

Estimation of Quantum Fisher Information via Stein's Identity in Variational Quantum Algorithms

Mourad Halla

Deutsches Elektronen-Synchrotron DESY, Platanenallee 6, 15738 Zeuthen, Germany

The Quantum Fisher Information Matrix (QFIM) plays a crucial role in quantum optimization algorithms such as Variational Quantum Imaginary Time Evolution and Quantum Natural Gradient Descent. However, computing the full QFIM incurs a quadratic computational cost of $O(d^2)$ with respect to the number of parameters d , limiting its scalability for high-dimensional quantum systems. To address this limitation, stochastic methods such as the Simultaneous Perturbation Stochastic Approximation (SPSA) have been employed to reduce computational complexity to a constant [Quantum **5**, 567 (2021)]. In this work, we propose an alternative estimation framework based on Stein's identity that also achieves constant computational complexity. Furthermore, our method reduces the quantum resources required for QFIM estimation compared to the SPSA approach. We provide numerical examples using the transverse-field Ising model and the lattice Schwinger model to demonstrate the feasibility of applying our method to realistic quantum systems.

1 Introduction

Quantum computing has emerged as a transformative framework for tackling computational challenges that exceed the reach of classical algorithms. Among the leading strategies in this field, *Variational Quantum Algorithms (VQAs)* [1, 2, 3], which use a hybrid quantum-classical optimization scheme, have garnered significant attention due to their applicability in areas such as quantum chemistry, materials science, and high-energy physics. A particularly impactful algorithm within this class is the *Variational Quantum Eigensolver (VQE)* [4], which is designed to approximate ground-state energies of quantum systems, making it especially well-suited for near-term noisy intermediate-scale quantum (NISQ) devices.

Optimization strategies play a pivotal role in the efficiency of VQAs, influencing the quality of the approximated solutions and the quantum resources required to reach them. Conventional optimization algorithms, such as gradient descent (GD), often struggle to navigate the intricate nature of quantum optimization landscapes, as these landscapes in VQAs are typically non-convex and prone to noise and barren plateaus. This challenge necessitates the development of more sophisticated optimization techniques to enhance convergence speed and improve the accuracy of approximated solutions. Some efficient quantum optimization algorithms include Quantum Natural Gradient (QNG) [5] and Variational Quantum Imaginary Time Evolution (VarQITE) [6] that use an update rule based on the Quantum Fisher Information Matrix (QFIM).

The QFIM serves as a fundamental metric that characterizes the local curvature of a quantum state manifold, providing a natural geometric framework for parameter optimization. Unlike classical approaches that rely on Euclidean distances in parameter space, the QFIM captures how quantum states evolve under small parameter variations, ensuring an optimization trajectory aligned with the true structure of quantum state space. Moreover, QFIM-based techniques have broader implications beyond optimization in VQAs, contributing to quantum sensing and quantum metrology (see the review in [7, 8, 9]).

Despite this, standard methods like the parameter-shift rule [10] for computing the QFIM scale quadratically as $O(d^2)$ with the number of parameters d , making them impractical for high-dimensional quantum systems. This limitation necessitates alternative approaches for efficiently computing the QFIM. In [5], diagonal and block-diagonal approximation were proposed to reduce computational complexity, but these methods result in the loss of parameter correlations in the QFIM. More efficient approaches that preserve these correlations involve the use of the simultaneous perturbation stochastic approximation (SPSA) [12] method to approximate the QFIM, reducing the complexity cost to a constant [11]. In this work, we propose an alternative and efficient method based on Stein’s identity, which also reduces the quantum computational complexity from $O(d^2)$ to a constant without losing parameter correlations.

Stein’s identity provides a powerful framework for estimating Hessian information in stochastic optimization, particularly in zeroth-order (ZO) settings where direct gradient and Hessian computations are infeasible. Recent studies have leveraged it to develop more efficient Hessian approximation methods, reducing per-iteration complexity compared to traditional second-order techniques such as the simultaneous perturbation stochastic approximation (2SPSA) [14, 15]. Employing a perturbation-based approach enables unbiased gradient and Hessian estimation while requiring fewer ZO queries than 2SPSA, thereby improving convergence efficiency. In this work, we use a similar idea to estimate the QFIM.

The remainder of the manuscript is organized as follows: Section 2 provides a mathematical overview of Stein’s identity, demonstrating its application in estimating gradients, the Hessian, and the QFIM. Section 3 presents a brief overview of variational quantum algorithms and the Quantum Natural Gradient. Section 4 includes numerical examples using VQE to compute the ground-state energy of the transverse-field Ising model and the lattice Schwinger model. Finally, Section 5 provides a summary and outlook.

2 Theoretical Results

In this section, we provide a theoretical overview of gradient and Hessian estimation using Stein’s identity, and subsequently integrate the Stein’s identity method into the QFIM framework. As Stein’s identity has been proposed as an alternative to Simultaneous Perturbation Stochastic Approximation (SPSA) for reducing computational complexity, we first review the SPSA method before introducing the Stein-based approach.

2.1 Gradient and Hessian Estimation via SPSA Methods

Simultaneous Perturbation Stochastic Approximation (SPSA) [12] provides an efficient method for estimating gradients using stochastic perturbations, significantly reducing computational costs compared to finite-difference approaches. Given a differentiable objective function $f: \mathbb{R}^d \rightarrow \mathbb{R}$, SPSA estimates the gradient using a two-point perturbation method.

Let $\Delta \sim \{\pm 1\}^d$ be a random perturbation vector drawn from the Rademacher distribution (each coordinate independently takes ± 1 with probability $1/2$). Then the gradient

is estimated as

$$\nabla f(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{\Delta}} \left[\frac{f(\boldsymbol{\theta} + c \boldsymbol{\Delta}) - f(\boldsymbol{\theta} - c \boldsymbol{\Delta})}{2c} \boldsymbol{\Delta} \right], \quad (1)$$

where $c > 0$ is a small finite displacement. This method requires only two function evaluations per iteration, and $\boldsymbol{\Delta}$ is independent of $\boldsymbol{\theta}$, making it scalable to high-dimensional problems.

