

Measurement Cost of Metric-Aware Variational Quantum Algorithms

Barnaby van Straaten and Bálint Koczor^{*}

Department of Materials, University of Oxford, Parks Road, Oxford OX1 3PH, United Kingdom



(Received 18 May 2020; accepted 28 June 2021; published 10 August 2021)

We consider metric-aware quantum algorithms that use a quantum computer to efficiently estimate both a matrix and a vector object. For example, the recently introduced quantum natural gradient approach uses the Fisher matrix as a metric tensor to correct the gradient vector for the codependence of the circuit parameters. We rigorously characterize and upper bound the number of measurements required to determine an iteration step to a fixed precision, and propose a general approach for optimally distributing samples between matrix and vector entries. Finally, we establish that the number of circuit repetitions needed for estimating the quantum Fisher information matrix is asymptotically negligible for an increasing number of iterations and qubits.

DOI: [10.1103/PRXQuantum.2.030324](https://doi.org/10.1103/PRXQuantum.2.030324)

I. INTRODUCTION

With quantum computers rising as realistic technologies, attention has turned to how such machines could perform as variational tools [1–24]. This results in a hybrid model with an iterative loop: a classical processor determines how to update the parameters describing a family of quantum states (parameterized ansatz states), while a quantum coprocessor generates and performs measurements on that state (via an ansatz circuit). This is of particular interest in the context of noisy, intermediate-scale quantum devices [25], because complex ansatz states can be prepared with shallow circuits [26–29]. Such shallow circuits will potentially enable obtaining a useful value before the era of resource-intensive quantum fault-tolerance methods. As such, variational quantum algorithms promise to solve key problems that are intractable to classical computers, such as finding ground states [2,4,6,11,30]—as relevant in quantum chemistry and in materials science—or approximately solving combinatorial problems [1] and beyond.

Despite their potential power, variational algorithms might require an extremely large number of quantum-circuit repetitions—optimally using quantum resources will therefore have a crucial economic importance. Attention has recently been focused on statistical aspects of these variational quantum algorithms [31–36], such as the effect of shot noise and the reduction of their measurement costs. It is our aim in this work to establish general

scaling results by rigorously characterizing the number of measurements required to obtain a single iteration step in the case of so-called metric-aware quantum algorithms. Let us first introduce basic notions.

A. Variational quantum algorithms

We consider variational quantum algorithms that typically aim to prepare a parameterized quantum state $\rho(\theta) := \Phi(\theta)\rho_0$ where we model via a mapping $\Phi(\theta)$ that acts on the computational zero state ρ_0 of N qubits and depends continuously on the parameters θ_i with $i \in \{1, 2, \dots, v\}$. This mapping can in general contain nonunitary elements, such as measurements [37,38], but in many applications one assumes that it acts (approximately) as a unitary circuit that decomposes into a product of individual quantum gates. These gates typically act on a small subset of the system, e.g., one- and two-qubit gates.

Recently, a novel variational algorithm was proposed for simulating real-time quantum evolution using shallow quantum circuits [8] and was further generalized to imaginary time and natural gradient evolutions [37,39] that can be used as optimizers of variational quantum eigensolvers (VQEs) [2,4,16,40]. This was shown to significantly outperform other approaches, such as simple gradient descent, in terms of convergence speed and accuracy according to numerical simulations [37,39,41].

In this work, we consider generalizations of the aforementioned techniques as variational algorithms that need to estimate the following two objects: (a) a positive-semidefinite, symmetric matrix, which is usually the quantum Fisher information that characterizes sensitivity with respect to parameters θ_k ; (b) a vector object that is in many applications the gradient vector of the loss function. Examples of such algorithms are provided in Refs. [37,39,

^{*}balint.koczor@materials.ox.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

42–44], and we refer to them in the following as metric-aware quantum algorithms. The metric tensor typically depends only on the parameter values, while the vector object additionally depends on, e.g., a Hermitian observable \mathcal{H} that in typical scenarios represents the Hamiltonian of a physical system and decomposes into a polynomially increasing number r_h of Pauli terms.

B. Quantum natural gradient

To be more concrete, in the following we focus on one prominent algorithm, the recently introduced quantum natural gradient approach [37,44] that is equivalent to imaginary time evolution when quantum circuits are noiseless and unitary [37,39]. This approach can be used as a VQE optimizer when minimizing the expectation value $E(\theta) := \text{Tr}[\rho(\theta)\mathcal{H}]$ over the parameters θ . However, the approach generalizes to any Lipschitz continuous mapping as an objective function [37].

In particular, natural gradient descent governs the evolution of the ansatz parameters according to the update rule [37]

$$\theta(t+1) = \theta(t) - \lambda \mathbf{F}_Q^{-1} \underline{g}, \quad (1)$$

where t is an index and λ is a step size. Here the inverse of the positive-semidefinite, symmetric quantum Fisher information matrix $\mathbf{F}_Q \in \mathbb{R}^{v \times v}$ corrects the gradient vector

$g_k := \partial_k E(\theta)$ for the codependence of the parameters, and both objects can be estimated efficiently using a quantum computer, while the inverse \mathbf{F}_Q^{-1} is computed by a classical processor.

We discuss different protocols for estimating the matrix $[\mathbf{F}_Q]_{kl}$ and vector g_k entries for both pure (idealized, perfect quantum gates) and mixed quantum states (via imperfect quantum gates or nonunitary elements as measurements) in the Appendix. We now highlight two results. (a) We derive the general upper bound $[\mathbf{F}_Q]_{kl} \leq r_g^2$, where r_g is the maximal number of Pauli terms into which generators of ansatz gates can be decomposed (Lemma 1). This bound is a generalization of what is known as the Heisenberg limit in quantum metrology [45] (see also Refs. [46–48]). (b) The matrix \mathbf{F}_Q might be ill conditioned and the inversion in Eq. (1) requires a regularization. We use the simple variant of Tikhonov regularization $\tilde{\mathbf{F}}_Q^{-1} := [\mathbf{F}_Q + \eta \text{Id}]^{-1}$ in the following; we derive analytical lower and upper bounds on the singular values of this inverse matrix in Appendix B 3 (Lemma 3) using (a).

II. UPPER BOUNDS ON THE MEASUREMENT COST

To motivate our approach, we illustrate in Fig. 1 [green line in (a)] how naively using the same number of measurements for estimating each matrix and vector entry,




FIG. 1. Exact numerical simulations: a 12-qubit ansatz circuit with 84 parameters is initialized at a good approximation of the ground state of a spin-chain Hamiltonian (see Appendix A). Natural gradient evolution from Eq. (1) is simulated with a regularization parameter $\eta = 10^{-1}$ (10^{-5}); see the solid (dashed) lines. (a) Measurement overhead κ from Result 1 (red [49]) at every iteration step t of the natural gradient evolution. This quantifies how much more it costs to estimate the natural gradient vector $\underline{v}(t)$ than it would cost to estimate the gradient vector $\underline{g}(t)$ assuming the same precision ϵ . The overhead κ converges to its constant (black) asymptotic approximation. Optimally distributing measurements (blue) via Result 3 significantly reduces sampling costs. However, naively (green [49]) using a fixed number of measurements for estimating each matrix and vector element results in a substantial overhead. (b) Multiplying $\underline{g}(t)$ (red) with the inverse of \mathbf{F}_Q results in $\underline{v}(t)$ (blue) whose norm might be orders of magnitude larger. (c) In practice, a relative precision is required, such that ϵ is proportional to the vector norms; see the text. Carefully setting the regularization parameter η significantly improves the practical applicability: solid lines with $\eta = 10^{-1}$ result in a sampling cost of $\underline{v}(t)$ comparable to (green shows $\kappa = 1$) or even smaller than $\underline{g}(t)$. Refer to Sec. IV for a remark about mitigating the initial high overheads seen in graphs (a) and (c).

such as in Ref. [41], can result in impractical sampling costs.

