\eta \left(\frac{\partial f}{\partial \theta} \right)^T. \quad (10)$$

This result is based on the l^2 geometry, where a shift in any direction in the parameter space is equally weighted. However, the objective function may not be equally sensitive to changes in different parameters and, therefore, a more adequate notion of distance would measure the step length $\Delta \theta$ by weighting the changes on each parameter. This is addressed by a method called natural gradient descent [99], which endows the parameter space with a suitable metric \mathcal{G} that induces the norm $\|\cdot\|_{\mathcal{G}} = \sqrt{\langle \cdot, \mathcal{G} \cdot \rangle}$. Then, the increment is stated as

$$\Delta \theta = \arg \min_{\Delta \theta \in \mathbb{R}^{2p}} \left\{ \left\langle \left(\frac{\partial f}{\partial \theta} \right)^T, \Delta \theta \right\rangle + \frac{1}{2\eta} \|\Delta \theta\|_{\mathcal{G}}^2 \right\}, \quad (11)$$

which leads to the natural gradient descent rule

$$\theta_{k+1} = \theta_k - \eta [\mathcal{G}(\theta_k)]^{-1} \left(\frac{\partial f}{\partial \theta}(\theta_k) \right)^T. \quad (12)$$

The quantum natural method, which takes \mathcal{G} as the Fubini-Study metric tensor, is particularly useful for improving convergence rates for optimization problems in quantum applications [75]. The Fubini-Study metric tensor is proportional to the quantum Fisher information matrix, so its calculation can be very expensive when many variables are involved. This problem was addressed by Gacon *et al.* [78] by taking advantage of the similarity between Eqs. (3) and (12), along with the possibility of writing the Fubini-Study metric tensor as

$$\mathcal{G}(\theta) = -\frac{1}{2} \left[\frac{\partial}{\partial \theta} \left(\frac{\partial F(\theta', \theta)}{\partial \theta} \right)^T \right] \bigg|_{\theta'=\theta}, \quad (13)$$

where $F(\theta', \theta)$ is the fidelity between two quantum states parametrized with the variables θ' and θ , respectively. In particular, the Fubini-Study metric tensor was approximated according to Eq. (13) using the stochastic approximation of the Hessian employed by the 2SPSA algorithm. In this manner, they proposed the quantum natural SPSA (QN-SPSA) algorithm, which avoids the curse of dimensionality.

In order to reuse the equations already presented for the 2SPSA method, we will abuse notation and denote \mathcal{H} the Hessian estimate of the Fubini-Study metric, yielding

$$[\mathcal{H}_k]_{ij} = -\frac{\delta^2 F_k(\theta_k)}{4b_k \tilde{b}_k \Delta_{k,i} \tilde{\Delta}_{k,j}}, \quad (14)$$

where

$$\begin{aligned} \delta^2 F_k(\boldsymbol{\theta}) = & F(\boldsymbol{\theta}, \boldsymbol{\theta} + b_k \boldsymbol{\Delta}_k + \tilde{b}_k \tilde{\boldsymbol{\Delta}}_k) \\ & - F(\boldsymbol{\theta}, \boldsymbol{\theta} + b_k \boldsymbol{\Delta}_k) \\ & - F(\boldsymbol{\theta}, \boldsymbol{\theta} - b_k \boldsymbol{\Delta}_k + \tilde{b}_k \tilde{\boldsymbol{\Delta}}_k) \\ & + F(\boldsymbol{\theta}, \boldsymbol{\theta} - b_k \boldsymbol{\Delta}_k), \end{aligned} \quad (15)$$

and $\boldsymbol{\Delta}$ and $\tilde{\boldsymbol{\Delta}}$ are two vectors of p components randomly sampled from the set $\{\pm 1\}$.

Following the same logic as in the 2SPSA algorithm, the simultaneous perturbation stochastic approximation of the Hessian in Eq. (14) must be conditioned by the procedure on the system of Eqs. (5). Let us note that while we are using the 2SPSA discretization scheme and update rule, this is a first-order method, as the conditioner \mathcal{H}_k comes not from a second-order expansion on the target function but only from a different metric in the parameter space.

Requiring only two measurements of the objective function makes this algorithm especially suitable for problems where the metric tensor can be efficiently approximated. That is when evaluating the fidelity F between two known pure quantum states requires marginal resources compared to the potentially expensive target function f .

An iteration of the QN-SPSA method is given by Eqs. (15), (14), (5), (2), and (4) and requires a total of two objective function evaluations and four fidelity evaluations.

B. Complex-variable methods

Now we formulate the problem of optimizing real-valued functions of complex variables. In the case of quantum mechanics, most of the functions that interest us depend on complex variables and their complex conjugates. Consequently, these functions do not satisfy the Cauchy-Riemann conditions and lack a Taylor-series expansion. This can be solved by resorting to the real and imaginary parts of the complex variables. Wirtinger calculus [86], however, allows us to define a derivative, the Wirtinger derivative, that exists even for nonholomorphic functions. We consider a function $f: \boldsymbol{\mu} \in \mathbb{C}^{2p} \rightarrow \mathbb{R}$ with $\boldsymbol{\mu} = (z, z^*)^T$, which can be expressed in a power series for a complex increment $\Delta\boldsymbol{\mu} = (\Delta z, \Delta z^*)^T$,

$$f(\boldsymbol{\mu} + \Delta\boldsymbol{\mu}) = f(\boldsymbol{\mu}) + \frac{\partial f}{\partial \boldsymbol{\mu}} \Delta\boldsymbol{\mu} + \frac{1}{2} \Delta\boldsymbol{\mu}^\dagger \mathcal{H} \Delta\boldsymbol{\mu} + \dots, \quad (16)$$

where

$$\mathcal{H} = \frac{\partial}{\partial \boldsymbol{\mu}} \left(\frac{\partial f}{\partial \boldsymbol{\mu}} \right)^\dagger \quad (17)$$

is the complex Hessian of the function f [87], the symbol (\dagger) denotes the conjugate transpose, and differentiation with respect to $\boldsymbol{\mu}$ is defined by

$$\frac{\partial f}{\partial \boldsymbol{\mu}} = \left(\frac{\partial f}{\partial z}, \frac{\partial f}{\partial z^*} \right), \quad (18)$$

where the complex variables z and z^* are considered to be independent. Let us note that the inner product between any column two vectors $\boldsymbol{\mu} = (z \ z^*)^T$ and $\boldsymbol{\mu}' = (z' \ z'^*)^T$, with

$z, z' \in \mathbb{C}^p$, is a real number,

$$\boldsymbol{\mu}^\dagger \boldsymbol{\mu}' = (z^* \ z) \begin{pmatrix} z' \\ z'^* \end{pmatrix} = 2\text{Re}\{z^\dagger z'\}. \quad (19)$$

1. CSPSA

Performing a first-order approximation on $|\Delta\boldsymbol{\mu}|$ from Eq. (16), that is,

$$f(\boldsymbol{\mu} + \Delta\boldsymbol{\mu}) - f(\boldsymbol{\mu}) \approx \frac{\partial f}{\partial \boldsymbol{\mu}} \Delta\boldsymbol{\mu}, \quad (20)$$

we obtain that the largest decrease of the function f is achieved by a perturbation $\Delta\boldsymbol{\mu}$ in the direction of $-(\partial f / \partial \boldsymbol{\mu})^\dagger$. This provides the complex equivalent to the gradient descent update rule, which is given by the expression

$$\boldsymbol{\mu}_{k+1} = \boldsymbol{\mu}_k - \eta \left(\frac{\partial f}{\partial \boldsymbol{\mu}} \right)^\dagger, \quad (21)$$

where $\eta \in \mathbb{R}^+$ is the learning rate. The above equation yields a stochastic approximation [19] used to introduce the CSPSA algorithm given by the iterative rule

$$\mathbf{z}_{k+1} = \mathbf{z}_k - a_k \mathbf{g}_k(\mathbf{z}_k), \quad (22)$$

where $a_k = a/(k+A)^s$. The gradient estimator is now given by

$$\mathbf{g}_k(\mathbf{z}) = \frac{f(\mathbf{z} + b_k \boldsymbol{\Delta}_k) - f(\mathbf{z} - b_k \boldsymbol{\Delta}_k)}{2b_k} \begin{pmatrix} 1/\Delta_{k,1}^* \\ \vdots \\ 1/\Delta_{k,p}^* \end{pmatrix}, \quad (23)$$

where $b_k = b/k^t$ and $\boldsymbol{\Delta}_k$ is a random vector with p components uniformly generated from the set $\{\pm 1, \pm i\}$, with i the imaginary unit. To keep the notation simple, we have omitted the dependency of \mathbf{g}_k on \mathbf{z}^* . Consequently, we write $\mathbf{g}_k(\mathbf{z}, \mathbf{z}^*)$ as $\mathbf{g}_k(\mathbf{z})$ and similarly for other functions.

An iteration of the CSPSA method is given by Eqs. (23) and (22) and requires a total of two objective function evaluations.

2. 2CSPSA

To obtain a second-order iterative rule, we add up to second-order terms on $|\Delta\boldsymbol{\mu}|$ from expansion (16) and consider the problem of finding the perturbation $\Delta\boldsymbol{\mu}$ that minimizes $f(\boldsymbol{\mu} + \Delta\boldsymbol{\mu})$. This is done by taking $\partial f(\boldsymbol{\mu} + \Delta\boldsymbol{\mu}) / \partial \Delta\boldsymbol{\mu} = 0$, which reduces to the equation

$$\left[\frac{\partial}{\partial \boldsymbol{\mu}} \left(\frac{\partial f}{\partial \boldsymbol{\mu}} \right)^\dagger \right] \Delta\boldsymbol{\mu} = - \left(\frac{\partial f}{\partial \boldsymbol{\mu}} \right)^\dagger. \quad (24)$$

This can be rewritten in terms of \mathbf{z} and \mathbf{z}^* as

$$\begin{pmatrix} \mathcal{H}_{zz} & \mathcal{H}_{zz^*} \\ \mathcal{H}_{z^*z} & \mathcal{H}_{z^*z^*} \end{pmatrix} \begin{pmatrix} \Delta z \\ \Delta z^* \end{pmatrix} = - \begin{pmatrix} [\partial f / \partial z]^\dagger \\ [\partial f / \partial z^*]^\dagger \end{pmatrix}, \quad (25)$$

where the elements of the block matrix are

$$\mathcal{H}_{zz} = \frac{\partial}{\partial z} \left(\frac{\partial f}{\partial z} \right)^\dagger, \quad (26)$$

$$\mathcal{H}_{zz^*} = \frac{\partial}{\partial z} \left(\frac{\partial f}{\partial z^*} \right)^\dagger, \quad (27)$$

$$\mathcal{H}_{z^*z} = \mathcal{H}_{zz^*}^\dagger, \text{ and} \quad (28)$$

$$\mathcal{H}_{z^*z^*} = \mathcal{H}_{zz}^*. \quad (29)$$

The system of Eqs. (25) has the solution

$$\Delta \mathbf{z} = (\mathcal{H}_{zz} - \mathcal{H}_{z^*z} \mathcal{H}_{z^*z^*}^{-1} \mathcal{H}_{zz^*})^{-1} \times \left\{ \mathcal{H}_{z^*z} \mathcal{H}_{z^*z^*}^{-1} \left(\frac{\partial f}{\partial z^*} \right)^\dagger - \left(\frac{\partial f}{\partial z} \right)^\dagger \right\}, \quad (30)$$

which is the update step corresponding to a Newton algorithm [87]. While this solution requires a large number of operations, it is customary to use a block-diagonal approximation, $\mathcal{H}_{z^*z^*} \approx 0$, yielding a pseudo-Newton method [100] with

$$\Delta \mathbf{z} = -\mathcal{H}_{zz}^{-1} \left(\frac{\partial f}{\partial z} \right)^\dagger, \quad (31)$$

which also has the advantage of being operationally independent of z^* in practice.