The second-order SPSA (2SPSA) extends SPSA to approximate the Hessian matrix $\mathbf{H}(\boldsymbol{\theta})$ without explicitly computing all d^2 entries. Two independent perturbation vectors $\boldsymbol{\Delta}_1$ and $\boldsymbol{\Delta}_2 \sim \{\pm 1\}^d$ are sampled, leading to the Hessian estimate

$$\hat{\mathbf{H}} = \mathbb{E}_{\boldsymbol{\Delta}_1, \boldsymbol{\Delta}_2} \left[\frac{\delta f}{2c^2} \frac{\boldsymbol{\Delta}_1 \boldsymbol{\Delta}_2^\top + \boldsymbol{\Delta}_2 \boldsymbol{\Delta}_1^\top}{2} \right], \quad (2)$$

where

$$\delta f = f(\boldsymbol{\theta} + c \boldsymbol{\Delta}_1 + c \boldsymbol{\Delta}_2) - f(\boldsymbol{\theta} + c \boldsymbol{\Delta}_1) - f(\boldsymbol{\theta} - c \boldsymbol{\Delta}_1 + c \boldsymbol{\Delta}_2) + f(\boldsymbol{\theta} - c \boldsymbol{\Delta}_1). \quad (3)$$

This approach reduces the $O(d^2)$ computational cost of explicit Hessian estimation while maintaining reasonable accuracy. The 2SPSA method requires four function evaluations and two perturbation vectors $\boldsymbol{\Delta}_1$ and $\boldsymbol{\Delta}_2$. In the following sections, we will see that Stein's identity requires two or three function evaluations and a single perturbation vector.

2.2 Gradient and Hessian Estimation via Stein's Identity

Stein's identity provides a fundamental mathematical tool for estimating gradients and Hessians, particularly in optimization problems where derivative information is inaccessible or computationally expensive. In this section, we provide an overview of this method and its application in estimating gradients and Hessians. For a more detailed mathematical treatment and rigorous proofs, we refer the reader to [13, 14, 15].

A central result of Stein's method is summarized in the following proposition:

Proposition 1 (First-Order and Second-Order Stein's Identity [13]). *Let $\mathbf{X} \in \mathbb{R}^d$ represent a d -dimensional random vector with an underlying probability density function $p(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$.*

- i) If $p(\mathbf{x})$ is differentiable, and $q : \mathbb{R}^d \rightarrow \mathbb{R}$ is a differentiable function such that $\mathbb{E} \{ \nabla q(\mathbf{X}) \}$ exists, then the following identity holds:*

$$\mathbb{E} \left\{ q(\mathbf{X}) [p(\mathbf{X})]^{-1} \nabla p(\mathbf{X}) \right\} = -\mathbb{E} \{ \nabla q(\mathbf{X}) \}. \quad (4)$$

- ii) If $p(\mathbf{x})$ and $q(\mathbf{x})$ are twice differentiable functions such that $\mathbb{E} \{ \nabla^2 q(\mathbf{X}) \}$ exists, then:*

$$\mathbb{E} \left\{ q(\mathbf{X}) [p(\mathbf{X})]^{-1} \nabla^2 p(\mathbf{X}) \right\} = \mathbb{E} \left\{ \nabla^2 q(\mathbf{X}) \right\}. \quad (5)$$

When considering a multivariate standard normal random vector $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, where $\mathbf{I}$ represents the identity matrix, we have $\nabla p(\mathbf{x}) = -\mathbf{x}p(\mathbf{x})$ and $\nabla^2 p(\mathbf{x}) = (\mathbf{x}\mathbf{x}^\top - \mathbf{I})p(\mathbf{x})$. Using these expressions, equations (4) and (5) in Proposition 1 take the form:

$$\mathbb{E} \{ \mathbf{X} q(\mathbf{X}) \} = \mathbb{E} \{ \nabla q(\mathbf{X}) \}, \quad (6)$$

$$\mathbb{E} \left\{ (\mathbf{X}\mathbf{X}^\top - \mathbf{I}) q(\mathbf{X}) \right\} = \mathbb{E} \left\{ \nabla^2 q(\mathbf{X}) \right\}. \quad (7)$$

For the case of $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$ where $\mathbf{\Sigma}$ is an arbitrary positive definite covariance matrix, the gradient and Hessian of the probability density function are given by $\nabla p(\mathbf{x}) = -\mathbf{\Sigma}^{-1}\mathbf{x}p(\mathbf{x})$ and $\nabla^2 p(\mathbf{x}) = (\mathbf{\Sigma}^{-1}\mathbf{x}\mathbf{x}^T\mathbf{\Sigma}^{-1} - \mathbf{\Sigma}^{-1})p(\mathbf{x})$. Under these conditions, equations (4) and (5) reduce to:

$$\mathbb{E} \left\{ \mathbf{\Sigma}^{-1} \mathbf{X} q(\mathbf{X}) \right\} = \mathbb{E} \left\{ \nabla q(\mathbf{X}) \right\}, \quad (8)$$

$$\mathbb{E} \left\{ (\mathbf{\Sigma}^{-1} \mathbf{X} \mathbf{X}^T \mathbf{\Sigma}^{-1} - \mathbf{\Sigma}^{-1}) q(\mathbf{X}) \right\} = \mathbb{E} \left\{ \nabla^2 q(\mathbf{X}) \right\}. \quad (9)$$

To approximate the gradient and Hessian of a function $f(\boldsymbol{\theta})$ in optimization problems, we define, for any $c > 0$ and $\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$,

$$f_c(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{u}} [f(\boldsymbol{\theta} + c\mathbf{u})]. \quad (10)$$

The function $f_c(\boldsymbol{\theta})$ is thus C^∞ . Stein's identity offers various approximation orders, requiring one, two, or three function evaluations to estimate either the unbiased gradient or the unbiased Hessian. The results are summarized in the following lemma:

Lemma 1 (Stein's Identity-Based Estimator [14]). *Let $\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and consider $f_c(\boldsymbol{\theta})$ defined by (10). Then the gradient $\nabla f_c(\boldsymbol{\theta})$ and Hessian $\nabla^2 f_c(\boldsymbol{\theta})$ can be estimated as follows:*

i) Single-Evaluation Estimator:

$$\nabla f_c(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{u}} \{ c^{-1} f(\boldsymbol{\theta} + c\mathbf{u}) \mathbf{u} \}, \quad (11)$$

$$\nabla^2 f_c(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{u}} \{ c^{-2} f(\boldsymbol{\theta} + c\mathbf{u}) (\mathbf{u} \mathbf{u}^T - \mathbf{I}) \}. \quad (12)$$

ii) Two-Evaluation Estimator:

$$\nabla f_c(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{u}} \{ (2c)^{-1} (f(\boldsymbol{\theta} + c\mathbf{u}) - f(\boldsymbol{\theta} - c\mathbf{u})) \mathbf{u} \}, \quad (13)$$

$$\nabla^2 f_c(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{u}} \{ c^{-2} (f(\boldsymbol{\theta} + c\mathbf{u}) - f(\boldsymbol{\theta})) (\mathbf{u} \mathbf{u}^T - \mathbf{I}) \}. \quad (14)$$

iii) Three-Evaluation Estimator:

$$\nabla^2 f_c(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{u}} \{ (2c^2)^{-1} (f(\boldsymbol{\theta} + c\mathbf{u}) + f(\boldsymbol{\theta} - c\mathbf{u}) - 2f(\boldsymbol{\theta})) (\mathbf{u} \mathbf{u}^T - \mathbf{I}) \}. \quad (15)$$