In particular, we aim to reduce the error due to shot noise (finite sampling) ϵ of the vector $\underline{v} := \tilde{\mathbf{F}}_Q^{-1} \underline{g}$ in the update rule in Eq. (1). We first express how the error in the matrix and vector entries propagates to the parameter-update rule in Eq. (1). We quantify this error as the expected Euclidean distance $\langle \|\Delta v\|^2 \rangle = \epsilon^2$, and this translates to the condition $\sum_{k=1}^v \text{Var}[v_k] = \epsilon^2$, where $\text{Var}[v_k]$ is the variance of a single vector entry.

We derive an analytical formula in Lemma 2 in Appendix B: we express the error ϵ in terms of the variances $\text{Var}\{[\mathbf{F}_Q]_{kl}\}$ and $\text{Var}[g_l]$ of the measurements used to estimate the matrix and vector entries, respectively, as

$$\epsilon^2 = \sum_{k,l=1}^v a_{kl} \text{Var}\{[\mathbf{F}_Q]_{kl}\} + \sum_{k=1}^v b_k \text{Var}[g_k]. \quad (2)$$

The coefficients a_{kl} and b_k describe how the error of $[\mathbf{F}_Q]_{kl}$ and g_k propagates through matrix inversion and subsequent vector multiplication into the precision ϵ . We remark that these results are completely general and can be applied to any quantum algorithm that requires the estimation of both an inverse matrix and a vector object, such as a Hessian-based optimization.

We derive general upper bounds on the variances $\text{Var}\{[\mathbf{F}_Q]_{kl}\}$ and $\text{Var}[g_l]$ for different experimental strategies in the Appendix A; the error ϵ^2 in Eq. (2) is reduced proportionally when repeating measurements. In the following, we assume that N_F measurements are assigned to estimate the full matrix \mathbf{F}_Q , while N_g measurements are used to estimate the gradient vector \underline{g} [49]. We now state an upper bound on them in terms of the precision ϵ .

Theorem 1. *To reduce the uncertainty of the vector $\underline{v} = \tilde{\mathbf{F}}_Q^{-1} \underline{g}$ due to shot noise to a precision ϵ , the number of samples to estimate the matrix \mathbf{F}_Q in Eq. (1) is upper bounded as*

$$N_F \leq 2\epsilon^{-2} v^4 \text{Spc}[\tilde{\mathbf{F}}_Q^{-1}]^2 \|\underline{g}\|_\infty^2 f_F, \quad (3)$$

while sampling the gradient has a cost upper bounded by

$$N_g \leq 2\epsilon^{-2} v^2 \text{Spc}[\tilde{\mathbf{F}}_Q^{-1}] \text{Spc}[\mathcal{H}] f_g. \quad (4)$$

The overall measurement cost of determining the natural gradient vector is $N_F + N_g$. Here $\text{Spc}[A]$ denotes the average squared singular values of a matrix $A \in \mathbb{C}^{d \times d}$ via its Hilbert-Schmidt or Frobenius norm as $\text{Spc}[A] := \|A\|^2/d$ and $\|\underline{g}\|_\infty$ is the absolute largest entry in the gradient vector.

The constant factors f_F and f_g in Theorem 1 are specific to the experimental setup used to estimate the matrix or

vector entries. For example, for $r_g = 1$, the factor simplifies as $f_F \leq 2$. The upper bounds in Theorem 1 crucially depend on the regularization and we prove that $\text{Spc}[\tilde{\mathbf{F}}_Q^{-1}] \leq \eta^{-2}$; see Lemma 3 in Appendix B 3. The product $\text{Spc}[\mathcal{H}] f_g$ is a constant that reflects the complexity of estimating the expected value of the fixed \mathcal{H} (and can be reduced with advanced techniques that simultaneously estimate commuting terms [34,36,50–53]). It is interesting to note that the sampling cost of the gradient vector N_g depends on the metric tensor via $\text{Spc}[\tilde{\mathbf{F}}_Q^{-1}]$ (and vice versa). Let us illustrate this point in an example where one of the entries in $\tilde{\mathbf{F}}_Q^{-1}$ is extremely large in absolute value and therefore via the matrix-vector product it magnifies both the mean and the variance of the gradient entries. Indeed, reducing such a magnified variance to our fixed precision ϵ requires an increased number of measurements in the gradient vector. We finally remark that Theorem 1 is quite general and the upper bounds apply to all metric-aware quantum algorithms [37,39,42–44] up to minor modifications.

We establish in the following that in many cases sampling the gradient vector N_g dominates the overall cost of the natural gradient approach as $N_F + N_g \approx N_g$. Before doing so, let us first bound the sampling cost of the natural gradient vector *relative to* the sampling cost N_{smp} of the gradient vector that would be used in simple gradient descent optimizations. Note that the difference between N_g and N_{smp} is that the latter corresponds to the scenario when we fix the metric tensor as the identity matrix $\mathbf{F}_Q := \text{Id}_v$ and thus the precision is $\epsilon^2 := \langle \|\Delta \underline{g}\|^2 \rangle$.

Theorem 2. *Determining the natural gradient vector to the same precision ϵ as the gradient vector requires a sampling overhead $\kappa := (N_F + N_g)/N_{\text{smp}}$. This overhead is upper bounded in general,*

$$\kappa \leq \eta^{-2} + y \quad \text{and} \quad \kappa \approx \text{Spc}[\tilde{\mathbf{F}}_Q^{-1}] + y,$$

up to the potentially vanishing term $y = N_F/N_{\text{smp}}$, as in Results 1 and 2. Here η is either a regularization parameter or the smallest singular value of \mathbf{F}_Q . The second equality establishes an approximation as a constant factor that is valid, e.g., when the evolution is close to the optimal point.

III. SCALING AS A FUNCTION OF THE ITERATIONS

Theorem 1 establishes that the sampling cost N_F of the matrix \mathbf{F}_Q depends on the norm of the gradient vector, which is expected to decrease polynomially during an optimization. In a typical scenario we expect that, even if initially estimating the matrix dominates the sampling costs, asymptotically sampling the vector \underline{g} dominates the costs.

Result 1. *The upper bound in Theorem 1 results in the growth rate $N_F + N_g = \mathcal{O}[\|g(t)\|_\infty^2] + N_g$ when viewed as a function of iterations or steps t . Assuming polynomial convergence via $\|g(t)\|_\infty = \mathcal{O}(t^{-c})$ with $c > 0$, the natural gradient vector requires only a constant sampling overhead asymptotically as*

$$\kappa = (N_F + N_g)/N_{\text{smp}} = \mathcal{O}(\text{Spc}[\tilde{\mathbf{F}}_Q^{-1}] + t^{-2c}),$$

when compared to the gradient vector via Theorem 2. We remark that convergence is guaranteed under mild continuity conditions [31].

We have numerically simulated the natural gradient evolution from Eq. (1) and determined its overhead κ . This quantifies how much more it costs at every iteration step t to estimate the natural gradient vector $\underline{v}(t)$ than it would cost to estimate the gradient vector $\underline{g}(t)$ assuming the same precision ϵ . Figure 1 [red line in (a)] shows how this sampling overhead converges to its constant asymptotic approximation as the average squared singular values $\text{Spc}[\tilde{\mathbf{F}}_Q^{-1}] \approx 10^6$ ($10^{1.5}$) in Fig. 1 [black line in (a)]. Figure 1 [dashed line in (a)] also demonstrates that under-regularizing the inverse (via $\eta = 10^{-5}$) results in unfeasible sampling costs. In fact, carefully increasing the regularization parameter (as $\eta = 10^{-1}$) reduces the sampling cost by several orders of magnitude without significantly affecting the performance: both evolutions decrease the gradient norm with a similar rate; compare the solid and dashed red lines in Fig. 1(b).

It is striking that the overhead plotted in Fig. 1(a) can be very high initially; while the focus of the present paper is on the asymptotic costs with respect to time and size, it is worth noting that this high initial cost could be straightforwardly mitigated by, e.g., only occasionally updating a low-rank approximation of the metric tensor. This may be expected to have little impact on the convergence rate since in the early phase the advantage of using natural gradient is typically less pronounced.