Analogous to the 2SPSA method, in the stochastic approximation, we take the descent direction given by Eq. (31). Thereby, we define the 2CSPSA algorithm by means of the update rule

$$\mathbf{z}_{k+1} = \mathbf{z}_k - \bar{a}_k [\bar{\mathcal{H}}_k(\mathbf{z}_k)]^{-1} \mathbf{g}_k(\mathbf{z}_k), \quad (32)$$

where $\bar{a}_k = 1/(k + A)^s$, $\mathbf{g}_k(\mathbf{z})$ is given by Eq. (23), and $\bar{\mathcal{H}}_k$ is a modified version of the simultaneous perturbation stochastic approximation for the partial complex Hessian \mathcal{H}_{zz} at the k th iteration. Similar to the system of Eqs. (5) for the real-variable case, $\bar{\mathcal{H}}_k$ is computed through the sequence

$$\mathcal{H}'_k = \frac{\mathcal{H}_k + [\mathcal{H}_k]^\dagger}{2}, \quad (33a)$$

$$\mathcal{H}''_k = \frac{k}{k+1} \mathcal{H}''_{k-1} + \frac{1}{k+1} \mathcal{H}'_k, \quad (33b)$$

$$\bar{\mathcal{H}}_k = \sqrt{\mathcal{H}''_k} + \varepsilon I, \quad (33c)$$

where, in execution order, Eq. (33a) makes the Hessian approximation Hermitian as the exact Hessian, then Eq. (33b) stabilizes the estimator by introducing inertia from previous iterations, starting from an identity at the zeroth iteration, that is, $\mathcal{H}''_0 = I$, and finally Eq. (33c) with $0 < \varepsilon \ll 1$ ensures positive definiteness. Note that the regularization equation (33c) is still valid in the complex-variable case since its input, \mathcal{H}''_k , has real eigenvalues due to the previous Hermitization in Eq. (33a).

In this case, the components of the simultaneous perturbation stochastic approximation of the partial complex Hessian \mathcal{H}_{zz} are given by

$$[\mathcal{H}_k(\mathbf{z})]_{ij} = \frac{\delta^2 f_k(\mathbf{z})}{2b_k \tilde{b}_k \Delta_{k,i}^* \tilde{\Delta}_{k,j}}, \quad (34)$$

where

$$\delta^2 f_k(\mathbf{z}) = f(\mathbf{z} + b_k \mathbf{\Delta}_k + \tilde{b}_k \tilde{\mathbf{\Delta}}_k) - f(\mathbf{z} + b_k \mathbf{\Delta}_k) - f(\mathbf{z} - b_k \mathbf{\Delta}_k + \tilde{b}_k \tilde{\mathbf{\Delta}}_k) + f(\mathbf{z} - b_k \mathbf{\Delta}_k), \quad (35)$$

and $\mathbf{\Delta}$ and $\tilde{\mathbf{\Delta}}$ are two random vectors, each composed by p elements uniformly generated from the set $\{\pm 1, \pm i\}$.

The 2CSPSA method requires the inversion and regularization of a $p \times p$ Hermitian complex matrix. In contrast, the analog 2SPSA optimization of an equivalent problem would require the inversion and regularization of a $2p \times 2p$ symmetric real matrix.

An iteration of this method is given by Eqs. (35), (34), (33), (23), and (32) and requires a total of four objective function evaluations.

3. QN-CSPSA

The natural gradient method was adapted [88] for a complex parameter space by posing the usual natural gradient update rule [Eq. (12)] with the relevant metric \mathcal{G} and using an invertible linear transformation W to move back and forth between the real and complex parametrizations such that

$$W \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} = \begin{pmatrix} \mathbf{x} + i\mathbf{y} \\ \mathbf{x} - i\mathbf{y} \end{pmatrix} := \begin{pmatrix} \mathbf{z} \\ \mathbf{z}^* \end{pmatrix}, \quad (36)$$

where $\mathbf{x}, \mathbf{y} \in \mathbb{R}^p$. However, continuously moving between parametrizations is undesirable, and therefore here we present a natively complex implementation of the natural gradient method for quantum applications, which proceeds analogously to the QN-SPSA method.

The complex gradient descent rule [Eq. (21)] can be obtained as a solution to the optimization problem

$$\Delta \boldsymbol{\mu} = \arg \min_{\Delta \boldsymbol{\mu} \in \mathbb{C}^{2p}} \left\{ \left\langle \left(\frac{\partial f}{\partial \boldsymbol{\mu}} \right)^\dagger, \Delta \boldsymbol{\mu} \right\rangle + \frac{1}{2\eta} \|\Delta \boldsymbol{\mu}\|_2^2 \right\}, \quad (37)$$

where $\eta \in \mathbb{R}^+$ is the learning rate, $\langle \boldsymbol{\mu}, \boldsymbol{\mu}' \rangle = \boldsymbol{\mu}^\dagger \boldsymbol{\mu}'$ is the inner product for two complex vectors $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$, respectively, and $\|\cdot\|_2 = \sqrt{\langle \cdot, \cdot \rangle}$ is the l^2 norm. As in the real case, to require the parameter update to remain small in the space endowed with metric \mathcal{G} , the l^2 norm is replaced in Eq. (37) by $\|\cdot\|_{\mathcal{G}} = \sqrt{\langle \cdot, \mathcal{G} \cdot \rangle}$. This leads to the optimization problem

$$\Delta \boldsymbol{\mu} = \arg \min_{\Delta \boldsymbol{\mu} \in \mathbb{C}^{2p}} \left\{ \left\langle \left(\frac{\partial f}{\partial \boldsymbol{\mu}} \right)^\dagger, \Delta \boldsymbol{\mu} \right\rangle + \frac{1}{2\eta} \|\Delta \boldsymbol{\mu}\|_{\mathcal{G}}^2 \right\}, \quad (38)$$

which has the solution

$$\Delta \boldsymbol{\mu} = -\eta \mathcal{G}^{-1} \left(\frac{\partial f}{\partial \boldsymbol{\mu}} \right)^\dagger, \quad (39)$$

where \mathcal{G} is an Hermitian matrix.

In the case that the optimization space is the set of pure quantum states, the metric \mathcal{G} can be chosen proportional to the quantum Fisher complex information matrix [101], that is,

$$\mathcal{G} = -\frac{1}{2} \left[\frac{\partial}{\partial \boldsymbol{\mu}} \left(\frac{\partial F(\boldsymbol{\mu}', \boldsymbol{\mu})}{\partial \boldsymbol{\mu}} \right)^\dagger \right] \bigg|_{\boldsymbol{\mu}' = \boldsymbol{\mu}}, \quad (40)$$

where $F(\boldsymbol{\mu}', \boldsymbol{\mu})$ is the fidelity between two states parametrized with variables $\boldsymbol{\mu}'$ and $\boldsymbol{\mu}$, respectively.

Considering, as in the 2CSPSA case, a block-diagonal approximation of \mathcal{G} , the first row of Eq. (39) yields

$$\Delta \mathbf{z} = -\eta \mathcal{G}_{zz}^{-1} \left(\frac{\partial f}{\partial \mathbf{z}} \right)^\dagger, \quad (41)$$

where \mathcal{G}_{zz} is the top left block of \mathcal{G} .

Given the Hessian form of \mathcal{G}_{zz} and considering the similarity of Eqs. (41) and (31), we can borrow the discretization scheme from 2CSPSA to approximate \mathcal{G}_{zz} . Denoting \mathcal{H}_k as the simultaneous perturbation stochastic approximation of \mathcal{G}_{zz} at iteration k allows us to reuse the equations already presented for 2CSPSA, giving

$$[\mathcal{H}_k]_{ij} = -\frac{\delta^2 F_k(\mathbf{z}_k)}{4b_k \tilde{b}_k \Delta_{k,i}^* \tilde{\Delta}_{k,j}}, \quad (42)$$

with

$$\begin{aligned} \delta^2 F_k(\mathbf{z}) = & F(\mathbf{z}, \mathbf{z} + b_k \mathbf{\Delta}_k + \tilde{b}_k \tilde{\mathbf{\Delta}}_k) \\ & - F(\mathbf{z}, \mathbf{z} + b_k \mathbf{\Delta}_k) \\ & - F(\mathbf{z}, \mathbf{z} - b_k \mathbf{\Delta}_k + \tilde{b}_k \tilde{\mathbf{\Delta}}_k) \\ & + F(\mathbf{z}, \mathbf{z} - b_k \mathbf{\Delta}_k), \end{aligned} \quad (43)$$

and conditioning as in the system of Eqs. (33). As before, $\mathbf{\Delta}$ and $\tilde{\mathbf{\Delta}}$ are two random vectors, each composed by p elements uniformly generated from the set $\{\pm 1, \pm i\}$.

An iteration of this method, which we call quantum natural CSPSA (QN-CSPSA), is given by Eqs. (43), (42), (33), (23), and (32) and requires a total of two objective function evaluations and four fidelity evaluations.

C. Method improvements

In the previous sections, optimization methods were presented in their vanilla form. It is possible, however, to introduce further modifications that can improve their convergence properties. In particular, we will address two typical modifications, blocking and resampling, and two extra variations we propose: an alternative Hessian postprocessing procedure and a scalar approximation for the preconditioned methods.

1. Blocking

This technique consists of blocking the progression of the method if the updated parameters \mathbf{z}_{k+1} fail to fulfill a given criterion. Conventionally, the updated variable is required to improve the value of the objective function with respect to the previous iteration plus some fixed non-negative tolerance,

$$f(\mathbf{z}_{k+1}) < f(\mathbf{z}_k) + \delta. \quad (44)$$

The tolerance δ is usually set as twice the approximate standard deviation of the noise in the objective function evaluation, which can be estimated by collecting several evaluations at the initial value of the parameters [72].

Regardless of whether the step is accepted, the Hessian estimate \mathcal{H}_k'' from Eqs. (5b) and (33b) must be updated at every iteration.

2. Resampling

This technique is also known as gradient or Hessian averaging. It consists in computing the random estimators for the gradient and Hessian N_R times per iteration to perform the corresponding variable update using the average of these estimators. This practice is recommended in noisy environments [72].

Note that the authors of QN-SPSA [78] implement resampling only for the Hessian estimator with the premise that evaluating the metric is cheaper than evaluating the objective function, which could lead to a better convergence rate with little increment on the experimental resources. However, here we stick to the convention stated in Ref. [72], which is also implemented on QISKIT [102].

3. Postprocessing

Several postprocessing procedures have been proposed to improve the stability of the 2SPSA algorithm [60]. We consider two alternatives: the original proposal given by Eqs. (5), and the procedure given by

$$\mathcal{H}'_k = \frac{\mathcal{H}_k + [\mathcal{H}_k]^\dagger}{2}, \quad (45a)$$

$$\mathcal{H}''_k = \sqrt{\mathcal{H}_k'^2 + \varepsilon I}, \quad (45b)$$

$$\overline{\mathcal{H}}_k = \frac{k}{k+1} \overline{\mathcal{H}}_{k-1} + \frac{1}{k+1} \mathcal{H}''_k. \quad (45c)$$

4. Scalar preconditioning approximation

Preconditioned methods, such as 2SPSA, 2CSPSA, QN-SPSA, and QN-CSPSA, adaptively adjust the descent direction and magnitude by adding a preconditioner to the stochastic approximation. However, these methods can exhibit numerical instabilities due to the inversion of a possibly ill-conditioned Hessian estimation. Postprocessing procedures can partially mitigate these issues, but these methods still lack consistency in numerical simulations compared to first-order methods. Most likely, these problems are caused by an inadequate adjustment of the descent direction. We consider these problems most likely induced by an inadequate adjustment of the descent direction.