To provide flexibility for QFIM estimation in the next section, we extend Lemma 1 to the case where $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$, with an arbitrary positive-definite covariance matrix $\mathbf{\Sigma}$. In this work, we focus on the special case where $\mathbf{\Sigma} = b^2 \mathbf{I}$, and summarize our results for Hessian estimation in the following corollary:

Corollary 1. *Let $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, b^2 \mathbf{I})$, and define $\mathbf{X} = c\mathbf{Y}$, with $\mathbf{Y} \sim \mathcal{N}(\mathbf{0}, \frac{b^2}{c^2} \mathbf{I})$. Assume $b, c > 0$. Then the two-evaluation estimator of the Hessian of the smoothed function $f_c(\boldsymbol{\theta})$ is given by:*

$$\nabla^2 f_c(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{Y}} \left\{ \frac{c^2}{b^4} [f(\boldsymbol{\theta} + c\mathbf{Y}) - f(\boldsymbol{\theta})] (\mathbf{Y} \mathbf{Y}^T - \frac{b^2}{c^2} \mathbf{I}) \right\}. \quad (16)$$

and the three-evaluation estimator is given by:

$$\nabla^2 f_c(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{Y}} \left\{ \frac{c^2}{2b^4} [f(\boldsymbol{\theta} + c\mathbf{Y}) + f(\boldsymbol{\theta} - c\mathbf{Y}) - 2f(\boldsymbol{\theta})] (\mathbf{Y} \mathbf{Y}^T - \frac{b^2}{c^2} \mathbf{I}) \right\}. \quad (17)$$

Proof. The result follows from Stein's second-order identity (9), applied to the transformation $\mathbf{X} = c\mathbf{Y}$. Given that $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, b^2 \mathbf{I})$, substituting $\mathbf{X} = c\mathbf{Y}$ ensures that $\mathbf{Y} \sim \mathcal{N}(\mathbf{0}, \frac{b^2}{c^2} \mathbf{I})$. $\square$

In the next section, we extend a similar approach used for Hessian estimation to the Quantum Fisher Information Matrix.

2.3 QFIM Estimation via Stein's Identity

Similar to the approach proposed by Gacon [11], which employs 2SPSA [12] to estimate the QFIM, we propose an alternative method based on Stein's identity for QFIM estimation. We begin with a brief discussion of the QFIM and refer to [7, 8, 9] for a general overview and its role in variational quantum algorithms.

Let $|\psi(\boldsymbol{\theta})\rangle$ be a parameterized pure quantum state in an n -qubit Hilbert space, where $\boldsymbol{\theta} \in \mathbb{R}^d$ represents a set of d trainable parameters. Fix a small displacement $\delta\boldsymbol{\theta} \in \mathbb{R}^d$ and define the overlap function:

$$\mathcal{F}(\boldsymbol{\theta}) := |\langle\psi(\boldsymbol{\theta}) | \psi(\boldsymbol{\theta} + \delta\boldsymbol{\theta})\rangle|^2. \quad (18)$$

The Fubini–Study metric tensor is given by

$$F_{ij}(\boldsymbol{\theta}) = -\frac{1}{2} \partial_i \partial_j \mathcal{F}(\boldsymbol{\theta}) \Big|_{\delta\boldsymbol{\theta}=0}, \quad (19)$$

so that the QFIM equals $4 F_{ij}(\boldsymbol{\theta})$. Explicitly,

$$F_{ij}(\boldsymbol{\theta}) = \text{Re} \left[\langle \partial_i \psi(\boldsymbol{\theta}) | \partial_j \psi(\boldsymbol{\theta}) \rangle - \langle \partial_i \psi(\boldsymbol{\theta}) | \psi(\boldsymbol{\theta}) \rangle \langle \psi(\boldsymbol{\theta}) | \partial_j \psi(\boldsymbol{\theta}) \rangle \right], \quad (20)$$

where ∂_i and ∂_j denote partial derivatives with respect to θ_i and θ_j . See Appendix A for detailed computations. Evaluating (20) using methods such as the parameter-shift rule requires $O(d^2)$ computations, rendering it impractical for large-scale VQA optimization.

We now examine how this issue can be addressed using Stein's identity. At first glance, we observe that (19) resembles a Hessian, which allows us to apply the results from Section 2.2 for its estimation. For practical purposes, we focus on the two-evaluation and three-evaluation methods, presented in (16) and (17), respectively. Furthermore, assuming that the rows of $\mathbf{Y}$ are independent and identically distributed (i.i.d.) random vectors drawn from $\mathcal{N}(\mathbf{0}, \frac{b^2}{c^2} \mathbf{I})$. Higher accuracy can be achieved by generating N independent perturbation vectors $\mathbf{Y}_i$.

Let $c > 0$ and $\mathbf{Y} \sim \mathcal{N}(\mathbf{0}, \frac{b^2}{c^2} \mathbf{I})$. Based on (18), define the smoothed function

$$\mathcal{F}_c(\boldsymbol{\theta}) := \mathbb{E}_{\mathbf{Y}} [\mathcal{F}(\boldsymbol{\theta} + c \mathbf{Y})]. \quad (21)$$

The two-evaluation Stein estimator for the metric tensor is

$$\hat{\mathbf{F}} = -\frac{c^2}{2b^4N} \sum_{i=1}^N [\mathcal{F}(\boldsymbol{\theta} + c \mathbf{Y}_i) - \mathcal{F}(\boldsymbol{\theta})] (\mathbf{Y}_i \mathbf{Y}_i^\top - \frac{b^2}{c^2} \mathbf{I}), \quad (22)$$

and the three-evaluation estimator is

$$\hat{\mathbf{F}} = -\frac{c^2}{4b^4N} \sum_{i=1}^N [\mathcal{F}(\boldsymbol{\theta} + c \mathbf{Y}_i) + \mathcal{F}(\boldsymbol{\theta} - c \mathbf{Y}_i) - 2\mathcal{F}(\boldsymbol{\theta})] (\mathbf{Y}_i \mathbf{Y}_i^\top - \frac{b^2}{c^2} \mathbf{I}). \quad (23)$$