Recall that Fig. 1(a) via Result 1 assumes a constant precision ϵ throughout the evolution that is not practical. In fact, one would require a relative precision such that $\epsilon = \epsilon_0 \|g(t)\|$ in the case of the gradient vector and $\epsilon = \epsilon_0 \|\underline{v}(t)\|$ in the case of the natural gradient vector, for some fixed ϵ_0 . In particular, using a moderate regularization of the inverse as $\eta = 0.1$, the cost of estimating $\underline{v}(t)$ is comparable or even smaller than estimating $\underline{g}(t)$; see the solid red line in Fig. 1(c).

We finally stress that in Fig. 1(a) we do not actually compare the overall performance of the simple and natural gradient methods, but only their per-iteration (per-epoch) costs. We therefore conclude that the natural gradient optimization requires overall less samples to converge (i.e., asymptotically constant overhead but faster convergence rate) when compared to simple gradient descent; see also

Refs. [37,39,41]. Moreover, we prove in the following that even the significant initial overheads in Figs. 1(a)–1(c) do in many practical applications asymptotically vanish for an increasing number of qubits.

IV. SCALING WITH THE SYSTEM SIZE

Let us now consider how the upper bounds in Theorem 1 scale with the number of qubits N . First, we consider the general growth rate $\nu = \mathcal{O}[Na(N)]$ of the number of parameters ν , where $a(N)$ is the depth of the ansatz circuit. For example, $\text{polylog}(N)$ -depth circuits constitute a very general class of ansätze via $a(N) = \mathcal{O}[x \log(N)^y]$ for some $x, y > 0$. Second, we establish that the spectral quantity scales with the number of qubits as $\text{Spc}[\tilde{\mathbf{F}}_Q^{-1}] = \mathcal{O}[N^{-s} a^{-s}(N)]$ with $0 \leq s \leq 2$; see Lemma 3 in Appendix B 3. Third, if \mathcal{H} decomposes into a number r_h of Pauli terms that grow polynomially (e.g., N^4 in the case of chemistry applications) then we obtain a polynomial growth rate $\text{Spc}[\mathcal{H}]f_g = \mathcal{O}(N^b)$ in Theorem 1 [Eq. (4)] with some $b \geq 1$. We finally obtain the growth rates

$$N_F = \mathcal{O}[N^{4-2s} a^{4-2s}(N) \|g\|_\infty^2], \quad (5)$$

$$N_g = \mathcal{O}[N^{2-s+b} a^{2-s}(N)]. \quad (6)$$

Note that the vector norm $\|g\|_\infty^2$ might in general also depend on the number of qubits, e.g., exponentially vanishing gradients in the case of barren plateaus [54–56] that would result in an exponentially decreasing relative sampling cost of the metric tensor. One may also think of scenarios where the gradient norm grows; however, one could then in practice decrease the inverse precision ϵ^{-1} proportionally as typical at the initial stages of an optimization. To simplify our discussion, we assume that the gradient norm $\|g\|_\infty$ is fixed (bounded), e.g., the evolution is initialized in a close vicinity of the optimal parameters as a good classical guess is known. This also encompasses scenarios where the optimization is near termination approaching a fixed, e.g., chemical, precision. We summarize the resulting measurement cost in the following result.

Result 2. *Assume that the number of Pauli terms in the Hamiltonian grows polynomially, implying that $\text{Spc}[\mathcal{H}]f_g = \mathcal{O}(N^b)$ for some $b \geq 1$, and that the gradient norm $\|g\|_\infty = \mathcal{O}(N^1)$ is bounded. The relative sampling cost of the matrix \mathbf{F}_Q vanishes for general $\text{polylog}(N)$ -depth circuits when $b > 2-s$ and, following Theorem 2, determining the natural gradient vector requires at most a constant overhead asymptotically,*

$$\kappa = (N_F + N_g)/N_{\text{smp}} = \mathcal{O}(\text{Spc}[\tilde{\mathbf{F}}_Q^{-1}] + N^{2-b}),$$

when compared to the gradient vector.

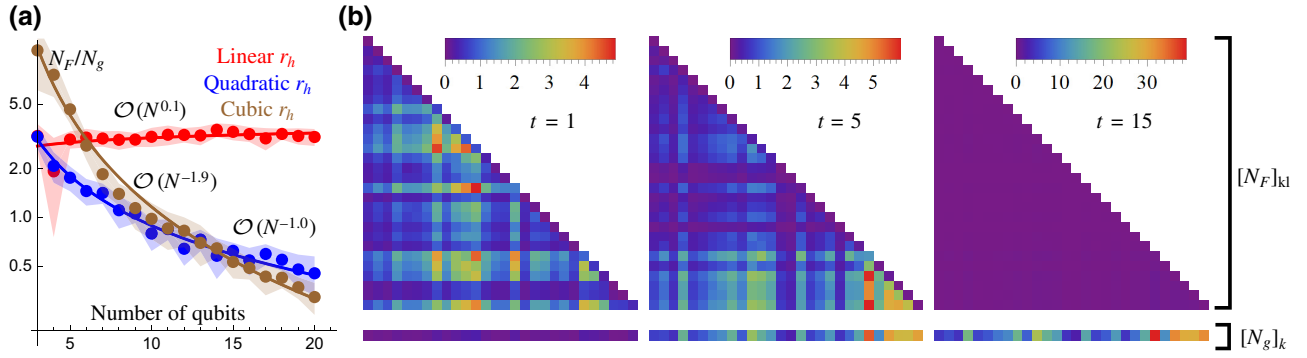


FIG. 2. (a) Sampling cost N_F of the matrix \mathbf{F}_Q relative to the sampling cost N_g of the gradient vector \mathbf{g} from Theorem 1 for an increasing number of qubits determined at randomly selected ansatz parameters with a fixed $\|\mathbf{g}\| = 0.1$. The relative sampling cost N_F/N_g vanishes asymptotically if the number r_h of Pauli terms grows quadratically (blue), cubically (brown), or beyond as established in Eqs. (5)–(6); see also Result 2. Shading represents the standard deviation; see the Appendix. (b) Color maps showing the optimal number of measurements, $[N_F]_{kl}$ and $[N_g]_k$ assigned to individual elements of the Fisher matrix and gradient, respectively, after 1, 5, 15 iterations. For clarity, here we have scaled the data by a constant factor such that if measurements are distributed uniformly then every matrix and vector element would receive one measurement, i.e., we set the total number of measurements N_{opt} to be the number of degrees of freedom. Most of the measurements are assigned to the gradient vector for an increasing number of iterations, as established in Result 1.

Note that Result 2 guarantees a vanishing sampling cost of the matrix \mathbf{F}_Q when the number of terms in the Hamiltonian grows faster than quadratically, i.e., $b > 2$. We have explicitly calculated the growth rates b in the case of three example Hamiltonians in the Appendix A as $b = 1, 2, 3$, respectively, and plot the relative sampling costs N_F/N_g in Fig. 2. We remark that this result can be applied to the general class of metric-aware quantum algorithms [37,39,42–44].

V. OPTIMAL MEASUREMENT DISTRIBUTION

So far, we have assumed that N_F (N_g) measurements are distributed uniformly among the ν^2 (ν) matrix (vector) entries [49]. However, the overall number of samples $N_F + N_g$ (from Theorem 1) needed to obtain the vector $\underline{v} = \tilde{\mathbf{F}}_Q^{-1} \mathbf{g}$ to a precision ϵ can be minimized by distributing samples between the elements of \mathbf{F}_Q and \mathbf{g} optimally [34]. The matrix $[N_F]_{kl}$ and the vector $[N_g]_k$ entries represent the number of measurements assigned to individual elements in \mathbf{F}_Q and in \mathbf{g} , respectively. The number of samples required is reduced to $N_{\text{opt}} = \Sigma^2/\epsilon^2$ with $N_{\text{opt}} \leq N_F + N_g$. We now state explicit expressions for determining Σ , $[N_F]_{kl}$, and $[N_g]_k$.