It has been suggested [74] to replace the Hessian estimation with a scalar function of its eigenvalues. Thereby, the descent direction is chosen according to the first-order gradient estimator while retaining the descent magnitude adaptivity from the preconditioner. Following these considerations, we propose a scalar approximation to the Hessian estimates (7), (14), (34), and (42). Specifically, we omit the stochastic perturbations $\mathbf{\Delta}_k$ and $\tilde{\mathbf{\Delta}}_k$ presented in the Hessian estimates to only adjust the descent magnitude and preserve the first-order descent direction. Namely, we approximate the Hessian estimate of the second-order methods by

$$\mathcal{H}'_k = \frac{\delta^2 f_k(\mathbf{z})}{2b_k \tilde{b}_k}, \quad (46)$$

and the Hessian estimates for quantum natural optimizers by

$$\mathcal{H}' = -\frac{\delta^2 F_k}{4b_k \tilde{b}_k}. \quad (47)$$

From this procedure, we consider a new set of second-order and quantum natural methods where the computational complexity is reduced. Namely, the number of classical operations on each iteration is reduced from $O(p^3)$ to $O(p)$, where p is the number of variables.

III. APPLICATIONS

We study the performance of the above optimization methods by comparing the rate of convergence of the objective function towards the minimum as a function of the number of iterations. We consider three important applications: variational quantum eigensolver, quantum control, and quantum state estimation. We use the variational quantum eigensolver to obtain the ground-state energy of the Heisenberg Hamiltonian, which is a ubiquitous and relatively simple model that describes the interactions within a chain of spins. For quantum control, we implement the GRAPE method [95], which approximates a control pulse by a sequence of constant-intensity pulses. The control parameters of this pulse are optimized to find the best implementation of a given state, even in the presence of noise [9]. Finally, for quantum state estimation, we implement SGQT [17], based on minimizing the infidelity between an unknown state and a known parametrized state.

The studied optimization methods are stochastic. We use ensembles of numerical simulations to estimate the mean, standard deviation, median, and interquartile range of the objective function. Measurements are simulated by sampling a multinomial distribution with various numbers of trials. In the figures below, only the upper half of the standard deviation is shown.

We test the optimization methods considering different configurations and look for the ones that offer the best performance. The configurations we tested are all possible combinations of the following alternatives: with or without blocking, resampling with $N_R = 1, 2, 5$, the two basic postprocessing procedures, Eqs. (33) or (45), and standard, asymptotic or static set of gain coefficients. The standard set is given by $a = 3$, $b = 0.1$, $A = 0$, $s = 0.602$, and $t = 0.101$, the asymptotic set by $a = 3$, $b = 0.1$, $A = 0$, $s = 1$, and $t = 1/6$, and the static set by $a = 0.01$, $b = 0.01$, $A = 0$, $s = 0$ and $t = 0$.

For clarity, we consider simulations with two groups of methods: (i) vanilla methods and (ii) improved methods, that is, the vanilla methods implemented with the improvements proposed in Sec. II C. The reason behind this separation lies in the drastic increase in resources required to perform blocking and resampling, and it could be useful to be able to discriminate when it is really worth swapping resources for better results.

Tables I to VI provide the value of the statistical indicators; median, interquartile range (IQR), mean, and standard deviation (STD), obtained through numerical simulations for the best configuration of each optimization method on each application. These values were used to determine the algorithm with the best performance after 700, 1000, and 5000 iterations for the variational quantum eigensolver, quantum control of quantum states, and self-guided quantum tomography, respectively. We indicate the gain coefficients and equations used for postprocessing for each vanilla method. In the case of improved methods, we also indicate the amount of resampling and the use of blocking.

We have created a freely available library [103] that contains the codes in the JULIA programming language [104] that implements all of the optimization methods.

A. Variational quantum eigensolver

The search for the ground state and its energy E_0 of a Hamiltonian is a problem of great interest in areas such as computational chemistry and condensed matter physics. This is because much of the phenomenology and properties of quantum systems can be studied from the ground state and its energy. However, finding this eigenstate in large systems is not a trivial task. It is often infeasible due to the exponential growth of the dimension of the Hilbert space with respect to the number of subsystems. For large systems, the Rayleigh-Ritz method [105,106] is a useful tool since it is limited to searching a parametrized subset of the original Hilbert space to reduce the computational cost of optimization. A further reduction in computational cost is achieved using the VQE method [107]. This consists of performing the Rayleigh-Ritz method with the help of a classical and a quantum computer, which makes it a promising tool for the current generation of quantum technologies.

The goal is to find the eigenstate $|\psi_0\rangle$ associated with the lowest eigenvalue E_0 of a Hamiltonian. This ground state can be characterized as the solution to the optimization problem

$$E_0 = \min_{|\psi\rangle} \frac{\langle\psi|H|\psi\rangle}{\langle\psi|\psi\rangle}. \quad (48)$$

The Rayleigh-Ritz method provides an estimate of E_0 by parametrizing the trial states as $|\psi(\theta)\rangle$ and optimizing over the vector θ of parameters. The underlying idea is that the subset defined by the parametrization must have a smaller dimension than the total Hilbert space to reduce the computational cost.

The VQE method considers the generic Hermitian Hamiltonian operator $H = \sum_{i=1}^n h_i \sigma_i$ and the trial states parametrization


$$|\psi(\theta)\rangle = R_N(\theta_N) \cdots R_1(\theta_1)|\mathbf{0}\rangle, \quad (49)$$

where $R_i(\theta_i)$ are quantum gates parametrized by θ_i and applied one after the other to the initial state $|\mathbf{0}\rangle$. This parametrization corresponds to a variational quantum circuit. The average energy $\langle\psi(\theta)|H|\psi(\theta)\rangle$ can then be computed by individually measuring each term $\langle\psi(\theta)|\sigma_i|\psi(\theta)\rangle$ on a quantum computer and adding the results weighted with their respective coefficients h_i . Thereafter, the values of $\langle\psi(\theta)|H|\psi(\theta)\rangle$ are used by a suitable optimization method running on a classic computer.

In general, the VQE method uses SPSA as the optimization algorithm due to its robustness against noise, which suggests that the optimization methods presented here can also be used. In order to evaluate the performance of these various methods, we use as a testing ground the problem of finding the ground-state energy of the Heisenberg Hamiltonian, which models the magnetic interaction of a ferromagnetic lattice. This is given by the expression

$$H_H = j \sum_{\langle m,n \rangle} \sum_{k=x,y,z} \sigma_m^k \sigma_n^k + h \sum_m \sigma_m^z, \quad (50)$$

where j and h are dimensionless coupling constants between neighboring sites and with an external magnetic field, respectively, $\sigma_m^x, \sigma_m^y, \sigma_m^z$ are the single-qubit Pauli operators acting on the m th lattice site, and $\langle m, n \rangle$ indicates that the sum is performed on the nearest neighbors in the lattice. To parametrize the trial states we use the complex parametric single-qubit

FIG. 1. Entangling gate U_{ENT} .

gate

$$W(z) = e^{-i(z\sigma_+ + z^*\sigma_-)}, \quad (51)$$

where z is a complex parameter and $\sigma_{\pm} = \sigma^x \pm i\sigma^y$. This gate can be implemented experimentally by a sequence of three real parameter gates.

The parametrization used for the trial states is given by

$$|\psi(\mathbf{z})\rangle = \prod_{q=1}^N W_q^d(z_q^d) U_{\text{ENT}} \times \cdots \times \prod_{q=1}^N W_q^1(z_q^1) U_{\text{ENT}} \prod_{q=1}^N W_q^0(z_q^0) |\mathbf{0}\rangle, \quad (52)$$

where W_q^l corresponds to applying W on the q th qubit, l indicates the layer of the circuit, and U_{ENT} is the three-qubit entangling gate depicted in Fig. 1, and similarly for z_q^l .

To evaluate the performance of the different algorithms we consider the Heisenberg Hamiltonian (50) with $h = 0.3$ and $j = 1$ for a ring of ten qubits with periodic boundary conditions $q_{i+10} = q_i$. The trial states are parametrized by Eq. (52) with $d = 1$ entangling layers, as depicted by the circuit in Fig. 2. Each algorithm is simulated considering 10^2 randomly selected initial states according to a Haar distribution, which allows estimating statistical indicators such as mean, median, standard deviation, and interquartile range. The measurements required by each optimization method are simulated with an ensemble size of 2×10^4 . The standard gain coefficients used in first-order algorithms are $b = 0.1$, and a follows a calibration based on Ref. [63].

Figure 3 displays the best performance of the vanilla algorithms as a function of the number of iterations. All methods delivered the best results using postprocessing equations (45). This figure shows that the best performers are the real and complex first-order and the quantum natural complex algorithms, which exhibit an almost indistinguishable behavior in mean, median, standard deviation, and interquartile range. These algorithms converge to a minimum at approximately 2×10^2 iterations, after which they become approximately constant. The second-order algorithms exhibit a slower convergence, reaching a similar value only after 7×10^2 iterations.




FIG. 2. Parametric circuit used to implement the VQE for the Heisenberg Hamiltonian of three qubits. U_{ENT} is an entangling gate depicted in Fig. 1.

Figure 4 displays the best performance of the improved methods. In this case, the best performers are the first-order methods, the scalar version of second-order methods, and the scalar version of quantum natural methods. These exhibit an almost indistinguishable convergence in mean and median as well as similar dispersion. In particular, a minimum is achieved at approximately 10^2 iterations, after which the energy becomes nearly constant. 2CSPSA and 2SPSA scalar methods use postprocessing of Eqs. (33) and (45), respectively, and standard gains. QN-CSPSA and QN-SPSA scalar methods use postprocessing of Eqs. (33) and (45), respectively, and asymptotic gains. All best performers use resampling with $N_R = 5$.

From Figs. 3 and 4 we conclude that the best performance in the variational quantum eigensolver is achieved using SPSCA, CSPSA, QN-CSPSA, and QN-CSPSA scalar methods, which does not significantly differ in their vanilla or improved versions. Blocking and resampling lead to a clear improvement of the second-order methods, delivering results similar to the best performers.

First-order algorithms provide the best performance for this particular problem. Nonetheless, it is crucial to note that this level of performance was attained through a resource-intensive search for gain coefficients. In the absence of such a search, the first-order algorithms performed poorly. To bypass the calibration of the gain coefficients, quantum natural algorithms can be applied while achieving a performance close to the calibrated CSPSA. In particular, the scalar quantum natural CSPSA algorithm also reduces the classical computational cost.

B. Quantum control

Quantum control theory lays a firm theoretical foundation for developing a series of systematic methods that allow the manipulation and control of quantum systems. In particular, the search for an optimized time evolution that allows guiding the system from an initial state to a desired final state is of great interest. Quantum control theory has already achieved significant successes in physical chemistry [108], atomic and molecular physics [109], and quantum optics [110], and has also contributed to understanding fundamental aspects of quantum mechanics [111]. In recent years, the development of the general principles of quantum control theory has been recognized as an essential requirement for the current and future applications of quantum technologies.

A particular problem in quantum control is the precise engineering of quantum states, that is, whether a quantum system can be driven into a given state. This problem has practical importance since it is closely related to the universality of quantum computing and the possibility of achieving transformations at the atomic or molecular scale. An important research problem is that of finite-dimensional quantum systems, for which the controllability criteria can be expressed in terms of parameters included in the Hamiltonian of the system.

The quantum state control problem [112] consists in identifying an appropriate set of time-dependent control parameters $u_k(t)$ in such a way that its controlled change in time guides the evolution of the system from an initial state $|\psi_0\rangle$ to a pre-determined objective state $|\psi_f\rangle$. The control parameters enter




FIG. 3. The mean (top row) and median (bottom row) of the energy (in arbitrary units) as a function of the number of iterations obtained through the VQE for the Heisenberg Hamiltonian in a ten-qubit ring configuration using vanilla optimization algorithms. The shaded areas represent the variance (top row) and the interquartile range (bottom row). The dashed line indicates the exact minimum. The statistics are obtained from a sample of 10^2 randomly generated states to estimate the minimum energy. The measurements in each circuit were estimated with 2×10^4 shots. The values of the gain coefficients and postprocessing class can be found in Table I.

in the Hamiltonian as coefficients in a linear combination of operators, that is,

$$H(t) = H_0 + \frac{1}{2} \sum_k (u_k(t)C_k + u_k^*(t)C_k^\dagger), \quad (53)$$

where the set $\{C_k\}$ is a base of operators and we allow the possibility of complex control parameters. In order to obtain this set of parameters, it is necessary to solve the time-dependent Schrödinger equation. Unfortunately, solutions of the Schrödinger equation for a time-dependent Hamiltonian cannot generally be obtained analytically. However, it is possible in certain cases to use techniques developed in the area of adiabatic control [113–116].