Both estimators (22) and (23) are unbiased for $\nabla^2 \mathcal{F}_c(\boldsymbol{\theta})$, i.e., $\mathbb{E}[\hat{\mathbf{F}}] = \nabla^2 \mathcal{F}_c(\boldsymbol{\theta})$. We emphasize that by letting the smoothing parameter $c \rightarrow 0$ while independently increasing the Monte Carlo sample size $N \rightarrow \infty$, the bias of (22)–(23) with respect to the true metric tensor $\nabla^2 \mathcal{F}(\boldsymbol{\theta})$ vanishes asymptotically as $O(c^2)$. In practice, one chooses c and N so that this residual bias is negligible compared to statistical fluctuations or shot noise.

In equations (22) and (23), the computational complexity is reduced from $O(d^2)$ to a constant, making it independent of the number of parameters d . Compared to the

2SPSA QFIM estimator proposed by Gacon [11], which requires four circuit evaluations, our method based on Stein’s identity offers greater flexibility. Users can choose between the two-circuit evaluation QFIM estimator, (22), or the three-circuit evaluation QFIM estimator, (23). Furthermore, the 2SPSA QFIM method requires two independent perturbation vectors, Δ_1 and Δ_2 , whereas the Stein-based method requires only a single perturbation vector, $\mathbf{Y}$. The approximation error in the QFIM estimates scales as $O(N^{-1/2})$.

The Fubini–Study metric, or QFIM, involves evaluating the squared overlap between quantum states $|\psi(\boldsymbol{\theta})\rangle$ and $|\psi(\boldsymbol{\theta} + \delta\boldsymbol{\theta})\rangle$, i.e., $|\langle\psi(\boldsymbol{\theta}) | \psi(\boldsymbol{\theta} + \delta\boldsymbol{\theta})\rangle|^2$. The estimation method employed in this work prepares the state $U^\dagger(\boldsymbol{\theta} + \delta\boldsymbol{\theta})U(\boldsymbol{\theta})|0\rangle$, where $U(\boldsymbol{\theta})$ is a parameterized unitary, and measures the probability of $|0\rangle$, which directly yields the overlap. This approach maintains the original circuit width n , where n is the number of qubits, but doubles the circuit depth to $2m$, with m representing the depth of $U(\boldsymbol{\theta})$.

In this manuscript, we utilize (22) and (23) as examples within the framework of the Quantum Natural Gradient (QNG) applied to the Variational Quantum Eigensolver (VQE). We leave the application to imaginary time evolution for future work. In the following section, we provide an overview of QNG.

3 Quantum Natural Gradient and VQE

In the Variational Quantum Eigensolver (VQE), the objective function is typically given by

$$\mathcal{L}(\boldsymbol{\theta}) = \langle 0 | U^\dagger(\boldsymbol{\theta}) O U(\boldsymbol{\theta}) | 0 \rangle, \quad (24)$$

where O is a Hermitian operator, and the parameterized unitary is

$$U(\boldsymbol{\theta}) = \prod_{\ell=1}^p W_\ell \exp(i\theta_\ell X_\ell), \quad (25)$$

with W_ℓ and X_ℓ being fixed unitary and Hermitian operators, respectively.

The parameter vector $\boldsymbol{\theta}$ is iteratively updated to minimize the objective function $\mathcal{L}(\boldsymbol{\theta})$. In standard gradient descent, the update rule is given by

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \eta \nabla \mathcal{L}(\boldsymbol{\theta}_k), \quad (26)$$

where $\eta > 0$ is a user-defined learning rate (step size), and $\nabla \mathcal{L}(\boldsymbol{\theta}_k)$ denotes the Euclidean gradient.

In the Quantum Natural Gradient (QNG) method, the Euclidean gradient is replaced by the Riemannian gradient $\mathbf{F}^{-1}(\boldsymbol{\theta}_k) \nabla \mathcal{L}(\boldsymbol{\theta}_k)$, where $\mathbf{F}^{-1}(\boldsymbol{\theta}_k)$ is the inverse of the Fubini–Study metric tensor. The update rule then becomes

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \eta \mathbf{F}^{-1}(\boldsymbol{\theta}_k) \nabla \mathcal{L}(\boldsymbol{\theta}_k). \quad (27)$$

To incorporate previous stochastic estimates of the metric tensor, we replace $\mathbf{F}$ in the update rule with $\bar{\mathbf{F}}_k$:

$$\bar{\mathbf{F}}_k = \frac{k}{k+1} \bar{\mathbf{F}}_{k-1} + \frac{1}{k+1} \hat{\mathbf{F}}_k, \quad (28)$$

where the metric estimate $\bar{\mathbf{F}}_k$ incorporates all previous samples. This recursive averaging introduces a small geometric error, as it combines metric tensors evaluated at slightly different parameter points, $\boldsymbol{\theta}_{k-1}$ and $\boldsymbol{\theta}_k$. However, for sufficiently small step sizes, this error remains negligible in practice (see Remark 1 in Appendix B).

To ensure the metric remains positive semi-definite and invertible, we apply the regularization strategy described in [11]. In particular, the metric defined in equation (22) or (23) is replaced by

$$\hat{\mathbf{F}}_{\text{reg}} = \frac{\sqrt{\hat{\mathbf{F}}^\top \hat{\mathbf{F}}} + \beta \mathbf{I}}{1 + \beta}, \quad (29)$$

where β is a tunable regularization coefficient selected by empirical testing. If β is too small, the adjusted metric may fail to remain positive semi-definite; if β is too large, the contribution of the metric is effectively erased, reducing the quantum natural gradient update to a first-order update and canceling the curvature information carried by the metric.

Unlike the Moore–Penrose pseudo-inverse, the regularization in equation (29) avoids uncontrolled amplification of noise-dominated eigenvalues and prevents discontinuous jumps in the inverse metric when small eigenvalues cross zero, thereby preserving a smooth, positive-definite Riemannian metric that yields stable natural-gradient updates.

In addition to the regularization of the metric, a blocking mechanism may be enforced by setting $\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k$ if the evaluation of the noisy objective function at $\boldsymbol{\theta}_{k+1}$ is substantially higher than at $\boldsymbol{\theta}_k$ by a user-specified constant. This ensures stability in the update rule and prevents divergence.