Result 3. *Measurements are distributed optimally when the number of samples for determining individual elements of the matrix and gradient are given by*

$$[N_F]_{kl} = \epsilon^{-2} \Sigma \sqrt{a_{kl} \text{Var}\{[\mathbf{F}_Q]_{kl}\}}, \quad (7)$$

$$[N_g]_k = \epsilon^{-2} \Sigma \sqrt{b_k \text{Var}[g_k]}, \quad (8)$$

respectively. Here $\text{Var}[\cdot]$ is the variance of a single measurement of the corresponding element and we explicitly define Σ via the coefficients a_{kl} and b_k as

$$\Sigma := \sum_{k,l=1}^{\nu} \sqrt{a_{kl} \text{Var}\{[\mathbf{F}_Q]_{kl}\}} + \sum_{k=1}^{\nu} \sqrt{b_k \text{Var}[g_k]}. \quad (9)$$

Furthermore, the symmetry of the Fisher matrix can be explicitly included just by modifying the coefficients a_{kl} , as discussed in the Appendix C.

We remark that this result is completely general and can be applied to any of the metric-aware quantum algorithms [37,39,42–44].

Figures 1(a) and 1(c) (blue line) shows how the optimal distribution of samples reduces the measurement overhead across the entire evolution—most significantly for small regularization parameters [Figs. 1(a) and 1(c), $\eta = 10^{-5}$], in which case some matrix elements might be crucially larger than others. Moreover, Result 3 automatically takes into account the decreasing sampling cost of the matrix, as established in Results 1 and 2. This is illustrated in Fig. 2(b); for the first few iterations, far from convergence, the bulk of the measurements are directed to the matrix, and comparatively few go to the elements of the gradient [Fig. 2(b), $t = 1$]. However, close to convergence, consistent with Result 1, the gradient takes the majority of the measurements [Fig. 2(b), $t = 20$].

VI. DISCUSSION AND CONCLUSION

In this work we have established general upper bounds on the sampling cost of metric-aware variational quantum

algorithms (e.g., natural gradient). We have analyzed how this sampling cost scales for increasing iterations in Result 1 and for increasing qubit numbers in Result 2. The latter establishes that the relative measurement cost of the matrix object \mathbf{F}_Q is asymptotically negligible in many practically relevant scenarios, such as in case of quantum chemistry applications.

Natural gradient has been shown to outperform other optimization approaches in numerical simulations [37,39,41]. We prove in this work that, for both an increasing number of iterations and number of qubits, the sampling overhead *per iteration (per epoch)* of the natural gradient approach is constant asymptotically when compared to simple gradient descent. The most important implication of our results is therefore that the *overall cost* of natural gradient is lower since it converges to the optimum faster.

We finally establish a general technique that optimally distributes measurements when estimating matrix and vector entries, further reducing the cost of general metric-aware quantum algorithms. Let us finally remark on the generality of our results: our techniques are immediately applicable to other problems beyond metric-aware approaches, for example, to Hessian-based optimizations via Eq. (2) as detailed in Appendix D.

ACKNOWLEDGMENTS

B.K. acknowledges funding received from EU H2020-FETFLAG-03-2018 under Grant No. 820495 (AQTION). The authors thank Simon C. Benjamin and Natalia Ares for their support, stimulating ideas, and useful comments on this manuscript. Numerical simulations in this work used the QuEST and QuESTlink quantum simulation packages [57,58]. The authors would like to acknowledge use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work. We thank Patrick Coles and Andrew Arrasmith for their useful comments.

APPENDIX A: DETERMINING VARIANCES

1. Pauli decompositions

Let us denote the set of Hermitian matrices of dimension d as $\text{Herm}[\mathbb{C}^{d \times d}]$. The Hamiltonian $\mathcal{H} \in \text{Herm}[\mathbb{C}^{d \times d}]$ of a qubit system in general decomposes into a sum over Pauli-operator strings via

$$\mathcal{H} = \sum_{l=1}^{r_h} h_l P_l \quad \text{with} \quad \mathbb{R} \ni h_l := \text{Tr}[\mathcal{H} P_l]/d, \quad (\text{A1})$$

where the $P_l \in \text{Herm}[\mathbb{C}^{d \times d}]$ are tensor products of single-qubit Pauli operators that act on an N -qubit system and form an orthonormal basis of the Hilbert-Schmidt operator space, and $d = 2^N$ is the dimensionality. We denote by $r_h \in \mathbb{N}$ the Pauli rank, i.e., the number of nonzero

Pauli components in the Hamiltonian. Note that, in general, $r_h \leq 4^N$.

In the following derivations we assume for simplicity that ansatz circuits U_c are unitary and decompose into a product of individual gates

$$U_c(\underline{\theta}) = U_v(\theta_v) \cdots U_2(\theta_2) U_1(\theta_1) \quad (\text{A2})$$

that typically act on a small subset of the system, e.g., one- and two-qubit gates. We assume in Eq. (A2) for ease of notation that each quantum gate depends on an individual parameter θ_i with $i = \{1, 2, \dots, v\}$.

Individual gates $U_k(\theta_k) \in SU(d)$ of the quantum circuit from Eq. (A2) are in general of the form $U_k(\theta_k) := \exp[-i\theta_k G_k]$ and their generators $G_k \in \text{Herm}[\mathbb{C}^{d \times d}]$ decompose into a sum of Pauli strings, resulting in

$$U_k(\theta_k) = \exp[-i\theta_k G_k] = \exp\left[-i\theta_k \sum_{l=1}^{r_g^{(k)}} g_{kl} P_l\right] \quad \text{with} \\ \mathbb{R} \ni g_{kl} := \text{Tr}[G_k P_l]/d,$$

where $r_g^{(k)} \in \mathbb{N}$ is the Pauli rank of the generator G_k . We additionally assume that $g_{kl} \leq 1/2$ for simplicity—but any other upper bound could be specified. It follows in general that the derivative $\partial_k U_k(\theta_k)$ decomposes into a sum of $r_g^{(k)}$ unitary operators as

$$\partial_k U_k(\theta_k) = -i \sum_{l=1}^{r_g^{(k)}} g_{kl} P_l U_k(\theta_k). \quad (\text{A3})$$

For ease of notation, in the following we consider circuits via Eq. (A2) that decompose into gates $U_k(\theta_k)$ with Pauli rank $r_g = 1$. This is naturally the case for a wide variety of ansatz circuits, e.g., circuits that consist of single-qubit rotations and two-qubit ZZ or XX evolution gates, as depicted in Fig. 3. This assumption results in a simplified structure of the gates as $U_k(\theta_k) := \exp[-i\theta_k P_k/2]$ and their derivatives as

$$\partial_k U_k(\theta_k) = -\frac{i}{2} P_k U_k(\theta_k), \quad (\text{A4})$$

where P_k is the Pauli generator of the gate $U_k(\theta_k)$. This construction simplifies our following derivations; however, the generalization to arbitrary parameterized gates straightforwardly follows from linearity of Eq. (A3).

We finally define the partial derivative of the circuit in Eq. (A2) using our simplified ansatz as

$$D_k := 2i\partial_k U_c(\underline{\theta}) = U_v(\theta_v) \cdots P_k U_k(\theta_k) \cdots U_2(\theta_2) U_1(\theta_1),$$

which itself is unitary via $[D_k]^\dagger = [D_k]^{-1}$ (and we omit its explicit dependence on the parameters $\underline{\theta}$) and $P_l P_l^\dagger = \text{Id}_d$.

where the second equation uses the decomposition of the Hamiltonian into Pauli operators from Eq. (A1) via denoting the matrix elements $M_{kl} := \text{Im}\langle 0|[D_k]^\dagger P_l U_c|0\rangle$. These matrix elements can be estimated by using an ancilla qubit via the circuits in Fig. 2 of Ref. [8] and the corresponding proof can be found in footnote [53] of Ref. [8]; see also Ref. [43]. The probability p of measuring this ancilla qubit in the $|\pm\rangle$ basis with outcome $+1$ determines the matrix elements via $(2p_{kl}-1) = M_{kl}$ for every Pauli component in the Hamiltonian P_l . This finally yields the explicit form of the gradient vector

$$g_k = \partial_k E(\underline{\theta}) = - \sum_{l=1}^{r_h} h_l (2p_{kl}-1) \quad (\text{A8})$$

in terms of the measurement probabilities $0 \leq p_{kl} \leq 1$. Note that each probability p_{kl} is estimated by sampling a binomial distribution that has a variance $\sigma_{kl}^2 = p_{kl}(1-p_{kl})$. It follows that the variance of the gradient components are determined by these individual variances via

$$\text{Var}[g_k] = 4 \sum_{l=1}^{r_h} h_l^2 \sigma_{kl}^2 = 4 \sum_{l=1}^{r_h} h_l^2 p_{kl}(1-p_{kl}). \quad (\text{A9})$$