To overcome this problem, we use the GRAPE method [95], originally introduced in nuclear magnetic resonance spectroscopy and proposed to design a pulse sequence that drives the evolution toward the optimum of a predefined

objective function. This method allows us to compute the evolution of a time-dependent Hamiltonian through a sequence H_m of time-independent Hamiltonians. The total evolution time T is divided into a number M of time intervals $\Delta t_m = t_{m+1} - t_m$ ($m = 0, \dots, M-1$), which are normally of equal length so that in each interval the control parameters $u_k(t)$ are approximately constant. In each time interval, the evolution is given by

$$U_m = e^{-i\Delta t_m H_m}, \quad (54)$$

where $H_m = H(t_m^*)$ with $t_m^* \in [t_m, t_{m+1}]$. A classical optimization algorithm is used to obtain the values of the control parameters that lead to the optimum of the objective function. The evolution of the system at time T is thus approximated by the sequence

$$U_{\text{GRAPE}} = \prod_{m=0}^{M-1} U_m \quad (55)$$

TABLE I. Best configuration and statistical indicators for each vanilla method applied to the variational quantum eigensolver.

Method	Gains	Postprocessing	Median	IQR	Mean	SD
SPSA	Standard	–	–6.48	1.11	–6.46	0.56
CSPA	Standard	–	–5.93	1.12	–6.44	0.57
2SPSA	Standard	Eqs. (45)	–5.76	1.17	–5.61	0.85
2CSPA	Standard	Eqs. (45)	–4.64	1.96	–4.82	1.19
scalar 2SPSA	Standard	Eqs. (45)	–5.84	1.08	–5.84	0.80
scalar 2CSPA	Standard	Eqs. (45)	–5.17	1.49	–5.01	1.08
QN-SPSA	Asymptotic	Eqs. (45)	–5.86	0.85	–6.13	0.48
QN-CSPA	Standard	Eqs. (45)	–5.92	1.10	–6.33	0.53
scalar QN-SPSA	Asymptotic	Eqs. (45)	–5.90	1.07	–6.34	0.53
scalar QN-CSPA	Standard	Eqs. (45)	–5.93	1.10	–6.39	0.55

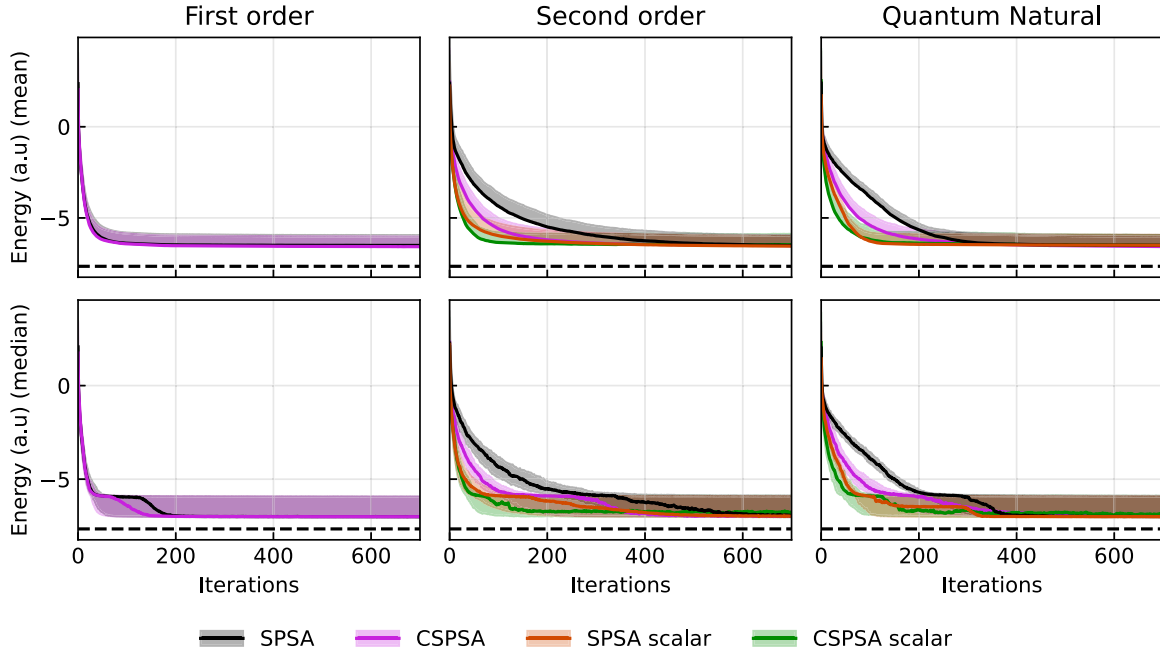


FIG. 4. The mean (top row) and median (bottom row) of the energy as a function of the number of iterations obtained through the VQE for the Heisenberg Hamiltonian in a ten-qubit ring configuration using improved optimization algorithms. The shaded areas represent the variance (top row) and the interquartile range (bottom row). The dashed line indicates the exact minimum. The statistics are obtained from a sample of 10^2 randomly generated states to estimate the minimum energy. The measurements in each circuit were estimated with 2×10^4 shots. The values of the gain coefficients, postprocessing class, and the setting of resampling and blocking can be found in Table II.

and the state of the system at time T is

$$|\tilde{\psi}_f\rangle = U_{\text{GRAPE}}|\psi_0\rangle. \quad (56)$$

Once a propagator has been computed for a set of control parameters, all that remains is to choose an objective function to compare the target state with the state given by the evolution for a given set of control parameter values. In our case, we use the infidelity that is given by

$$I(|\tilde{\psi}_f\rangle, |\psi_f\rangle) = 1 - |\langle\tilde{\psi}_f|\psi_f\rangle|^2, \quad (57)$$

which is minimized with an optimization algorithm. The original GRAPE proposal uses the descending gradient algorithm. The dimension of the search space is given by $N_p M$, where N_p is the number of parameters, and therefore can be very large.

To test the optimization methods introduced here, we turn to the quantum control of a five-qubit system, where we aim at preparing the target state $|\psi_f\rangle = |0\rangle^{\otimes 5}$ by controlling the

evolution generated by the Heisenberg Hamiltonian given by

$$H_H(t) = -\frac{1}{2} \sum_{k=x,y,z} J_k(t) \sum_{\langle m,n \rangle} \sigma_m^k \sigma_n^k, \quad (58)$$

which depends on the three complex coupling constants $J_x(t)$, $J_y(t)$, and $J_z(t)$. These play the role of control parameters whose values are driven by the quantum control method to approach the desired target state.

After applying the GRAPE method for the evolution of the system, the final state is

$$|\tilde{\psi}_f\rangle = \prod_{i=1}^{M-1} e^{-i\Delta t_m H_H(t_m^*)} |\psi_0\rangle, \quad (59)$$

where $|\psi_0\rangle$ is an initial five-qubit state and H_H contains the control parameters.

TABLE II. Best configuration and statistical indicators for each method with improvements applied to variational quantum eigensolver.

Method	Gains	Postprocessing	Resampling	Blocking	Median	IQR	Mean	SD
SPSA	Standard	—	5	No	−7.00	1.12	−6.58	0.56
CSPA	Standard	—	5	No	−7.00	1.12	−6.51	0.57
2SPSA	Standard	Eqs. (5)	2	Yes	−6.94	1.11	−6.51	0.56
2CSPA	Standard	Eqs. (33)	5	Yes	−6.98	1.13	−6.52	0.56
scalar 2SPSA	Standard	Eqs. (45)	5	Yes	−6.97	1.09	−6.55	0.30
scalar 2CSPA	Standard	Eqs. (33)	5	Yes	−6.78	1.15	−6.47	0.58
QN-SPSA	Asymptotic	Eqs. (45)	2	Yes	−7.00	1.12	−6.50	0.56
QN-CSPA	Asymptotic	Eqs. (33)	5	Yes	−6.98	1.11	−6.57	0.54
scalar QN-SPSA	Asymptotic	Eqs. (45)	5	Yes	−7.00	1.25	−6.50	0.57
scalar QN-CSPA	Asymptotic	Eqs. (33)	5	Yes	−6.88	1.13	−6.51	0.57

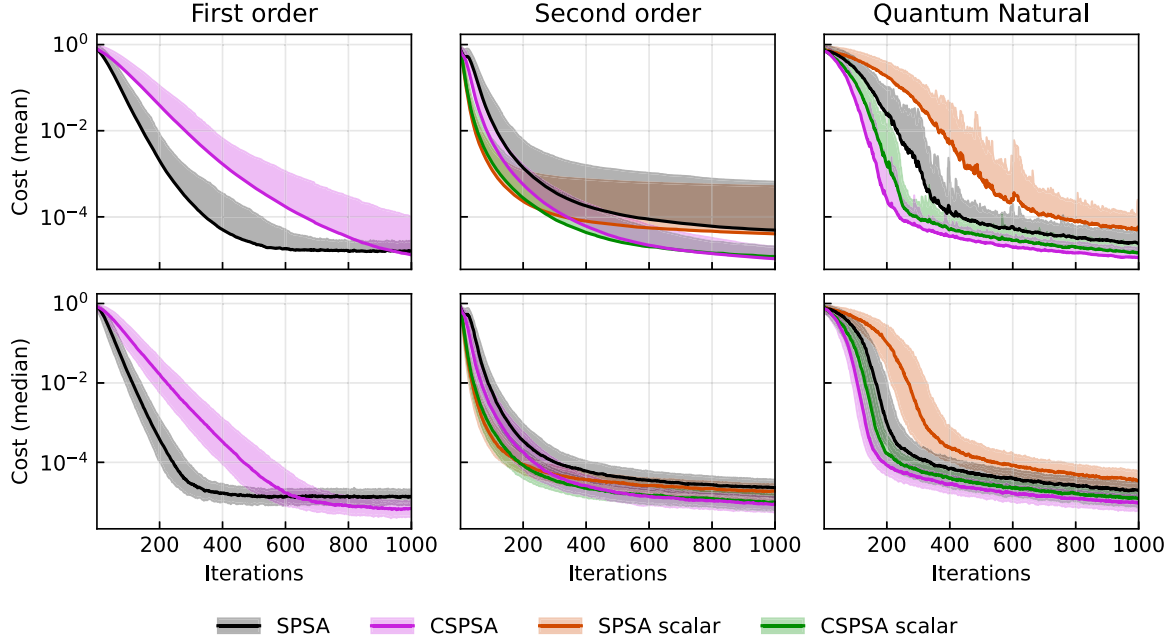


FIG. 5. The mean (top row) and median (bottom row) of the infidelity as a function of the number of iterations obtained through the GRAPE method applied to the quantum control of a five-qubit pure state and vanilla optimization algorithms. The shaded areas represent the variance (top row) and the interquartile range (bottom row). The values of the infidelity are obtained by simulating a measurement process with a sample size of 2^{13} and 25 iterations of GRAPE, 10^4 shots per measurement, and 10^3 iterations, which are generated through uniformly distributed initial states $|\psi_0\rangle$. The values of the gain coefficients and postprocessing class can be found in Table III.

Our performance study is based on numerical simulations where we implement GRAPE with each of the methods reviewed or proposed here. For a given optimization method, we start by choosing an initial state $|\psi_0\rangle$ from a Haar-uniform distribution on which we apply the GRAPE method with $M = 25$ and 10^3 iterations. Therefore, the dimension of the complex search space is 75, with the real search space being twice as large. The measurements required by the optimization method are simulated with an ensemble of size 2^{13} . This procedure is repeated 10^4 times to obtain estimates of relevant statistical indicators such as mean, median, standard deviation, and interquartile range, as functions of the number of iterations. The gain parameters used in the numerical simulations are shown in Tables III and IV.