For the simulations in the next section, we use in the update rule (27) the metric estimators (22) and (23), regularized according to (29). For the gradient, we employ the two-function evaluation estimator given in equation (13):

$$\hat{\mathbf{g}}_k = \frac{1}{N} \sum_{i=1}^N (2c)^{-1} (\mathcal{L}(\boldsymbol{\theta}_k + c\mathbf{u}_{k,i}) - \mathcal{L}(\boldsymbol{\theta}_k - c\mathbf{u}_{k,i})) \mathbf{u}_{k,i}, \quad (30)$$

where $\mathbf{u}_{k,i} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. This choice of metric tensor and gradient estimation offers good practical accuracy while significantly reducing the required number of resamplings N per iteration k .

4 Numerical Results

To demonstrate the practical applicability of our approach, we conducted numerical simulations of the VQE algorithm using the open-source software PennyLane [18].

4.1 Example 1: Transverse Field Ising Model

The Transverse Field Ising Model (TFIM) with open boundary conditions is described by the Hamiltonian:

$$H = J \sum_{i=1}^{N-1} \sigma_i^z \sigma_{i+1}^z + h \sum_{i=1}^N \sigma_i^x, \quad (31)$$

where J is the coupling constant, h is the transverse field strength, and σ_i^z, σ_i^x are Pauli matrices acting on site i . The first term represents nearest-neighbor spin interactions along the z -axis, while the second term introduces quantum fluctuations via the transverse field along the x -axis.

In this example, we investigate the case with $J = -1$ and $h = -2$, and approximate the ground state of H using the hardware-efficient ansatz. This ansatz constructs the wavefunction with a layered quantum circuit that combines parameterized single-qubit rotations (R_Y) and entangling controlled-NOT (CNOT) gates (see Fig. 1).




Figure 1: Hardware-efficient ansatz with two layers, using R_Y rotations and CNOT gates.

In the VQE simulation, we benchmark the following optimizers: Gradient Descent (GD), Quantum Natural Gradient (QNG), Simultaneous Perturbation Stochastic Approximation (SPSA), Quantum Natural SPSA (QNPSA), Stein Optimizer (using equation (30) for the gradient without the Fubini–Study metric), Quantum Natural Stein Optimizer 2 (QNSTEIN2, using the gradient from equation (30) and the metric from equation (22)), and Quantum Natural Stein Optimizer 3 (QNSTEIN3, using the gradient from equation (30) and the metric from equation (23)). For all stochastic optimizers (SPSA, QNPSA, STEIN, QNSTEIN2, and QNSTEIN3), resampling is performed at each optimization step k , with $N = 10$ samples used for both the gradient and the metric (see N in (30), (22), and (23)). The learning rate for all optimizers is fixed at 0.01. The regularization parameter ϵ is set to 10^{-2} for QNPSA, QNSTEIN2, and QNSTEIN3, and to 10^{-1} for QNG. The finite-difference step for SPSA and QNPSA is fixed at 0.05, while in QNSTEIN2 and QNSTEIN3, parameters $b = 2$ and $c = 0.05$ are used. All simulations use 8192 shots, and each optimizer is run for up to 300 steps.




Figure 2: Energy error as a function of iteration steps for different optimization methods—SPSA, QNPSA, STEIN, QNSTEIN2, and QNSTEIN3—applied to 3-layer circuits with varying numbers of qubits (12, 17, 20). The results are averaged over 30 different random initializations.




Figure 3: Same as in Figure 2, but this time with 12 qubits fixed and varying layers ($L = 2, 3, 6$).




Figure 4: Same as in Figure 3, but here we plot the energy error as a function of the number of quantum circuits required for convergence.

The performance of the optimizers is highlighted in Figures 2–4. Figure 2 shows the energy error (i.e., the difference between the obtained and target energies) as a function of iteration for a fixed three-layer circuit with varying numbers of qubits, indicating that the QNSTEINs tend to exhibit an overall better convergence rate in these benchmarks. Figure 3 supports this observation for a fixed 12-qubit system with varying numbers of layers. Figure 4 illustrates the quantum resources required for convergence, showing that the QNSTEINs achieve target energies with the fewest circuit evaluations, thereby reducing quantum resource usage. Across all simulations, there is no significant difference between SPSA and STEIN in terms of convergence behavior and quantum resource consumption.

4.2 Example 2: Schwinger Model

The Schwinger Model [16] is a (1+1)-dimensional quantum field theory that serves as a fundamental framework for studying quantum electrodynamics (QED) in reduced dimensions. Although it is lower-dimensional, the model retains essential features of more complex gauge theories, such as confinement, charge screening, and chiral symmetry breaking. These characteristics make it an important system for investigating non-perturbative effects in quantum field theory. Additionally, its lattice formulation facilitates efficient mapping onto quantum hardware, making it a promising candidate for exploring quantum simulations of gauge theories.

The dynamics of the Schwinger Model on a lattice are described by the Kogut-Susskind Hamiltonian. After mapping the Hamiltonian to qubits using the Jordan-Wigner transformation, the Hamiltonian for our quantum computing optimization task is:

$$H = \frac{x}{2} \sum_{n=0}^{N-2} (\sigma_n^x \sigma_{n+1}^x + \sigma_n^y \sigma_{n+1}^y) + \frac{\mu}{2} \sum_{n=0}^{N-1} [1 + (-1)^n \sigma_n^z] + \sum_{n=0}^{N-2} \left(l + \frac{1}{2} \sum_{k=0}^n (-1)^k \sigma_k^z \right)^2, \quad (32)$$

where the operators $\sigma_n^x, \sigma_n^y, \sigma_n^z$ represent the Pauli matrices applied to the qubit at site n . The parameter $x = 1/(g^2 a^2)$ is related to the coupling constant g and lattice spacing a , while $\mu = 2m/g^2 a$ is the dimensionless fermion mass term, with m being the fermion mass. The parameter l is a background electric field contribution, associated with the zero-mode of the gauge field.

To find the ground state of the Hamiltonian (32) using VQE, we employ the ansatz shown in Figure 5 [17], with parameters set to $l = 0$, $x = 1$, and $\mu = 0.5$. We benchmark the same optimizers and use the same hyperparameters as in Example 1; however, in this case, all optimizers run for up to 200 steps, with $N = 15$ samples and shot noise set to 10024.