Reexpressing this variance in terms of the matrix elements via $p_{kl} = (M_{kl}+1)/2$ yields the simplified form

$$\text{Var}[g_k] = \sum_{l=1}^{r_h} h_l^2 (1-[M_{kl}]^2). \quad (\text{A10})$$

This expression is related directly to the parameterized quantum state $|\psi(\underline{\theta})\rangle$ via the expectation value as $M_{kl} = -2\text{Re}\langle \partial_k \psi(\underline{\theta}) | P_l | \psi(\underline{\theta}) \rangle$. In complete generality, i.e., when gates decompose into a linear combination of at most r_g Pauli terms, the variance of the gradient entries is upper bounded [via Eq. (A9)] as

$$\text{Var}[g_k] \leq r_g \sum_{l=1}^{r_h} h_l^2 = r_g \text{Spc}[\mathcal{H}], \quad (\text{A11})$$

where $\text{Spc}[\mathcal{H}]$ follows from the Hilbert-Schmidt scalar product as

$$\begin{aligned} \text{Spc}[\mathcal{H}] &:= \|\mathcal{H}\|^2/d := \text{Tr}[\mathcal{H}\mathcal{H}]/d \\ &= \sum_{k,l=1}^{r_h} h_k h_l \text{Tr}[P_k P_l]/d = \sum_{l=1}^{r_h} h_l^2, \end{aligned}$$

via Eq. (A1) and recalling that $\text{Tr}[P_k P_l] = d\delta_{kl}$, where δ_{kl} is the Kronecker delta and $d = 2^N$.

So far, we have assumed that each term in the Hamiltonian is estimated separately from outcomes of independent

ancilla measurements and the above variance therefore corresponds to overall r_h measurements. Indeed, advanced techniques could be used for simultaneously measuring commuting terms in the Hamiltonian (possibly without an ancilla qubit), reducing the overall number of shots [34,36,50–53] and we would like to take this into account in our final result. We conclude by stating the upper bound on the variance $\text{Var}[g_k]$ of a single measurement to estimate the gradient entry g_k as

$$\text{Var}[g_k] \leq \text{Spc}[\mathcal{H}] f_g. \quad (\text{A12})$$

Here we have introduced the constant factor f_g . We can generally state the bounds $1 \leq f_g \leq r_g r_h$ as f_g depends on the system (type of gates via r_g) and on the measurement technique used for estimating terms in the Hamiltonian (number of commuting groups). Here the lower bound (best-case scenario $f_g = 1$) is saturated for Pauli gates ($r_g = 1$) and Hamiltonians from Eq. (A1) in which all terms commute and are measured simultaneously. The upper bound (worst-case scenario) is saturated by Hamiltonians from Eq. (A1) in which all r_h terms are estimated from separate measurements (all terms are noncommuting) and all terms have comparable strengths (optimally distributing samples does not reduce $\text{Var}[g_k]$). The factor f_g interpolates between these two extremal cases and will correspond to a value in the bounded range $1 \leq f_g \leq r_g r_h$. In most of this work we assume a fixed \mathcal{H} and therefore we can treat f_g and $\text{Spc}[\mathcal{H}]$ as constants. The only exception is our derivation in Result 2 where we make the mild, general assumption that the number r_h of terms in the Hamiltonian grows polynomially and therefore necessarily the product $\text{Spc}[\mathcal{H}] f_g = \mathcal{O}(N^b)$ grows in some polynomial order b with the number of qubits N . To illustrate this, we construct three example Hamiltonians in Appendix A5 and explicitly compute the polynomial order b in which the cost $\text{Spc}[\mathcal{H}] f_g$ of estimating g_k grows with the number of qubits N .

Let us now consider mixed quantum states, e.g., due to gate imperfections, via the eigendecomposition $\rho = \sum_n p_n |\psi_n\rangle\langle\psi_n|$. If the parametrization $\underline{\theta}$ is approximately unitary via $\partial p_n / \partial \theta_k \approx 0$ then gradient components of the expectation value $\text{Tr}[\rho(\underline{\theta})\mathcal{H}]$ can be expressed as

$$\begin{aligned} \frac{\partial}{\partial \theta_k} \text{Tr}[\rho(\underline{\theta})\mathcal{H}] &\approx \sum_n p_n \frac{\partial}{\partial \theta_k} [\langle \psi_n(\underline{\theta}) | \mathcal{H} | \psi_n(\underline{\theta}) \rangle] \\ &= \sum_n p_n [g_k]_n, \end{aligned} \quad (\text{A13})$$

where $[g_k]_n$ is the gradient that would be measured by the above protocol for the pure eigenstate $|\psi_n(\underline{\theta})\rangle$. The above discussed protocol therefore estimates the correct gradient for mixed states—as long as the parametrization is approximately unitary, such as in case of noisy gates. The same

upper bound holds for the variances via $\sum_n p_n = 1$ and $0 \leq p_n \leq 1$, and the bound is only saturated by pure states.

In summary, the variance of the gradient entries is upper bounded as $\text{Var}[g_k] \leq \text{Spc}[\mathcal{H}]f_g$, where f_g is a constant factor that depends only on the ansatz structure, on the particular quantum algorithm that is used to estimate the entries, and on the Hamiltonian. We remark that the above discussed protocol is used in other metric-aware quantum algorithms, and our bounds therefore apply to other vector objects used in these algorithms [37,39,42–44].

4. Components of the quantum Fisher information matrix

We now focus on determining variances of the quantum Fisher information entries $[\mathbf{F}_Q]_{kl}$. For pure states as $\rho = |\psi\rangle\langle\psi|$, entries of the quantum Fisher information can be expressed via the state-vector scalar products [37]

$$[\mathbf{F}_Q]_{kl} = 4\text{Re}[\langle\partial_k\psi|\partial_l\psi\rangle - \langle\partial_k\psi|\psi\rangle\langle\psi|\partial_l\psi\rangle]. \quad (\text{A14})$$

The second term in the above equation vanishes when the global phase evolution of $|\psi\rangle$ is zero [43], and an experimental protocol for measuring the remaining component $\text{Re}\langle\partial_k\psi|\partial_l\psi\rangle$ was used in Ref. [39] for simulating imaginary time evolution.