The results of the numerical simulations of the GRAPE method with the different optimization algorithms in the five-

qubit case are depicted in Figs. 5 and 6, which show the best results among the vanilla methods and the improved methods, respectively. Each figure shows the value of the mean (upper row) and median (lower row) infidelity as a function of the number of iterations together with the variance (upper row) and the interquartile range (lower row) as shaded areas.

Figure 5 shows the comparison between methods without using blocking and resampling (see Sec. II C), that is, the vanilla methods. Second-order methods exhibit the best mean performance, particularly 2CSPA and scalar 2CSPA. These are closely followed by their quantum natural counterparts. First-order methods initially offer a better convergence rate but stagnate after a certain number of iterations. Let us note that this is the only case among all applications where first-order SPSA and CSPA achieve their best performance using the static gain coefficients. In the

TABLE III. Best configuration and statistical indicators for each vanilla method applied to quantum control of quantum states.

Method	Gains	Postprocessing	Median	IQR	Mean	SD
SPSA	Static	—	1.35×10^{-5}	1.21×10^{-5}	1.60×10^{-5}	1.11×10^{-5}
CSPA	Static	—	6.84×10^{-6}	6.01×10^{-6}	1.31×10^{-5}	8.23×10^{-5}
2SPSA	Standard	Eqs. (45)	2.31×10^{-5}	2.34×10^{-5}	4.90×10^{-5}	6.02×10^{-4}
2CSPA	Standard	Eqs. (45)	8.79×10^{-6}	8.12×10^{-6}	1.07×10^{-5}	9.76×10^{-6}
scalar 2SPSA	Standard	Eqs. (45)	1.89×10^{-5}	1.71×10^{-5}	4.03×10^{-5}	4.75×10^{-4}
scalar 2CSPA	Standard	Eqs. (45)	9.73×10^{-6}	8.93×10^{-6}	1.18×10^{-5}	8.32×10^{-6}
QN-SPSA	Asymptotic	Eqs. (45)	1.93×10^{-5}	1.88×10^{-5}	2.36×10^{-5}	1.82×10^{-5}
QN-CSPA	Asymptotic	Eqs. (45)	9.44×10^{-6}	8.49×10^{-6}	1.12×10^{-5}	7.50×10^{-6}
scalar QN-SPSA	Asymptotic	Eqs. (45)	3.44×10^{-5}	3.73×10^{-5}	5.68×10^{-5}	1.22×10^{-4}
scalar QN-CSPA	Asymptotic	Eqs. (45)	1.24×10^{-5}	1.11×10^{-5}	1.47×10^{-5}	1.04×10^{-5}




FIG. 6. The mean (top row) and median (bottom row) of the infidelity as a function of the number of iterations obtained through the GRAPE method applied to the quantum control of a five-qubit pure state and improved optimization algorithms. The shaded areas represent the variance (top row) and the interquartile range (bottom row). The values of the infidelity are obtained by simulating a measurement process with a sample size of 2^{13} and 25 iterations of GRAPE, 10^4 shots per measurement, and 10^3 iterations, which are generated through uniformly distributed initial states $|\psi_0\rangle$. The values of the gain coefficients, postprocessing class, and the setting of resampling and blocking can be found in Table IV.

median, second-order methods exhibit a higher convergence rate, closely followed by the quantum natural methods. While first-order methods require a larger number of iterations, a similar value of infidelity is reached in all cases. However, they behave differently in mean and median, in contrast to the complex second-order and quantum natural methods exhibiting consistent statistical indicators.

Figure 6 shows the comparison between methods when we allow the usage of blocking and resampling, that is, the improved methods. The best performance in the mean and the median is attained by the first-order methods seconded by the quantum natural methods, which exhibit a slightly slower rate of convergence with similar standard deviation and interquartile range. Second-order methods are the worst performers. These are characterized by a lower precision in mean, large standard deviation, and a slower rate of convergence, with

the exception of scalar methods. Figure 6 also indicates that complex methods perform better than their real counterparts.

Generally, the first-order CSPA method with resampling and blocking obtains the best result, using Eqs. (45) for postprocessing, closely followed by the QN-CSPA method.

C. Quantum state estimation

Born's rule endows quantum mechanics with predictive power. According to this rule, the probability p_k of obtaining a result k in an experiment described by a positive operator-valued measure (POVM) $\{E_k\}$ when the quantum system is described by a quantum state ρ is given by the Hilbert-Schmidt inner product $p_k = \text{Tr}(\rho E_k)$. Therefore, the comparison between the theoretical predictions and the experimental results requires an accurate characterization of the

TABLE IV. Best configuration and statistical indicators for each method with improvements applied to quantum control of quantum states.

Method	Gains	Postprocessing	Resampling	Blocking	Median	IQR	Mean	SD
SPSA	Asymptotic	—	5	No	2.02×10^{-6}	1.75×10^{-6}	2.35×10^{-6}	1.49×10^{-6}
CSPA	Asymptotic	—	5	No	9.24×10^{-7}	8.16×10^{-7}	1.07×10^{-6}	6.95×10^{-7}
2SPSA	Standard	Eqs. (45)	2	No	2.64×10^{-5}	2.64×10^{-5}	3.19×10^{-5}	2.24×10^{-5}
2CSPA	Standard	Eqs. (45)	2	No	9.15×10^{-6}	8.45×10^{-6}	1.36×10^{-5}	5.57×10^{-5}
scalar 2SPSA	Standard	Eqs. (45)	5	No	1.43×10^{-5}	1.28×10^{-5}	1.58×10^{-5}	9.12×10^{-6}
scalar 2CSPA	Standard	Eqs. (45)	5	No	8.84×10^{-6}	7.68×10^{-6}	9.85×10^{-6}	5.84×10^{-6}
QN-SPSA	Asymptotic	Eqs. (45)	5	Yes	3.46×10^{-6}	2.61×10^{-6}	3.85×10^{-6}	2.22×10^{-6}
QN-CSPA	Asymptotic	Eqs. (45)	5	No	1.69×10^{-6}	1.43×10^{-6}	1.92×10^{-6}	1.15×10^{-6}
scalar QN-SPSA	Asymptotic	Eqs. (45)	5	Yes	5.92×10^{-6}	5.01×10^{-6}	6.85×10^{-6}	4.44×10^{-6}
scalar QN-CSPA	Asymptotic	Eqs. (45)	5	Yes	2.80×10^{-6}	2.24×10^{-6}	3.20×10^{-6}	1.95×10^{-6}




FIG. 7. The mean (top row) and median (bottom row) of the infidelity as a function of the number of iterations obtained by using SGQT to estimate six-qubit states and vanilla optimization algorithms. Shaded areas represent variance (top row) and interquartile range (bottom row). Statistical indicators are calculated from a sample of 10^2 Haar-uniform distributed pairs of unknown and initial guess states. Measurements of the infidelity are simulated with a binomial distribution with $N = 2 \times 10^4$ shots. The values of the gain coefficients and postprocessing class can be found in Table V.

quantum state ρ and of the experiment through the POVM $\{E_k\}$. This leads to the problem of estimating quantum states and processes. To do this, several quantum state estimation methods have been designed, most of them based on the postprocessing of experimental data acquired through the measurement of a fixed informationally complete POVM. Adaptive measurements have also been used to design quantum state estimation methods. Today, methods for estimating quantum states are an important tool for both quantum communications and quantum computing and have been used for the characterization of single-photon and continuous variable states [117–121], cavity fields [122], atomic ensembles [123–125], trapped ions [126,127], optical detectors [128–130], and for quantum key distribution [131].

Recently, the estimation of finite-dimensional pure unknown states has been formulated as an optimization problem [17]. According to this, the unknown state is characterized

as the minimizer of infidelity $I(|\psi\rangle, |\phi\rangle) = 1 - |\langle\psi|\phi\rangle|^2$, that is,

$$|\psi\rangle = \arg \left(\min_{|\phi\rangle \in \mathcal{H}} I(|\psi\rangle, |\phi\rangle) \right). \quad (60)$$

This suggests using optimization algorithms to minimize fidelity and estimate the unknown state $|\psi\rangle$, which has been called self-guided quantum tomography (SGQT). Gradient-based optimization is ruled out since it is not known how to measure the infidelity gradient with respect to the parameters entering the $|\phi\rangle$ state. However, infidelity can be measured by projecting the unknown state onto any basis containing the state $|\phi\rangle$. In this scenario, the optimization methods presented in the previous section can be used to experimentally implement the infidelity minimization according to SGQT. Initially, SGQT was based on the SPSA algorithm. Subsequently, CSPA was introduced in SGQT, obtaining an improvement

TABLE V. Best configuration and statistical indicators for each vanilla method applied to self-guided quantum tomography.

Method	Gains	Postprocessing	Median	IQR	Mean	SD
SPSA	Asymptotic	—	4.76×10^{-4}	6.48×10^{-5}	4.79×10^{-4}	5.50×10^{-5}
CSPA	Asymptotic	—	1.01×10^{-4}	2.00×10^{-5}	1.03×10^{-4}	1.40×10^{-5}
2SPSA	Standard	Eqs. (45)	6.17×10^{-4}	3.52×10^{-3}	3.55×10^{-3}	4.75×10^{-4}
2CSPA	Standard	Eqs. (45)	1.43×10^{-4}	8.15×10^{-4}	8.15×10^{-4}	1.02×10^{-4}
scalar 2SPSA	Standard	Eqs. (45)	5.22×10^{-4}	3.26×10^{-3}	3.29×10^{-3}	3.70×10^{-4}
scalar 2CSPA	Standard	Eqs. (45)	1.38×10^{-4}	7.60×10^{-4}	7.58×10^{-4}	9.87×10^{-5}
QN-SPSA	Standard	Eqs. (45)	6.68×10^{-3}	4.04×10^{-3}	6.72×10^{-3}	8.42×10^{-4}
QN-CSPA	Standard	Eqs. (45)	1.52×10^{-3}	9.43×10^{-4}	1.53×10^{-3}	1.93×10^{-4}
scalar QN-SPSA	Standard	Eqs. (45)	6.58×10^{-3}	4.00×10^{-3}	6.55×10^{-3}	8.29×10^{-4}
scalar QN-CSPA	Standard	Eqs. (45)	1.49×10^{-3}	9.69×10^{-4}	1.51×10^{-3}	1.94×10^{-4}




FIG. 8. The mean (top row) and median (bottom row) of the infidelity as a function of the number of iterations obtained by using SGQT to estimate six-qubit states and improved optimization algorithms. Shaded areas represent variance (top row) and interquartile range (bottom row). Statistical indicators are calculated from a sample of 10^2 Haar-uniform distributed pairs of unknown and initial guess states. Measurements of the infidelity are simulated with a binomial distribution with $N = 2 \times 10^4$ shots. The values of the gain coefficients, postprocessing class, and the setting of resampling and blocking can be found in Table VI.

in the rate of convergence and a lower dispersion in the sample of estimates. More recently, CSPA was combined with maximum likelihood estimation to achieve precision close to the Gill-Massar lower limit, which is the best achievable estimation accuracy for pure states. Estimating pure states through SPSA and CSPA has already been experimentally demonstrated [18].

We use pure-state estimation through SGQT to test the performance of the optimization methods proposed in the previous sections. After selecting a particular optimization method, we generate an unknown six-qubit pure state and an initial guess state from a Haar-uniform distribution. The optimization method is iterated 5×10^3 times and the fidelity values are obtained by simulating a measurement with binomial distribution on an ensemble size $N = 2 \times 10^4$. This procedure is repeated 10^2 times to generate estimates of relevant statistic indicators. The gain parameters used in the

numerical simulations of each method are shown in the Appendix.