Figure 5: (a) The ansatz used in the VQE algorithm to approximate the ground state of the Schwinger Hamiltonian, shown here with a single layer incorporating universal $SO(4)$ gates. (b) The decomposition of the two-qubit $SO(4)$ gate into single-qubit phase gates S and its conjugate transpose $S^\dagger$, the Hadamard gate H , and rotation gates (RX , RZ), along with entangling two-qubit CNOT operations.




Figure 6: Energy error as a function of iteration steps for different optimization methods—SPSA, QNSPSA, STEIN, QNSTEIN2, and QNSTEIN3—applied to 2-layer circuits with varying numbers of qubits (4, 6, 8). The Results are averaged over 30 random initializations of the variational parameters.




Figure 7: Same as in Figure 6, but here we plot the energy error as a function of the number of quantum circuits required for convergence.

As in the TFIM Hamiltonian, the results in Figures 6 and 7 show that the QNSTEIN optimizers also perform well in terms of convergence and in reducing the quantum resources required for the more complex, higher-energy physics Hamiltonian of the Schwinger model.

5 Conclusions and Outlook

In this work, we proposed an approach for estimating the Quantum Fisher Information Matrix (QFIM) within Variational Quantum Algorithms (VQAs) using Stein’s identity. The method significantly reduces computational complexity from $O(m^2)$ to a constant, effectively preserving parameter correlations and enhancing performance in quantum optimization tasks. Compared to the existing QFIM estimation method based on Simultaneous

Perturbation Stochastic Approximation (SPSA) proposed in [11], which also reduces the computational complexity to a constant, our Stein-based estimator requires fewer quantum circuit evaluations (function overlaps), thereby reducing the overall quantum resource requirements.

We validated our approach by benchmarking various optimizers in VQE simulations applied to the Transverse Field Ising Model and the lattice Schwinger Model. The results suggest that Quantum Natural Stein algorithms tend to achieve lower energy values with fewer iterations and reduced quantum resource consumption in these benchmarks, highlighting the potential benefits of integrating our Stein-based QFIM estimator into practical variational quantum algorithms.

Quantitatively, our estimators notably reduce the number of circuit evaluations per optimization step compared to Quantum Natural SPSA (QNPSA). Specifically, QNPSA typically requires two evaluations for gradients and four for the Fubini–Study metric, totaling six evaluations per step. In contrast, the two-evaluation Stein estimator (QNSTEIN2) uses two circuits for gradient and two for metric estimation, totaling four evaluations per step—a $1.5\times$ improvement. The three-evaluation Stein estimator (QNSTEIN3) employs two circuits for gradient estimation and three for metric estimation, totaling five circuits per step—a $1.2\times$ improvement. Considering metric evaluations alone, QNSTEIN2 achieves a $2\times$ reduction, and QNSTEIN3 achieves approximately a $1.33\times$ reduction. Although practical factors such as finite-shot noise, estimator variance, and parameter settings (c and b) may slightly affect these theoretical gains, appropriate tuning or modest increases in sampling generally allow the theoretical advantages to translate into practical benefits.

The presented framework is flexible and opens avenues for further developments, including exploring alternative perturbation distributions and adaptive resampling strategies to enhance accuracy and robustness. Moreover, this QFIM estimation technique may naturally extend to other quantum computational domains, such as quantum metrology.

Acknowledgements

This work was supported by the Ministry of Science, Research, and Culture of the State of Brandenburg through the Centre for Quantum Technologies and Applications (CQTA) at DESY (Germany) and by the German Ministry of Education and Research (BMBF) under project NiQ (Grant No. 13N16203). The author thanks Yibin Guo (DESY, Germany) for valuable discussions on the Schwinger model.



References

- [1] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, et al., “Variational Quantum Algorithms,” *Nature Reviews Physics*, 3, 625-644 (2021). <https://doi.org/10.1038/s42254-021-00348-9>.
- [2] J. R. McClean, J. Romero, R. Babbush and A. Aspuru-Guzik, “The theory of variational hybrid quantum-classical algorithms,” *New Journal of Physics*, 18, 023023 (2016). <https://doi.org/10.1088/1367-2630/18/2/023023>.
- [3] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, et al., “Noisy intermediate-scale quantum algorithms,” *Reviews of Modern Physics*, 94, 015004 (2022). <https://doi.org/10.1103/RevModPhys.94.015004>.
- [4] 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,” *Nature Communications* 5, 4213 (2014). <https://doi.org/10.1038/ncomms5213>.
- [5] J. Stokes, J. Izaac, N. Killoran and G. Carleo, “Quantum Natural Gradient,” *Quantum* 4, 269 (2020). <https://doi.org/10.22331/q-2020-05-25-269>.
- [6] S. McArdle, T. Jones, S. Endo, Y. Li, S. Benjamin and X. Yuan, “Variational ansatz-based quantum simulation of imaginary time evolution,” *npj Quantum Information* 5, 75 (2019). <https://doi.org/10.1038/s41534-019-0187-2>.
- [7] J. Jakob Meyer, “Fisher Information in Noisy Intermediate-Scale Quantum Applications,” *Quantum* 5, 539 (2021). <https://doi.org/10.22331/q-2021-09-09-539>.
- [8] J. Liu, H. Yuan, X. Lu and X. Wang, “Quantum Fisher information matrix and multiparameter estimation,” *J. Phys. A: Math. Theor.* 53, 023001 (2020). <https://doi.org/10.1088/1751-8121/ab5d4d>.
- [9] J. S. Sidhu and P. Kok, “Geometric perspective on quantum parameter estimation Available to Purchase,” *AVS Quantum Sci.* 2, 014701 (2020). <https://doi.org/10.1116/1.5119961>.
- [10] A. Mari, T. R. Bromley, and N. Killoran, “Estimating the gradient and higher-order derivatives on quantum hardware,” *Phys. Rev. A* 103, 012405 (2021). <https://doi.org/10.1103/PhysRevA.103.012405>.
- [11] J. Gacon, C. Zoufal, G. Carleo and S. Woerner, “Simultaneous Perturbation Stochastic Approximation of the Quantum Fisher Information,” *Quantum* 5, 567 (2021). <https://doi.org/10.22331/q-2021-10-20-567>.
- [12] J. C. Spall, “Multivariate stochastic approximation using a simultaneous perturbation gradient approximation,” *IEEE Transactions on Automatic Control* 37(3):332–341 (1992). <https://doi.org/10.1109/9.119632>.
- [13] C. Stein, P. Diaconis, S. Holmes and G. Reinert, “Use of exchangeable pairs in the analysis of simulations,” *Institute of Mathematical Statistics Lecture Notes – Monograph Series*, vol. 46, pp. 1–25, 2004. <https://doi.org/10.1214/lnms/1196283797>.
- [14] J. Zhu, “Hessian Estimation via Stein’s Identity in Black-Box Problems,” *Proceedings of the 2nd Mathematical and Scientific Machine Learning Conference*, PMLR 145, 1161–1178 (2022). <https://proceedings.mlr.press/v145/zhu22c.html>.
- [15] M. A. Erdogdu, “Newton-Stein Method: A Second Order Method for GLMs via Stein’s Lemma,” *Advances in Neural Information Processing Systems*, 28, 1216–1224 (2015). <https://papers.nips.cc/paper/5750-newton-stein-method-a-second-order-method-for-glms-via-steins-lemma>.
- [16] J. Schwinger, “Gauge invariance and mass. II,” *Physical Review* 128, 2425 (1962). <https://doi.org/10.1103/PhysRev.128.2425>.