We now propose a protocol that determines both terms in Eq. (A14). Assuming the simplified ansatz from Appendix A 1, our protocol allows us to evaluate the coefficients by measuring an ancilla qubit

$$\begin{aligned} A_{kl} &= 4\text{Re}\langle\partial_k\psi|\partial_l\psi\rangle = \text{Re}\langle 0|[D_k]^\dagger D_l|0\rangle = 2[p_a]_{kl} - 1, \\ B_k &= 2\text{Re}\langle\partial_k\psi|\psi\rangle = \text{Re}\langle 0|[D_k]^\dagger U_c|0\rangle = 2[p_b]_k - 1, \\ C_k &= 2\text{Im}\langle\partial_k\psi|\psi\rangle = \text{Im}\langle 0|[D_k]^\dagger U_c|0\rangle = 2[p_c]_k - 1, \end{aligned}$$

using the circuits in Fig. 2 of Ref. [8]; see footnote [53] of Ref. [8] for a proof. These circuits allow for estimating the probabilities p_a, p_b , and p_c by sampling the ancilla qubit as a binomial distribution. The quantum Fisher information is then obtained as

$$\begin{aligned} [\mathbf{F}_Q]_{kl} &= A_{kl} + B_k B_l - C_k C_l = (2[p_a]_{kl} - 1) \\ &\quad + (2[p_b]_k - 1)(2[p_b]_l - 1) \\ &\quad - (2[p_c]_k - 1)(2[p_c]_l - 1). \end{aligned}$$

Since the probabilities p_a, p_b , and p_c are determined from binomial distributions, their variances are given by, e.g., $[\sigma_a^2]_{kl} = [p_a]_{kl}(1 - [p_a]_{kl})$. It follows that

$$\begin{aligned} \text{Var}\{[\mathbf{F}_Q]_{kl}\} &= 4[\sigma_a^2]_{kl} + 4[\sigma_b^2]_k B_l^2 + 4[\sigma_b^2]_l B_k^2 \\ &\quad + 4[\sigma_c^2]_k C_l^2 + 4[\sigma_c^2]_l C_k^2. \end{aligned}$$

Substituting $4[\sigma_b^2]_k = (1 - [B_l]^2)$ and $4[\sigma_c^2]_k = (1 - [C_l]^2)$, we can express the variances as

$$\begin{aligned} \text{Var}\{[\mathbf{F}_Q]_{kl}\} &= (1 - [A_{kl}]^2) + (1 - [B_k]^2)B_l^2 + (1 - [B_l]^2)B_k^2 \\ &\quad + (1 - [C_k]^2)C_l^2 + (1 - [C_l]^2)C_k^2, \end{aligned}$$

in terms of the estimated quantities A_{kl}, B_k , and C_k , where we have used the notation $(A_{kl} + 1)/2 = [p_a]_{kl}$ for example.

Note that the inequality $(1 - [B_k]^2)B_l^2 \leq 1/4$ is saturated when $B_k = 1/\sqrt{2}$ and in general $|A_{kl}|, |B_l|, |C_l| \leq 1$. Using this inequality, we can establish the general upper bound

$$\text{Var}\{[\mathbf{F}_Q]_{kl}\} \leq 2r_g^2, \quad (\text{A15})$$

when gates decompose into a linear combination of at most r_g Pauli terms.

When assuming noisy unitary circuits, Result 3 of Ref. [37] establishes that $[\mathbf{F}_Q]_{kl} \approx 2\text{Tr}[(\partial_k\rho)(\partial_l\rho)]$ and the approximation becomes exact for pure states as $\rho = |\psi\rangle\langle\psi|$. The Hilbert-Schmidt scalar products $\text{Tr}[(\partial_k\rho)(\partial_l\rho)]$ can be measured using the circuit based on SWAP tests from Ref. [43] and one can directly estimate the quantity $[\mathbf{F}_Q]_{kl} = 2p_{kl} - 1$ by measuring the probability p_{kl} of an ancilla qubit in the case when using the simplified ansatz from Appendix A 1, i.e., when gates decompose into single Pauli terms. We remark that this implementation requires more qubits when compared to the above introduced pure-state approach. However, it is preferable as it results in negligible approximation errors when gates are imperfect; see Ref. [37]. The variance follows as $\text{Var}\{[\mathbf{F}_Q]_{kl}\} = 4p_{kl}(1 - p_{kl}) = (1 - [\mathbf{F}_Q]_{kl}^2) \leq 1$ in the case of the simplified ansatz from Appendix A 1 and we have used $[\mathbf{F}_Q]_{kl} \leq 1$ from Lemma 1.

In complete generality, i.e., when gates decompose into a linear combination of at most r_g Pauli terms, the variance of the matrix entries is upper bounded as

$$\text{Var}\{[\mathbf{F}_Q]_{kl}\} \leq r_g^2. \quad (\text{A16})$$

In summary, the variance of the matrix entries are upper bounded as $\text{Var}\{[\mathbf{F}_Q]_{kl}\} \leq f_F$, where f_F is a constant factor that depends only on the ansatz structure and the approach used to estimate the matrix entries. We remark that the above discussed two protocols are used in other metric-aware quantum algorithms and our bounds therefore apply to other matrix objects estimated by these algorithms [37,39,42–44].

5. Numerical simulations

In our numerical simulations we use the ansatz illustrated in Fig. 3. This decomposes into repeated blocks. The first block B_1 consists of single-qubit X rotations, while the second block B_2 decomposes into nearest-neighbor

Pauli ZZ gates followed by single qubit Y and X rotations. Each gate depends on an individual parameter θ_k with $k \in \{1, \dots, v\}$. In our numerical simulations we use the ansatz structure $B_1 B_2 B_2$ that has a linearly growing number of parameters $v = \mathcal{O}(N)$ in the number of qubits via the constant depth $a(N) = \mathcal{O}(N^0)$.

In Fig. 1 we simulate the natural gradient approach for finding the ground-state energy of the spin-chain Hamiltonian

$$\mathcal{H} = \sum_{i=1}^{N-1} J[\sigma_x^{(i)} \sigma_x^{(i+1)} + \sigma_y^{(i)} \sigma_y^{(i+1)} + \sigma_z^{(i)} \sigma_z^{(i+1)}] + J[\sigma_x^{(1)} \sigma_x^{(N)} + \sigma_y^{(1)} \sigma_y^{(N)} + \sigma_z^{(1)} \sigma_z^{(N)}] + \sum_{i=1}^N \omega_i \sigma_z^{(i)}, \quad (\text{A17})$$

which contains identical couplings xx , yy , and zz between nearest neighbors with a constant that we set as $J = 1$. Here the $\sigma_\alpha^{(k)}$ represent Pauli matrices acting on qubit k with $\alpha = \{x, y, z\}$. We select on-site frequencies ω_i randomly according to a uniform distribution with values varying between -1 and 1 . The resulting Hamiltonian has a nontrivial, highly entangled ground state that we aim to approximate using the (not necessarily optimal) ansatz circuit shown in Fig. 3. We initialize the optimization at a point in parameter space close to the optimum and we set the step size as $\lambda = 0.2$.

In Fig. 2 we simulate various different Hamiltonians using the same technique. In particular, we use Eq. (A17) as the linearly scaling Hamiltonian in Fig. 2 (red). We define the quadratically scaling Hamiltonian in Fig. 2 (blue) as

$$\mathcal{H} = \sum_{k>l=1}^N J[\sigma_x^{(k)} \sigma_x^{(l)} + \sigma_y^{(k)} \sigma_y^{(l)} + \sigma_z^{(k)} \sigma_z^{(l)}] + \sum_{k=1}^N \omega_k \sigma_z^{(k)}, \quad (\text{A18})$$

while we chose the cubically scaling Hamiltonian in Fig. 2 (brown) as

$$\mathcal{H} = \sum_{l>k}^N \sum_{m>l}^N J \sigma_x^{(k)} \sigma_y^{(l)} \sigma_z^{(m)} + \sum_{k=1}^N \omega_k \sigma_z^{(k)}. \quad (\text{A19})$$

In our simulations we start the optimization at a random initial point in parameter space, i.e., $\underline{\theta}$ is selected randomly, and run the optimization until the gradient vector is such that $\|\underline{v}\| \approx 10^{-1}$. This ensures that we approximately randomly select points in parameter space for which the gradient norm is fixed, hence satisfying our assumption in Result 2. We compute the values of N_F and N_g at 25 instances of such randomly selected ansatz parameters. The dots, shading, and solid lines in Fig. 2 respectively

show the average, the standard deviation, and the fitting of the ratio N_F/N_g .

Let us now compute how the product $\text{Spc}[\mathcal{H}]f_g = \mathcal{O}(N^b)$ from Result 2 [which reflects the cost of estimating a gradient entry g_k via Eq. (A12)] grows with the number of qubits. Terms in the above Hamiltonians can be grouped into a constant number of commuting groups that can be measured simultaneously and, therefore, $f_g = \mathcal{O}(1)$ in Eq. (A12). Furthermore, the squared sum of the coefficients grows as $\text{Spc}[\mathcal{H}] = \sum_{l=1}^{r_h} h_l^2 = \mathcal{O}(N^b)$ with $b = 1, 2, 3$. We therefore conclude that our upper bound in Eq. (A12) grows as $\text{Var}[g_k] \leq \text{Spc}[\mathcal{H}]f_g = \mathcal{O}(N^b)$ and, indeed, the product in Result 2 grows as $\text{Spc}[\mathcal{H}]f_g = \mathcal{O}(N^b)$ with $b = 1, 2, 3$.