The results of the numerical simulations of SGQT with the different optimization algorithms are depicted in Figs. 7 and 8 that show the best results among the vanilla and improved methods, respectively. Each figure shows the value of the mean (upper row) and median (lower row) infidelity as a function of the number of iterations together with the variance (upper row) and the interquartile range (lower row) as shaded areas. In every figure, the first column contains the results of SPSA and CSPA. The second column contains the results obtained by the second-order algorithms, that is, 2SPSA, 2CSPA, scalar 2SPSA, and scalar 2CSPA. The third column contains the results obtained by the algorithms based on the quantum natural method, that is, QN-SPSA, QN-CSPA, scalar QN-SPSA, and scalar QN-CSPA.

TABLE VI. Best configuration and statistical indicators for each method with improvements applied to self-guided quantum tomography.

Method	Gains	Postprocessing	Resampling	Blocking	Median	IQR	Mean	SD
SPSA	Asymptotic	–	1	No	4.76×10^{-4}	6.48×10^{-5}	4.79×10^{-4}	5.50×10^{-5}
CSPA	Asymptotic	–	1	No	1.01×10^{-4}	2.00×10^{-5}	1.03×10^{-4}	1.40×10^{-5}
2SPSA	Standard	Eqs. (45)	5	No	6.45×10^{-4}	1.01×10^{-4}	6.39×10^{-4}	7.26×10^{-5}
2CSPA	Standard	Eqs. (45)	5	No	1.52×10^{-4}	3.12×10^{-5}	1.54×10^{-4}	1.90×10^{-5}
scalar 2SPSA	Standard	Eqs. (45)	5	No	1.27×10^{-3}	2.21×10^{-4}	1.27×10^{-3}	1.67×10^{-4}
scalar 2CSPA	Standard	Eqs. (45)	5	No	3.05×10^{-4}	5.05×10^{-5}	3.06×10^{-4}	4.04×10^{-5}
QN-SPSA	Standard	Eqs. (45)	5	No	1.17×10^{-3}	2.06×10^{-4}	1.19×10^{-3}	1.45×10^{-4}
QN-CSPA	Standard	Eqs. (45)	5	No	2.74×10^{-4}	4.53×10^{-5}	2.76×10^{-4}	3.33×10^{-5}
scalar QN-SPSA	Standard	Eqs. (45)	5	No	2.57×10^{-3}	4.68×10^{-4}	2.55×10^{-3}	3.45×10^{-4}
scalar QN-CSPA	Standard	Eqs. (45)	5	No	6.06×10^{-4}	1.15×10^{-4}	6.09×10^{-4}	7.60×10^{-5}

In Figs. 7 and 8 mean and median values for each algorithm are very close. Furthermore, the variance and interquartile range are very narrow, which shows the absence of outliers in the generated samples. Typically, all optimization algorithms are characterized by a sharp decrease in infidelity followed by an approximately linear asymptotic regime.

Figure 7 shows the comparison between every method without using the blocking and resampling improvements in Eqs. (45), as these methods largely increase the number of resources. The first-order methods offer the best performance, getting about an order-of-magnitude improvement over the other methods. In contrast, the second-order methods perform slightly better than their QN counterpart. The scalar approximation shows no improvements for the second-order and quantum natural methods. The complex methods show better convergence than their real counterparts by about an order of magnitude.

Figure 8 shows the comparison between methods when we allow the use of blocking and resampling. For first-order algorithms, gradient blocking and resampling show no improvement, while second-order and QN methods improve when the Hessian approximation is averaged five times per iteration. This improvement decreases the performance difference between the first-order and the other methods. Second-order methods still perform slightly better than QN methods. We also note that the improvement obtained by resampling is smaller for the scalar approximation.

In our simulations, blocking does not improve our results when considering the Hessian postprocessing equations (45). On the other hand, when considering postprocessing equations (5), the blocking shows great improvement which matches our median results, but with worse mean performance and with wider data variability.

From Figs. 7 and 8 we conclude that the first-order methods show better mean and median performance in the estimation of pure six-qubit states via SGQT, even without considering gradient resampling. In this scenario, second-order methods are not expected to work properly since the fidelity Hessian vanishes for pure states. This issue could lower both precision and convergence speed. However, the Hessian postprocessing allows us to mitigate this problem by perturbing the Hessian matrix with a weighted identity matrix. In this way, the best result achieved by the second-order methods uses postprocessing equations (45). Quantum-natural-based methods show similar behavior, albeit with slightly slower convergence.

First-order methods perform the best even without considering gradient resampling. In contrast, second-order and quantum natural methods need resampling improvement to stay competitive but require a much higher number of resources.

IV. CONCLUSIONS

In this article, we have exhaustively compared different stochastic optimization methods applied to real-valued functions of complex variables. We started by reviewing the theory of the SPSA algorithm and two of its variants: 2SPSA and QN-SPSA. These three methods use a simultaneous perturbation stochastic approximation of the gradient of the objective function to optimize it. SPSA is a first-order algorithm, while

2SPSA is a second-order algorithm. QN-SPSA is a first-order algorithm that preconditions considering a metric natural for the problem at hand. We also reviewed the CSPSA algorithm, which optimizes real functions of complex variables without resorting to the real and imaginary parts of complex variables. This is a more natural approach in quantum mechanics, where most functions have complex arguments. Using CSPSA as starting point, we proposed two optimization methods, 2CSPSA and QN-CSPSA, which are the complex field formulations of their real counterparts.

All the optimization methods presented here share the property that the number of evaluations (or measurements) of the objective function does not depend on the dimension of the optimization problem. This is an important advantage when the number of parameters on which the objective function depends is large. The number of objective function evaluations is constant at each iteration but different for each method. SPSA and CSPSA use two evaluations of the objective function per iteration. 2SPSA and 2CSPSA use four evaluations of the objective function since they are second-order methods. Finally, QN-SPSA and QN-CSPSA use two evaluations of the objective function plus the calculation of an approximation of a metric. If the metric is the Fubini-Study metric tensor, then the approximation is calculated by evaluating the fidelity with respect to four different pure states.

To assess the performance of the optimization methods, we have compared them in three important applications in quantum computing: a variational quantum eigensolver applied to the Heisenberg Hamiltonian of a ten-qubit ring, quantum control applied to a five-qubit pure quantum state, and quantum state estimation to reconstruct a six-qubit pure quantum state. These three applications have different objective functions that need to be measured in a quantum device and iteratively optimized to obtain a solution. In particular, we have compared the convergence rate as a function of the number of iterations. To do this, we have considered vanilla and improved algorithm versions.

Our simulations show several interesting results. The best performance is systematically achieved by the first-order CSPSA algorithm. In the case of the variational quantum eigensolver, improved first-order CSPSA and SPSA algorithms provide the best performance, exhibiting identical mean and median and similar standard deviation and interquartile range. In quantum control, improved CSPSA achieves better convergence in mean and median than all other algorithms, exhibiting a narrow standard deviation and interquartile range. This is also the case for state estimation, although in this case, the vanilla version of the CSPSA algorithm is almost indistinguishable from its improved version.

The second-best overall performance has mixed results. In the variational quantum eigensolver, the improved second-order and improved quantum natural algorithms lead to an almost indistinguishable performance, while in quantum control the improved quantum natural algorithms, particularly improved QN-CSPSA, are clearly second best. In this case, scalar second-order algorithms perform, in mean and median, similarly to quantum natural algorithms. In contrast, their nonscalar counterparts show much lower mean performance, indicating the presence of a large number of outliers. In the case of quantum state estimation, improved second-order

algorithms provide better performance than their quantum natural counterparts. In particular, second-order CSPSA achieves the second-best performance.

Generally, vanilla second-order algorithms lead to lower performance than vanilla quantum natural algorithms. This is mitigated by blocking and resampling in the improved versions of second-order algorithms, which offer performance close to that of improved quantum natural algorithms. Furthermore, complex algorithms perform better than their real counterparts, although the difference may be statistically insignificant in certain cases.

While optimizing a function, it may be possible that no information about the Hessian matrix is available *a priori*, either because of its high complexity or because it cannot be easily obtained analytically or numerically. For such cases, it would be desirable that second-order methods, which are based on the Hessian matrix approximation, would still be useful in the event that the Hessian matrix exhibits singularities. This is the case of quantum state estimation, where the Hessian vanishes identically. Nevertheless, second-order methods display a performance similar to first-order methods. Hessian postprocessing equations (45) ensure that the preconditioning matrix is proportional to the identity for a vanishing Hessian matrix. This leads to second-order methods working like first-order methods, albeit possibly with suboptimal gain coefficients.

The stochastic optimization methods studied here are defined through a set of gain parameters whose values specify the gain coefficients. These in turn control the step size and magnitude of the approximation of the gradient. In this way, the gain parameters are hyperparameters that allow controlling the algorithms' convergence rate. In principle, it is conceivable to find gain parameters that lead to the best convergence rate. This is, however, an expensive optimization problem whose solution might even depend on the optimizer of the objective function. Therefore, it is usual to resort to gain parameters that have proven to be good enough in practice. We have resorted to the standard gain parameters, which lead to a fast convergence in the regime of a small number of iterations, and to the asymptotic gain parameters, which lead to a fast convergence in the regime of a large number of iterations. Let us note that a change in the gain parameters affects not only the mean and median convergence but also the variance and interquartile range. We have also performed our simulations considering static gain coefficients, which only led to a significant improvement in the case of vanilla first-order methods applied to quantum control.

From numerical simulations with fewer qubits, we observed that the performance difference between quantum natural and first-order algorithms tends to narrow as the number of qubits increases. For the simulations reported here, the performance difference among these algorithms is small. This may indicate that quantum natural methods may outperform

first-order methods for a larger number of qubits. However, this advantage of quantum natural methods is obtained by increasing the number of measurements and the classical computational cost. In this scene, the scalar quantum natural methods proposed here might be a good alternative since, according to our results, they offer comparable performance at a reduced classical cost.

According to the applications considered here, vanilla first-order algorithms are efficient and reliable options for the most general case. If higher accuracy is needed, improved first-order algorithms are the straight choice. First-order methods may require careful calibration of the gain parameters, in which case the quantum natural algorithms are a suitable alternative. In addition, quantum natural algorithms show promising results for many qubits, while second-order algorithms do not exhibit a comparative advantage.

In our study of first- and second-order algorithms, we have considered a single source of noise, namely, the statistical character of quantum measurements. It is possible to consider other error sources, such as those affecting NISQ processors. However, first-order algorithms, real or complex, have convergence proofs that allow for certain types of errors affecting the evaluation of the target function. Thereby, it is expected that these algorithms will converge even in the presence of moderate noise, albeit with an increased number of iterations. The scenario in the case of the preconditioned algorithms is less clear due to the inversion of the approximated Hessian matrix. Therefore, a natural extension of this work would be to consider realistic noise sources and their impact on the convergence rate. Also, we have considered the performance as a function of the number of iterations. It is possible, however, to consider other valuable resources such as the number of measurements, evaluations, and circuits. These should also be considered in further studies of the real performance of optimization algorithms.

ACKNOWLEDGMENTS

This work was supported by ANID–Millennium Science Initiative Program ICN17_012. J.G. was supported by ANID Chile, National Doctoral Degree Scholarship No. 21202616. L.P. was supported by ANID-PFCHA/DOCTORADO-BECAS-CHILE/2019-772200275, the CSIC Interdisciplinary Thematic Platform (PTI+) on Quantum Technologies (PTI-QTEP+), and the Proyecto Sinérgico CAM 2020 Y2020/TCS-6545 (NanoQuCo-CM). L.Z. was supported by ANID-PFCHA/DOCTORADO-NACIONAL/2018-21181021, the Government of Spain (Severo Ochoa CEX2019-000910-S, TRANQI and European Union NextGenerationEU PRTR-C17.I1), Fundació Cellex, Fundació Mir-Puig and Generalitat de Catalunya (CERCA program). A.D. was supported by FONDECYT Grants No. 1231940 and No. 1230586.

-
- [1] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
 - [2] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn *et al.*, Quantum optimization using variational algorithms on

near-term quantum devices, *Quantum Sci. Technol.* **3**, 030503 (2018).

- [3] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).