- [17] Y. Guo , T. Angelides , K. Jansen and S. Kuehn, “Concurrent VQE for Simulating Excited States of the Schwinger Model,” (2024). <https://doi.org/10.48550/arXiv.2407.15629>.
- [18] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, S. Ahmed et al., “PennyLane: Automatic differentiation of hybrid quantum-classical computations,”. <https://doi.org/10.48550/arXiv.1811.04968>
- [19] M. Halla,, “Quantum Natural Gradient with Geodesic Corrections for Small Shallow Quantum Circuits ,” Phys.Scripta, 100, 055121 (2025). <https://doi.org/10.1088/1402-4896/add05e>.
- [20] M. Halla, “Modified Conjugate Quantum Natural Gradient ,” 2025. <https://doi.org/10.48550/arXiv.2501.05847>.

A Fubini-Study Metric

To derive the Fubini-Study metric tensor $F_{ij}(\boldsymbol{\theta})$, we investigate the infinitesimal distance between adjacent quantum states in the parameter space. This metric captures the local differential geometry of the quantum state manifold, reflecting how the quantum state $|\psi(\boldsymbol{\theta})\rangle$ responds to infinitesimal parameter variations [5, 7, 8, 9].

We begin with the normalization condition of the quantum state:

$$\langle \psi(\boldsymbol{\theta}) | \psi(\boldsymbol{\theta}) \rangle = 1. \quad (33)$$

Differentiating this condition with respect to θ^i yields:

$$\langle \psi(\boldsymbol{\theta}) | \partial_i \psi(\boldsymbol{\theta}) \rangle + \langle \partial_i \psi(\boldsymbol{\theta}) | \psi(\boldsymbol{\theta}) \rangle = 0. \quad (34)$$

Taking a further derivative with respect to θ^j leads to:

$$\langle \psi(\boldsymbol{\theta}) | \partial_i \partial_j \psi(\boldsymbol{\theta}) \rangle + \langle \partial_i \partial_j \psi(\boldsymbol{\theta}) | \psi(\boldsymbol{\theta}) \rangle + \langle \partial_i \psi(\boldsymbol{\theta}) | \partial_j \psi(\boldsymbol{\theta}) \rangle + \langle \partial_j \psi(\boldsymbol{\theta}) | \partial_i \psi(\boldsymbol{\theta}) \rangle = 0. \quad (35)$$

We now consider the perturbed quantum state $|\psi(\boldsymbol{\theta} + \delta\boldsymbol{\theta})\rangle$, where $\delta\boldsymbol{\theta}$ is a small displacement vector. Expanding this state in a Taylor series yields:

$$|\psi(\boldsymbol{\theta} + \delta\boldsymbol{\theta})\rangle = |\psi(\boldsymbol{\theta})\rangle + \partial_i |\psi(\boldsymbol{\theta})\rangle \delta\theta^i + \frac{1}{2} \partial_i \partial_j |\psi(\boldsymbol{\theta})\rangle \delta\theta^i \delta\theta^j + \mathcal{O}(\delta\theta^3). \quad (36)$$

Taking the inner product with the original state results in:

$$\langle \psi(\boldsymbol{\theta}) | \psi(\boldsymbol{\theta} + \delta\boldsymbol{\theta}) \rangle = 1 + \langle \psi(\boldsymbol{\theta}) | \partial_i \psi(\boldsymbol{\theta}) \rangle \delta\theta^i + \frac{1}{2} \langle \psi(\boldsymbol{\theta}) | \partial_i \partial_j \psi(\boldsymbol{\theta}) \rangle \delta\theta^i \delta\theta^j + \mathcal{O}(\delta\theta^3). \quad (37)$$

The fidelity between these quantum states, accurate to second order in $\delta\boldsymbol{\theta}$, is given by:

$$\begin{aligned} |\langle \psi(\boldsymbol{\theta}) | \psi(\boldsymbol{\theta} + \delta\boldsymbol{\theta}) \rangle|^2 &= 1 + [\langle \psi(\boldsymbol{\theta}) | \partial_i \psi(\boldsymbol{\theta}) \rangle + \langle \partial_i \psi(\boldsymbol{\theta}) | \psi(\boldsymbol{\theta}) \rangle] \delta\theta^i \\ &\quad + [\langle \partial_i \psi(\boldsymbol{\theta}) | \psi(\boldsymbol{\theta}) \rangle \langle \psi(\boldsymbol{\theta}) | \partial_j \psi(\boldsymbol{\theta}) \rangle] \delta\theta^i \delta\theta^j \\ &\quad + \frac{1}{2} [\langle \psi(\boldsymbol{\theta}) | \partial_i \partial_j \psi(\boldsymbol{\theta}) \rangle + \langle \partial_i \partial_j \psi(\boldsymbol{\theta}) | \psi(\boldsymbol{\theta}) \rangle] \delta\theta^i \delta\theta^j + \dots \end{aligned} \quad (38)$$

By incorporating Eqs. (34) and (35), this simplifies to:

$$\begin{aligned} |\langle \psi(\boldsymbol{\theta}) | \psi(\boldsymbol{\theta} + \delta\boldsymbol{\theta}) \rangle|^2 &= 1 + \left[\langle \partial_i \psi(\boldsymbol{\theta}) | \psi(\boldsymbol{\theta}) \rangle \langle \psi(\boldsymbol{\theta}) | \partial_j \psi(\boldsymbol{\theta}) \rangle \right. \\ &\quad \left. - \frac{1}{2} (\langle \partial_i \psi(\boldsymbol{\theta}) | \partial_j \psi(\boldsymbol{\theta}) \rangle + \langle \partial_j \psi(\boldsymbol{\theta}) | \partial_i \psi(\boldsymbol{\theta}) \rangle) \right] \delta\theta^i \delta\theta^j + \dots \end{aligned} \quad (39)$$