APPENDIX B: PROPAGATING VARIANCES

Lemma 2. *Let us define the regularized inverse $\tilde{\mathbf{F}}_Q^{-1} := [\mathbf{F}_Q + \eta \text{Id}]^{-1}$ of the Fisher information matrix for some regularization parameter $\eta \geq 0$ and recall that we have defined the error measure $\epsilon^2 := \sum_{k=1}^v \text{Var}[v_k]$ with $\underline{v} := \tilde{\mathbf{F}}_Q^{-1} \underline{g}$ in the main text. If the elements of \mathbf{F}_Q and \underline{g} are measured independently and their errors are sufficiently small then the error measure can be written in the form*

$$\epsilon^2 = \sum_{k,l=1}^v a_{kl} \text{Var}\{[\mathbf{F}_Q]_{kl}\} + \sum_{k=1}^v b_k \text{Var}[g_l], \quad (\text{B1})$$

$$\text{where } a_{kl} := \sum_{i,j=1}^v [\tilde{\mathbf{F}}_Q^{-1}]_{ik}^2 [\tilde{\mathbf{F}}_Q^{-1}]_{lj}^2 g_l^2,$$

$$b_k := \sum_{l=1}^v \{[\tilde{\mathbf{F}}_Q^{-1}]_{kl}\}^2. \quad (\text{B2})$$

Proof. Under the assumption that the elements are measured independently and are sufficiently small, it is appropriate to use the variance formula [59]; thus, we can write the error measure in terms of the variance of the elements in $\tilde{\mathbf{F}}_Q^{-1}$ and \underline{g} , yielding

$$\begin{aligned} \epsilon^2 &= \sum_{k=1}^v \text{Var}[v_k] \\ &= \sum_{k,l=1}^v \text{Var}\{[\tilde{\mathbf{F}}_Q^{-1}]_{kl}\} g_l^2 + \{[\tilde{\mathbf{F}}_Q^{-1}]_{kl}\}^2 \text{Var}[g_l]. \end{aligned} \quad (\text{B3})$$

Now we use the result derived in Ref. [60] to relate the variance of elements of $\tilde{\mathbf{F}}_Q^{-1}$ to elements of \mathbf{F}_Q , namely,

$$\text{Var}\{[\tilde{\mathbf{F}}_Q^{-1}]_{kl}\} = \sum_{i,j=1}^v [\tilde{\mathbf{F}}_Q^{-1}]_{ik}^2 \text{Var}\{[\mathbf{F}_Q]_{kl}\} [\tilde{\mathbf{F}}_Q^{-1}]_{lj}^2. \quad (\text{B4})$$

Here we have used the fact that $\text{Var}\{\tilde{\mathbf{F}}_Q^{-1}\} = \text{Var}\{\mathbf{F}_Q\}$. Substituting this result into the error metric and trivially rearranging yields

$$\begin{aligned} \epsilon^2 &= \sum_{k,l=1}^v \left[\sum_{i,j=1}^v [\tilde{\mathbf{F}}_Q^{-1}]_{ik}^2 \text{Var}\{[\mathbf{F}_Q]_{kl}\} [\tilde{\mathbf{F}}_Q^{-1}]_{lj}^2 \right] g_l^2 \\ &\quad + \{[\tilde{\mathbf{F}}_Q^{-1}]_{kl}\}^2 \text{Var}[g_l] \\ &= \sum_{k,l=1}^v \underbrace{\left[\sum_{i,j=1}^v [\tilde{\mathbf{F}}_Q^{-1}]_{ik}^2 [\tilde{\mathbf{F}}_Q^{-1}]_{lj}^2 g_l^2 \right]}_{a_{kl}} \text{Var}\{[\mathbf{F}_Q]_{kl}\} \\ &\quad + \sum_{k=1}^v \underbrace{\left[\sum_{l=1}^v \{[\tilde{\mathbf{F}}_Q^{-1}]_{kl}\}^2 \right]}_{b_k} \text{Var}[g_k], \end{aligned} \quad (\text{B5})$$

as required. \blacksquare

1. Proof of Theorem 1

Recall that Lemma 2 establishes the error propagation formula that we abbreviate as $\epsilon^2 = \epsilon_F^2 + \epsilon_g^2$ via

$$\epsilon_F^2 := \sum_{\alpha,\beta=1}^v a_{\alpha\beta} \text{Var}\{[\mathbf{F}_Q]_{\alpha\beta}\}, \quad \epsilon_g^2 := \sum_{l=1}^v b_l \text{Var}[g_l]. \quad (\text{B6})$$

The coefficients $a_{\alpha\beta}$ can be upper bounded as

$$\begin{aligned} a_{\alpha\beta} &= \sum_{k,l=1}^v g_l^2 [\tilde{\mathbf{F}}_Q^{-1}]_{k\alpha}^2 [\tilde{\mathbf{F}}_Q^{-1}]_{l\beta}^2 \\ &\leq \|g\|_\infty^2 \sum_{k=1}^v [\tilde{\mathbf{F}}_Q^{-1}]_{k\alpha}^2 \sum_{l=1}^v [\tilde{\mathbf{F}}_Q^{-1}]_{l\beta}^2 \quad \text{and} \\ b_l &= \sum_{k=1}^v [\tilde{\mathbf{F}}_Q^{-1}]_{kl}^2, \end{aligned}$$

where $\|g\|_\infty$ is the absolute largest element in the gradient vector. We assume that every matrix and vector element is assigned measurements uniformly as N_F/ν^2 and N_g/ν , where N_F and N_g are the overall numbers of measurements required to estimate the matrix and vector objects such that the vector \underline{v} is obtained to a precision ϵ . Using the upper bounds on the variances of individual gradient vector entries from Eq. (A11) and individual matrix entries from Eqs. (A16) and (A15), we derive the explicit bound

$$\begin{aligned} \text{Var}\{[\mathbf{F}_Q]_{\alpha\beta}\} &\leq V_F := \nu^2 N_F^{-1} f_F, \\ \text{Var}[g_l] &\leq V_G := \nu N_g^{-1} \text{Spc}[\mathcal{H}] f_g, \end{aligned}$$

where $\|\mathcal{H}\|$ is the Hilbert-Schmidt or Frobenius norm of the Hamiltonian and f_F, f_g are constant factors that depend

on the ansatz structure and the approach used to estimate the gradient or Fisher matrix; see Appendix A 1. For example, for the simplified ansatz ($r_g = 1$) in Appendix A 1, we obtain $f_F \leq 2$ and $1 \leq f_g \leq r_h$. Here the lower bound $f_g = 1$ is saturated when all terms in the Hamiltonian from Eq. (A1) commute and can be measured simultaneously, while the upper bound $f_g = r_h$ is saturated when all r_h terms in the Hamiltonian need to be estimated independently (because they do not commute) and their strengths are comparable; see Appendix A 3.

We use the above derived upper bounds and obtain

$$\epsilon_F^2 \leq V_F \|g\|_\infty^2 \sum_{\alpha,k=1}^v [\tilde{\mathbf{F}}_Q^{-1}]_{k\alpha}^2 \sum_{\beta,l=1}^v [\tilde{\mathbf{F}}_Q^{-1}]_{l\beta}^2 = V_F \|g\|_\infty^2 \|\tilde{\mathbf{F}}_Q^{-1}\|^4, \quad (\text{B7})$$

$$\epsilon_g^2 \leq V_G \sum_{k,l=1}^v [\tilde{\mathbf{F}}_Q^{-1}]_{kl}^2 = V_G \|\tilde{\mathbf{F}}_Q^{-1}\|^2, \quad (\text{B8})$$

where $\|\tilde{\mathbf{F}}_Q^{-1}\|$ is the Hilbert-Schmidt or Frobenius norm of the inverse matrix $\tilde{\mathbf{F}}_Q^{-1}$.