- [4] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann *et al.*, Noisy intermediate-scale quantum algorithms, *Rev. Mod. Phys.* **94**, 015004 (2022).
- [5] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and J. Tennyson, The variational quantum eigensolver: A review of methods and best practices, *Phys. Rep.* **986**, 1 (2022).
- [6] B. P. Lanyon, J. D. Whitfield, G. G. Gillett, M. E. Goggin, M. P. Almeida, I. Kassal, J. D. Biamonte, M. Mohseni, B. J. Powell, M. Barbieri *et al.*, Towards quantum chemistry on a quantum computer, *Nat. Chem.* **2**, 106 (2010).
- [7] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush *et al.*, Quantum Chemistry Calculations on a Trapped-Ion Quantum Simulator, *Phys. Rev. X* **8**, 031022 (2018).
- [8] Y. Nam, J.-S. Chen, N. C. Pienti, K. Wright, C. Delaney, D. Maslov, K. R. Brown, S. Allen, J. M. Amini, J. Apisdorf *et al.*, Ground-state energy estimation of the water molecule on a trapped-ion quantum computer, *npj Quantum Inf.* **6**, 33 (2020).
- [9] C. Ferrie and O. Moussa, Robust and efficient *in situ* quantum control, *Phys. Rev. A* **91**, 052306 (2015).
- [10] D. Lu, K. Li, J. Li, H. Katiyar, J. Park, G. Feng, T. Xin, H. Li, G. Long, A. Brodutch *et al.*, Enhancing quantum control by bootstrapping a quantum processor of 12 qubits, *npj Quantum Inf.* **3**, 45 (2017).
- [11] D. J. Egger and F. K. Wilhelm, Adaptive Hybrid Optimal Quantum Control for Imprecisely Characterized Systems, *Phys. Rev. Lett.* **112**, 240503 (2014).
- [12] X. Yuan, S. Endo, Q. Zhao, Y. Li, and S. C. Benjamin, Theory of variational quantum simulation, *Quantum* **3**, 191 (2019).
- [13] S. Endo, J. Sun, Y. Li, S. C. Benjamin, and X. Yuan, Variational Quantum Simulation of General Processes, *Phys. Rev. Lett.* **125**, 010501 (2020).
- [14] X. Wang, Z. Song, and Y. Wang, Variational quantum singular value decomposition, *Quantum* **5**, 483 (2021).
- [15] K. Wang, Z. Song, X. Zhao, Z. Wang, and X. Wang, Detecting and quantifying entanglement on near-term quantum devices, *npj Quantum Inf.* **8**, 52 (2022).
- [16] A. D. Muñoz-Moller, L. Pereira, L. Zambrano, J. Cortés-Vega, and A. Delgado, Variational Determination of Multiqubit Geometrical Entanglement in Noisy Intermediate-Scale Quantum Computers, *Phys. Rev. Appl.* **18**, 024048 (2022).
- [17] C. Ferrie, Self-Guided Quantum Tomography, *Phys. Rev. Lett.* **113**, 190404 (2014).
- [18] R. J. Chapman, C. Ferrie, and A. Peruzzo, Experimental Demonstration of Self-Guided Quantum Tomography, *Phys. Rev. Lett.* **117**, 040402 (2016).
- [19] A. Utreras-Alarcón, M. Rivera-Tapia, S. Niklitschek, and A. Delgado, Stochastic optimization on complex variables and pure-state quantum tomography, *Sci. Rep.* **9**, 16143 (2019).
- [20] L. Zambrano, L. Pereira, S. Niklitschek, and A. Delgado, Estimation of pure quantum states in high dimension at the limit of quantum accuracy through complex optimization and statistical inference, *Sci. Rep.* **10**, 12781 (2020).
- [21] M. Rambach, M. Qaryan, M. Kewming, C. Ferrie, A. G. White, and J. Romero, Robust and Efficient High-Dimensional Quantum State Tomography, *Phys. Rev. Lett.* **126**, 100402 (2021).
- [22] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature (London)* **549**, 195 (2017).
- [23] M. Benedetti, E. Grant, L. Wossnig, and S. Severini, Adversarial quantum circuit learning for pure state approximation, *New J. Phys.* **21**, 043023 (2019).
- [24] A. Patterson, H. Chen, L. Wossnig, S. Severini, D. Browne, and I. Rungger, Quantum state discrimination using noisy quantum neural networks, *Phys. Rev. Res.* **3**, 013063 (2021).
- [25] H. Chen, L. Wossnig, S. Severini, H. Neven, and M. Mohseni, Universal discriminative quantum neural networks, *Quantum Mach. Intell.* **3**, 1 (2021).
- [26] D. Concha, L. Pereira, L. Zambrano, and A. Delgado, Training a quantum measurement device to discriminate unknown non-orthogonal quantum states, *Sci. Rep.* **13**, 7460 (2023).
- [27] X. Xu, S. C. Benjamin, and X. Yuan, Variational Circuit Compiler for Quantum Error Correction, *Phys. Rev. Appl.* **15**, 034068 (2021).
- [28] E. Farhi, J. Goldstone, S. Gutmann, and L. Zhou, The quantum approximate optimization algorithm and the Sherrington-Kirkpatrick model at infinite size, *Quantum* **6**, 759 (2022).
- [29] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, Quantum Approximate Optimization Algorithm: Performance, Mechanism, and Implementation on Near-Term Devices, *Phys. Rev. X* **10**, 021067 (2020).
- [30] M. P. Harrigan, K. J. Sung, M. Neeley, K. J. Satzinger, F. Arute, K. Arya, J. Atalaya, J. C. Bardin, R. Barends, S. Boixo *et al.*, Quantum approximate optimization of non-planar graph problems on a planar superconducting processor, *Nat. Phys.* **17**, 332 (2021).
- [31] K. Kubo, Y. O. Nakagawa, S. Endo, and S. Nagayama, Variational quantum simulations of stochastic differential equations, *Phys. Rev. A* **103**, 052425 (2021).
- [32] M. Lubasch, J. Joo, P. Moinier, M. Kiffner, and D. Jaksch, Variational quantum algorithms for nonlinear problems, *Phys. Rev. A* **101**, 010301(R) (2020).
- [33] P. García-Molina, J. Rodríguez-Mediavilla, and J. J. García-Ripoll, Quantum Fourier analysis for multivariate functions and applications to a class of Schrödinger-type partial differential equations, *Phys. Rev. A* **105**, 012433 (2022).
- [34] D. Herman, C. Googin, X. Liu, A. Galda, I. Safro, Y. Sun, M. Pistoia, and Y. Alexeev, A survey of quantum computing for finance, *arXiv:2201.02773*.
- [35] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
- [36] A. Arrasmith, M. Cerezo, P. Czarnik, L. Cincio, and P. J. Coles, Effect of barren plateaus on gradient-free optimization, *Quantum* **5**, 558 (2021).
- [37] K. J. Sung, J. Yao, M. P. Harrigan, N. C. Rubin, Z. Jiang, L. Lin, R. Babbush, and J. R. McClean, Using models to improve optimizers for variational quantum algorithms, *Quantum Sci. Technol.* **5**, 044008 (2020).
- [38] O. Lockwood, An empirical review of optimization techniques for quantum variational circuits, *arXiv:2202.01389*.
- [39] W. Lavrijsen, A. Tudor, J. Muller, C. Iancu, and W. de Jong, Classical optimizers for noisy intermediate-scale quantum devices, in *2020 IEEE International Conference on Quantum*

- Computing and Engineering (QCE)* (IEEE, New York, 2020), pp. 267–277.
- [40] X. Bonet-Monroig, H. Wang, D. Vermetten, B. Senjean, C. Moussa, T. Bäck, V. Dunjko, and T. E. O'Brien, Performance comparison of optimization methods on variational quantum algorithms, *Phys. Rev. A* **107**, 032407 (2023).
 - [41] D. Saad, ed., *On-Line Learning in Neural Networks* (Cambridge University Press, Cambridge, UK, 1999).
 - [42] J. Duchi, E. Hazan, and Y. Singer, Adaptive subgradient methods for online learning and stochastic optimization, *J. Mach. Learn. Res.* **12**, 2121 (2011).
 - [43] T. Tieleman and G. Hinton, Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude, Coursera: Neural networks for machine learning **4**, 26 (2012).
 - [44] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
 - [45] T. Dozat, Incorporating Nesterov momentum into Adam, in *Proceedings of the 4th International Conference on Learning Representations* (Stanford University, 2016), pp. 1–4.
 - [46] S. J. Reddi, S. Kale, and S. Kumar, On the convergence of Adam and beyond, [arXiv:1904.09237](https://arxiv.org/abs/1904.09237).
 - [47] J. A. Nelder and R. Mead, A simplex method for function minimization, *Comput. J.* **7**, 308 (1965).
 - [48] M. J. D. Powell, An efficient method for finding the minimum of a function of several variables without calculating derivatives, *Comput. J.* **7**, 155 (1964).
 - [49] S. G. Nash, Newton-type minimization via the Lanczos method, *SIAM J. Numer. Anal.* **21**, 770 (1984).
 - [50] D. Kraft, *A software package for sequential quadratic programming*, Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt Köln: Forschungsbericht (Wiss. Berichtswesen d. DFVLR, 1988).
 - [51] R. Fletcher and C. M. Reeves, Function minimization by conjugate gradients, *Comput. J.* **7**, 149 (1964).
 - [52] M. J. D. Powell, A direct search optimization method that models the objective and constraint functions by linear interpolation, in *Advances in Optimization and Numerical Analysis* (Springer Netherlands, Dordrecht, 1994), pp. 51–67.
 - [53] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, A limited memory algorithm for bound constrained optimization, *SIAM J. Numer. Anal.* **16**, 1190 (1995).
 - [54] M. Lalee, J. Nocedal, and T. Plantenga, On the implementation of an algorithm for large-scale equality constrained optimization, *SIAM J. Numer. Anal.* **8**, 682 (1998).
 - [55] J. Nocedal and S. Wright, *Numerical Optimization* (Springer, New York, 2006).
 - [56] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in python, *Nat. Methods* **17**, 261 (2020).
 - [57] H. J. Kushner and D. S. Clark, *Stochastic Approximation Methods for Constrained and Unconstrained Systems* (Springer, New York, 1978), Vol. 26.
 - [58] H. J. Kushner and G. G. Yin, *Stochastic Approximation Algorithms and Applications* (Springer, New York, 1997).
 - [59] A. E. Albert and L. A. Gardner, *Stochastic Approximation and Non-Linear Regression* (MIT Press, Cambridge, MA, 2003).
 - [60] J. Spall, Introduction to stochastic search and optimization: Estimation, simulation, and control, *IEEE Trans. Neural Netw. Learn. Syst.* **18** (2007).
 - [61] S. Bhatnagar, H. L. Prasad, and L. A. Prashanth, *Stochastic Recursive Algorithms for Optimization* (Springer, London, 2013), Vol. 434.
 - [62] J. C. Spall, A stochastic approximation technique for generating maximum likelihood parameter estimates, in *Proceedings of the 1987 American Control Conference*, Minneapolis (IEEE, New York, 1987), pp. 1161–1167.
 - [63] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature (London)* **549**, 242 (2017).
 - [64] O. V. Borzenkova, G. I. Struchalin, A. S. Kardashin, V. V. Krasnikov, N. N. Skryabin, S. S. Straupe, S. P. Kulik, and J. D. Biamonte, Variational simulation of Schwinger's Hamiltonian with polarization qubits, *Appl. Phys. Lett.* **118**, 144002 (2021).
 - [65] I. Hamamura and T. Imamichi, Efficient evaluation of quantum observables using entangled measurements, *npj Quantum Inf.* **6**, 56 (2020).
 - [66] F. Escudero, D. Fernández-Fernández, G. Jaumà, G. F. Peñas, and L. Pereira, Hardware-efficient entangled measurements for variational quantum algorithms, [arXiv:2202.06979](https://arxiv.org/abs/2202.06979).
 - [67] P. Díez-Valle, D. Porras, and J. J. García-Ripoll, Quantum variational optimization: The role of entanglement and problem hardness, *Phys. Rev. A* **104**, 062426 (2021).
 - [68] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parameterized quantum circuits as machine learning models, *Quantum Sci. Technol.* **4**, 043001 (2019).
 - [69] S. Mangini, F. Tacchino, D. Gerace, C. Macchiavello, and D. Bajoni, Quantum computing model of an artificial neuron with continuously valued input data, *Mach. Learn.: Sci. Technol.* **1**, 045008 (2020).
 - [70] G. Agliardi and E. Prati, Optimal tuning of quantum generative adversarial networks for multivariate distribution loading, *Quantum Rep.* **4**, 75 (2022).
 - [71] Z. Hou, J.-F. Tang, C. Ferrie, G.-Y. Xiang, C.-F. Li, and G.-C. Guo, Experimental realization of self-guided quantum process tomography, *Phys. Rev. A* **101**, 022317 (2020).
 - [72] J. Spall, Adaptive stochastic approximation by the simultaneous perturbation method, *IEEE Trans. Autom. Control* **45**, 1839 (2000).
 - [73] C. Wang, An overview of SPSA: Recent development and applications, [arXiv:2012.06952](https://arxiv.org/abs/2012.06952).
 - [74] X. Zhu and J. C. Spall, A modified second-order SPSA optimization algorithm for finite samples, *Int. J. Adapt. Control Signal Process.* **16**, 397 (2002).
 - [75] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, Quantum natural gradient, *Quantum* **4**, 269 (2020).
 - [76] S. Amari, Natural gradient works efficiently in learning, *Neural Comput.* **10**, 251 (1998).
 - [77] T. Liang, T. Poggio, A. Rakhlin, and J. Stokes, Fisher-Rao metric, geometry, and complexity of neural networks, in *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research, Vol. 89, edited by K. Chaudhuri and M. Sugiyama (Machine Learning Research Press, Cambridge, MA, 2019), pp. 888–896.