Using the small-angle approximation for the Fubini-Study distance between projective rays:

$$d^2(P_\psi, P_\phi) = \arccos^2(|\langle\psi|\phi\rangle|) \approx 1 - |\langle\psi|\phi\rangle|^2 + \mathcal{O}((1 - |\langle\psi|\phi\rangle|^2)^2), \quad (40)$$

we obtain the following expression for the infinitesimal squared distance:

$$\begin{aligned} d^2(P_{\psi(\boldsymbol{\theta})}, P_{\psi(\boldsymbol{\theta}+\delta\boldsymbol{\theta})}) &= \left[\frac{1}{2} (\langle\partial_i\psi(\boldsymbol{\theta})|\partial_j\psi(\boldsymbol{\theta})\rangle + \langle\partial_j\psi(\boldsymbol{\theta})|\partial_i\psi(\boldsymbol{\theta})\rangle) \right. \\ &\quad \left. - \langle\partial_i\psi(\boldsymbol{\theta})|\psi(\boldsymbol{\theta})\rangle \langle\psi(\boldsymbol{\theta})|\partial_j\psi(\boldsymbol{\theta})\rangle \right] d\theta^i d\theta^j. \end{aligned} \quad (41)$$

The first term is evidently real, and the second is also real due to the normalization constraint, which implies:

$$\text{Re} [\langle\psi(\boldsymbol{\theta})|\partial_i\psi(\boldsymbol{\theta})\rangle] = 0. \quad (42)$$

Thus, the infinitesimal squared distance corresponds to the real part of the quantum geometric tensor:

$$d^2(P_{\psi(\boldsymbol{\theta})}, P_{\psi(\boldsymbol{\theta}+\delta\boldsymbol{\theta})}) = \text{Re} [\langle\partial_i\psi(\boldsymbol{\theta})|\partial_j\psi(\boldsymbol{\theta})\rangle - \langle\partial_i\psi(\boldsymbol{\theta})|\psi(\boldsymbol{\theta})\rangle \langle\psi(\boldsymbol{\theta})|\partial_j\psi(\boldsymbol{\theta})\rangle] d\theta^i d\theta^j. \quad (43)$$

This leads directly to the definition of the Fubini-Study metric tensor:

$$F_{ij}(\boldsymbol{\theta}) = \text{Re} [\langle\partial_i\psi(\boldsymbol{\theta})|\partial_j\psi(\boldsymbol{\theta})\rangle - \langle\partial_i\psi(\boldsymbol{\theta})|\psi(\boldsymbol{\theta})\rangle \langle\psi(\boldsymbol{\theta})|\partial_j\psi(\boldsymbol{\theta})\rangle]. \quad (44)$$

One of the most costly traditional methods for estimating the Fubini-Study metric, with a computational complexity of $O(d^2)$, is the parameter-shift rule [10]

$$\begin{aligned} F_{j_1, j_2}(\boldsymbol{\theta}) &= \frac{1}{4} \left[|\langle\psi(\boldsymbol{\theta})|\psi(\boldsymbol{\theta} + (\mathbf{e}_{j_1} + \mathbf{e}_{j_2})\pi/2)\rangle|^2 \right. \\ &\quad - |\langle\psi(\boldsymbol{\theta})|\psi(\boldsymbol{\theta} + (\mathbf{e}_{j_1} - \mathbf{e}_{j_2})\pi/2)\rangle|^2 \\ &\quad - |\langle\psi(\boldsymbol{\theta})|\psi(\boldsymbol{\theta} + (-\mathbf{e}_{j_1} + \mathbf{e}_{j_2})\pi/2)\rangle|^2 \\ &\quad \left. + |\langle\psi(\boldsymbol{\theta})|\psi(\boldsymbol{\theta} - (\mathbf{e}_{j_1} + \mathbf{e}_{j_2})\pi/2)\rangle|^2 \right]. \end{aligned} \quad (45)$$

Here, $\mathbf{e}_j$ denotes the unit vector along the θ_j axis.

B Tensor-averaging error

Remark 1 (Tensor-averaging error). *Strictly speaking, the update of $\bar{\mathbf{F}}_k$ in equation (27) mixes two Riemannian metric tensors that lie in different tangent spaces at $\boldsymbol{\theta}_{k-1}$ and $\boldsymbol{\theta}_k$. A fully geometric treatment would first pull back $\mathbf{F}(\boldsymbol{\theta}_{k-1})$ along the map $\boldsymbol{\theta}_{k-1} \rightarrow \boldsymbol{\theta}_k$ into the tangent space at $\boldsymbol{\theta}_k$, and only then average the two metrics in the same space.*

However, for a small learning rate η , it can be shown that

$$\|\boldsymbol{\theta}_k - \boldsymbol{\theta}_{k-1}\| = \|\eta \mathbf{F}(\boldsymbol{\theta}_{k-1})^{-1} \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1})\| = O(\eta), \quad (46)$$

and if $\mathbf{F}(\boldsymbol{\theta})$ is Lipschitz continuous with constant L , then

$$\|\mathbf{F}(\boldsymbol{\theta}_k) - \mathbf{F}(\boldsymbol{\theta}_{k-1})\| \leq L \|\boldsymbol{\theta}_k - \boldsymbol{\theta}_{k-1}\| = O(\eta). \quad (47)$$

Since the pullback operator itself deviates from the identity by $O(\eta)$ over such a small step, the combined geometric bias—from neglecting the pullback and averaging two nearby metrics—is $O(\eta)$ per step. For typical VQE learning rates ($\eta \ll 1$), this remains negligible compared to stochastic and shot-noise fluctuations.

C Additional Experiments and Figures

In Figure 8, we fix the resampling size $N = 5$ for QNSTEIN2 and QNSTEIN3, while varying the resampling size N for QNSPSA. We observe that QNSPSA requires a larger resampling size (greater than 5) to achieve a convergence rate comparable to the QNSTEIN optimizers. However, this increase in N significantly raises the quantum resources needed for convergence.




Figure 8: The resampling size $N = 5$ is fixed for QNSTEIN2 and QNSTEIN3, while N is varied for QNSPSA. The experimental conditions are the same as in Section 4.1 (12 qubits and 3 layers), and results are averaged over 30 different random initializations.