- [78] J. Gacon, C. Zoufal, G. Carleo, and S. Woerner, Simultaneous perturbation stochastic approximation of the quantum fisher information, *Quantum* **5**, 567 (2021).
- [79] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf, Quantum Fingerprinting, *Phys. Rev. Lett.* **87**, 167902 (2001).
- [80] Q. Wang, Z. Zhang, K. Chen, J. Guan, W. Fang, J. Liu, and M. Ying, Quantum algorithm for fidelity estimation, *IEEE Trans. Inf. Theory* **69**, 273 (2023).
- [81] P. Gokhale, O. Angiuli, Y. Ding, K. Gui, T. Tomesh, M. Suchara, M. Martonosi, and F. T. Chong, Minimizing state preparations in variational quantum eigensolver by partitioning into commuting families, [arXiv:1907.13623](https://arxiv.org/abs/1907.13623).
- [82] S. Bravyi, J. M. Gambetta, A. Mezzacapo, and K. Temme, Tapering off qubits to simulate fermionic Hamiltonians, [arXiv:1701.08213](https://arxiv.org/abs/1701.08213).
- [83] N. Yamamoto, On the natural gradient for variational quantum eigensolver, [arXiv:1909.05074](https://arxiv.org/abs/1909.05074).
- [84] D. Wierichs, C. Gogolin, and M. Kastoryano, Avoiding local minima in variational quantum eigensolvers with the natural gradient optimizer, *Phys. Rev. Res.* **2**, 043246 (2020).
- [85] B. van Straaten and B. Koczor, Measurement cost of metric-aware variational quantum algorithms, *PRX Quantum* **2**, 030324 (2021).
- [86] W. Wirtinger, Zur formalen theorie der funktionen von mehr komplexen veränderlichen, *Math. Ann.* **97**, 357 (1927).
- [87] K. Kreutz-Delgado, The complex gradient operator and the CR-calculus, [arXiv:0906.4835](https://arxiv.org/abs/0906.4835).
- [88] Y. Yao, P. Cussenot, R. A. Wolf, and F. Miatto, Complex natural gradient optimization for optical quantum circuit design, *Phys. Rev. A* **105**, 052402 (2022).
- [89] S. Zhang, Y. Xia, and J. Wang, A complex-valued projection neural network for constrained optimization of real functions in complex variables, *IEEE Trans. Neural Networks Learn. Syst.* **26**, 3227 (2015).
- [90] A. Hirose, *Complex-Valued Neural Networks* (Springer, Berlin, 2012), Vol. 400.
- [91] O. M. Smirnov and C. Tasse, Radio interferometric gain calibration as a complex optimization problem, *Mon. Not. R. Astron. Soc.* **449**, 2668 (2015).
- [92] J. M. Arrazola, T. R. Bromley, J. Izaac, C. R. Myers, K. Brádler, and N. Killoran, Machine learning method for state preparation and gate synthesis on photonic quantum computers, *Quantum Sci. Technol.* **4**, 024004 (2019).
- [93] N. Killoran, T. R. Bromley, J. M. Arrazola, M. Schuld, N. Quesada, and S. Lloyd, Continuous-variable quantum neural networks, *Phys. Rev. Res.* **1**, 033063 (2019).
- [94] J. Cortés-Vega, J. F. Barra, L. Pereira, and A. Delgado, Detecting entanglement of unknown states by violating the Clauser-Horne-Shimony-Holt inequality, *Quantum Inf. Process.* **22**, 203 (2023).
- [95] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms, *J. Magn. Reson.* **172**, 296 (2005).
- [96] L. Sorber, M. V. Barel, and L. D. Lathauwer, Unconstrained optimization of real functions in complex variables, *SIAM J. Opt.* **22**, 879 (2012).
- [97] H. Robbins and S. Monro, A stochastic approximation method, *Ann. Math. Stat.* **22**, 400 (1951).
- [98] J. Kiefer and J. Wolfowitz, Stochastic estimation of the maximum of a regression function, *Ann. Math. Stat.* **23**, 462 (1952).
- [99] S. Amari and S. Douglas, Why natural gradient? in *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP '98, Seattle, WA, USA* (IEEE, New York, 2002), pp. 1213–1216.
- [100] G. Yan and H. Fan, A Newton-like algorithm for complex variables with applications in blind equalization, *IEEE Trans. Signal Process.* **48**, 553 (2000).
- [101] M. Muñoz, L. Pereira, S. Niklitschek, and A. Delgado, Complex field formulation of the quantum estimation theory, [arXiv:2203.03064](https://arxiv.org/abs/2203.03064).
- [102] G. Aleksandrowicz *et al.*, QISKIT: An open-source framework for quantum computing (2021).
- [103] J. Gidi, [jgidi/complexspsa.jl](https://github.com/jgidi/complexspsa.jl): v0.2.1 (2023), [Zenodo](https://zenodo.org/record/7811111).
- [104] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, JULIA: A fresh approach to numerical computing, *SIAM Rev.* **59**, 65 (2017).
- [105] W. Ritz, Über eine neue methode zur lösung gewisser variationsprobleme der mathematischen physik, *J. Reine Angew. Math.* **135**, 1 (1909).
- [106] J. W. Rayleigh, In finding the correction for the open end of an organ-pipe, *Philos. Trans.* **161**, 75 (1870).
- [107] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [108] M. Shapiro and P. Brumer, *Quantum Control of Molecular Processes* (Wiley-VCH Verlag, Weinheim, Germany, 2011).
- [109] S. Chu, Cold atoms and quantum control, *Nature (London)* **416**, 206 (2002).
- [110] R. van Handel, J. K. Stockton, and H. Mabuchi, Modelling and feedback control design for quantum state preparation, *J. Opt. B: Quantum Semiclass. Opt.* **7**, S179 (2005).
- [111] S. Lloyd, Coherent quantum feedback, *Phys. Rev. A* **62**, 022108 (2000).
- [112] J. Werschnik and E. K. U. Gross, Quantum optimal control theory, *J. Phys. B* **40**, R175 (2007).
- [113] D. Burgarth, P. Facchi, H. Nakazato, S. Pascazio, and K. Yuasa, Generalized adiabatic theorem and strong-coupling limits, *Quantum* **3**, 152 (2019).
- [114] G. Quiroz, Robust quantum control for adiabatic quantum computation, *Phys. Rev. A* **99**, 062306 (2019).
- [115] B. B. Zhou, A. Baksic, H. Ribeiro, C. G. Yale, F. J. Heremans, P. C. Jerger, A. Auer, G. Burkard, A. A. Clerk, and D. D. Awschalom, Accelerated quantum control using superadiabatic dynamics in a solid-state lambda system, *Nat. Phys.* **13**, 330 (2017).
- [116] Z. Wu and H. Yang, Validity of the quantum adiabatic theorem, *Phys. Rev. A* **72**, 012114 (2005).
- [117] K. Laiho, M. Avenhaus, and C. Silberhorn, Characteristics of displaced single photons attained via higher order factorial moments, *New J. Phys.* **14**, 105011 (2012).
- [118] A. M. Brańczyk, D. H. Mahler, L. A. Rozema, A. Darabi, A. M. Steinberg, and D. F. V. James, Self-calibrating quantum state tomography, *New J. Phys.* **14**, 085003 (2012).
- [119] C. R. Müller, B. Stoklasa, C. Peuntinger, C. Gabriel, J. Řeháček, Z. Hradil, A. B. Klimov, G. Leuchs, C. Marquardt, and L. L. Sánchez-Soto, Quantum polarization tomography of bright squeezed light, *New J. Phys.* **14**, 085002 (2012).

- [120] A. Chiuri, L. Mazzola, M. Paternostro, and P. Mataloni, Tomographic characterization of correlations in a photonic tripartite state, *New J. Phys.* **14**, 085006 (2012).
- [121] S. Wallentowitz, B. Seifert, and S. Godoy, Local sampling of the quantum phase-space distribution of a continuous-wave optical beam, *New J. Phys.* **14**, 105019 (2012).
- [122] C. Sayrin, I. Dotsenko, S. Gleyzes, M. Brune, J. M. Raimond, and S. Haroche, Optimal time-resolved photon number distribution reconstruction of a cavity field by maximum likelihood, *New J. Phys.* **14**, 115007 (2012).
- [123] S. L. Christensen, J. B. Béguin, H. L. Sørensen, E. Bookjans, D. Oblak, J. H. Müller, J. Appel, and E. S. Polzik, Toward quantum state tomography of a single polariton state of an atomic ensemble, *New J. Phys.* **15**, 015002 (2013).
- [124] R. Rey-de Castro, R. Cabrera, D. I. Bondar, and H. Rabitz, Time-resolved quantum process tomography using Hamiltonian-encoding and observable-decoding, *New J. Phys.* **15**, 025032 (2013).
- [125] M. W. Mitchell, M. Koschorreck, M. Kubasik, M. Napolitano, and R. J. Sewell, Certified quantum non-demolition measurement of material systems, *New J. Phys.* **14**, 085021 (2012).
- [126] M. Guță, T. Kypraios, and I. Dryden, Rank-based model selection for multiple ions quantum tomography, *New J. Phys.* **14**, 105002 (2012).
- [127] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Chek-al kar, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt *et al.*, Scalable multiparticle entanglement of trapped ions, *Nature (London)* **438**, 643 (2005).
- [128] L. Zhang, A. Datta, H. B. Coldenstrodt-Ronge, X.-M. Jin, J. Eisert, M. B. Plenio, and I. A. Walmsley, Recursive quantum detector tomography, *New J. Phys.* **14**, 115005 (2012).
- [129] G. Brida, L. Ciavarella, P. Degiovanni, M. Genovese, L. Lolli, M. G. Mingolla, F. Piacentini, M. Rajteri, E. Taralli, and M. G. A. Paris, Quantum characterization of superconducting photon counters, *New J. Phys.* **14**, 085001 (2012).
- [130] A. Anis and A. I. Lvovsky, Maximum-likelihood coherent-state quantum process tomography, *New J. Phys.* **14**, 105021 (2012).
- [131] S. Watanabe, R. Matsumoto, and T. Uyematsu, Tomography increases key rates of quantum-key-distribution protocols, *Phys. Rev. A* **78**, 042316 (2008).