We now require that $\epsilon^2/2 =: \epsilon_F^2$ and $\epsilon^2/2 =: \epsilon_g^2$ as a possible choice to satisfy $\epsilon^2 = \epsilon_F^2 + \epsilon_g^2$. This results in the explicit bound on the number of measurements after substituting V_F and V_G as

$$\begin{aligned} N_F &\leq 2\nu^2 \|g\|_\infty^2 \|\tilde{\mathbf{F}}_Q^{-1}\|^4 \epsilon^{-2} f_F, \\ N_g &\leq 2\nu \|\tilde{\mathbf{F}}_Q^{-1}\|^2 \epsilon^{-2} \text{Spc}[\mathcal{H}] f_g. \end{aligned}$$

We introduce the notation $\text{Spc}[\tilde{\mathbf{F}}_Q^{-1}] := \|\tilde{\mathbf{F}}_Q^{-1}\|^2/\nu = (1/\nu) \sum_{k=1}^v \sigma_k^2(\tilde{\mathbf{F}}_Q^{-1})$ to denote the average of the squared singular values of $\tilde{\mathbf{F}}_Q^{-1}$. Note that, for example, the identity matrix yields $\text{Spc}[\text{Id}] = 1$ and we derive upper and lower bounds in general in Lemma 3.

We finally establish the upper bounds

$$\begin{aligned} N_F &\leq 2\nu^4 \|g\|_\infty^2 \text{Spc}[\tilde{\mathbf{F}}_Q^{-1}]^2 \epsilon^{-2} f_F, \\ N_g &\leq 2\nu^2 \text{Spc}[\tilde{\mathbf{F}}_Q^{-1}] \epsilon^{-2} \text{Spc}[\mathcal{H}] f_g. \end{aligned}$$

2. Proof of Theorem 2

Recall that in Appendix B 1 we defined the precision associated with the gradient vector in the natural gradient approach as $\epsilon_g^2 := \sum_{l=1}^v b_l \text{Var}[g_l]$. We also defined the total number of measurements N_g that needs to be assigned to determine the gradient vector \underline{g} to a precision ϵ_g^2 as

$$N_g := \frac{\nu}{\epsilon_g^2} \sum_{l=1}^v b_l \text{Var}[g_l],$$

since each gradient entry g_l receives N_g/ν samples.

In the limiting case $\tilde{\mathbf{F}}_Q^{-1} \rightarrow \text{Id}$ and $\text{Var}\{[\mathbf{F}_Q]_{kl}\} \rightarrow 0$ the natural gradient approach reduces to the simple gradient descent approach with $b_l = 1$. We can therefore define the total number of measurements N_{smpI} required to reconstruct the gradient vector in the simple gradient descent approach via $b_l = 1$ as

$$N_{\text{smpI}} := \frac{\nu}{\epsilon_g^2} \sum_{l=1}^{\nu} \text{Var}[g_l].$$

Let us start by explicitly writing the ratio of measurements as

$$\frac{N_g}{N_{\text{smpI}}} = \frac{\sum_{l=1}^{\nu} b_l \text{Var}[g_l]}{\sum_{l=1}^{\nu} \text{Var}[g_l]},$$

and let us consider the term

$$\begin{aligned} b_l &= \sum_{k=1}^{\nu} [\tilde{\mathbf{F}}_Q^{-1}]_{kl}^2 = \|\text{Col}_l[\tilde{\mathbf{F}}_Q^{-1}]\|^2 \\ &= \|\tilde{\mathbf{F}}_Q^{-1} B_l\|^2 \leq \|\tilde{\mathbf{F}}_Q^{-1}\|_{\infty}^2 = \sigma_{\max}(\tilde{\mathbf{F}}_Q^{-1})^2 \leq \eta^{-2}, \end{aligned}$$

where η is either a regularization parameter or the smallest singular value of \mathbf{F}_Q , $\text{Col}_l[\tilde{\mathbf{F}}_Q^{-1}]$ denotes the l th column vector of the matrix $\tilde{\mathbf{F}}_Q^{-1}$, and B_l is the l th standard basis vector with $\|B_l\| = 1$. Our general upper bound follows as

$$\frac{N_g}{N_{\text{smpI}}} = \frac{\sum_{l=1}^{\nu} b_l \text{Var}[g_l]}{\sum_{l=1}^{\nu} \text{Var}[g_l]} \leq \eta^{-2} \frac{\sum_{l=1}^{\nu} \text{Var}[g_l]}{\sum_{l=1}^{\nu} \text{Var}[g_l]} = \eta^{-2}.$$

We now establish an approximation under the assumption that $\text{Var}[g_l]$ does not significantly depend on the index l , e.g., when the gradient is vanishing close to an optimal point via $M_{kl} \rightarrow 0$ in Eq. (A10) as, e.g.,

$$\text{Var}[g_k] = \sum_{l=1}^{r_h} h_l^2 (1 - [M_{kl}]^2) \rightarrow \sum_{l=1}^{r_h} h_l^2 = \text{Spc}[\mathcal{H}]. \quad (\text{B9})$$

This results in

$$\begin{aligned} \frac{N_g}{N_{\text{smpI}}} &= \frac{\sum_{l=1}^{\nu} b_l \text{Var}[g_l]}{\sum_{l=1}^{\nu} \text{Var}[g_l]} \approx \frac{\sum_{l=1}^{\nu} b_l}{\nu} \\ &= \frac{\sum_{k,l=1}^{\nu} [\tilde{\mathbf{F}}_Q^{-1}]_{kl}^2}{\nu} = \frac{\|\tilde{\mathbf{F}}_Q^{-1}\|}{\nu} = \text{Spc}[\tilde{\mathbf{F}}_Q^{-1}]. \end{aligned}$$

3. Remarks on Theorem 2

We establish bounds in the case of the relative-precision scheme, i.e., when $\epsilon \propto \|g(t)\|$ and $\epsilon \propto \|v(t)\|$ in the case of the gradient and natural gradient vectors, respectively.

The upper bound follows via

$$\frac{\|g\|^2}{\|v\|^2} = \frac{\|g\|^2}{\|\tilde{\mathbf{F}}_Q^{-1} g\|^2} \leq \sigma_{\min}(\tilde{\mathbf{F}}_Q^{-1})^{-2}, \quad (\text{B10})$$

and a lower bound can be specified as

$$\frac{\|g\|^2}{\|v\|^2} = \frac{\|g\|^2}{\|\mathbf{F}_Q^{-1} g\|^2} \geq \sigma_{\max}(\tilde{\mathbf{F}}_Q^{-1})^{-2}, \quad (\text{B11})$$

and in complete generality

$$\frac{N_g}{N_{\text{smpI}}} \frac{\|g\|^2}{\|v\|^2} \leq \left[\frac{\sigma_{\max}(\tilde{\mathbf{F}}_Q^{-1})}{\sigma_{\min}(\tilde{\mathbf{F}}_Q^{-1})} \right]^2 =: \text{Cnd}[\mathbf{F}_Q^{-1}]^2, \quad (\text{B12})$$

and Lemma 3 establishes that $\text{Cnd}[\mathbf{F}_Q^{-1}] \leq \eta^{-1}(\nu r_g + \eta)$.

Lemma 3. Assuming the simple regularization $\tilde{\mathbf{F}}_Q^{-1} := (\mathbf{F}_Q + \eta \text{Id}_\nu)^{-1}$, the largest singular value of the inverse is upper bounded as $\sigma_{\max}(\tilde{\mathbf{F}}_Q^{-1}) \leq \eta^{-1}$ and the smallest singular value is lower bounded via $\sigma_{\min}(\tilde{\mathbf{F}}_Q^{-1}) \geq (\nu r_g + \eta)^{-1}$. Moreover, the bounds $(\nu r_g + \eta)^{-2} \leq \text{Spc}[\tilde{\mathbf{F}}_Q^{-1}] \leq \eta^{-2}$ and $\text{Cnd}[\tilde{\mathbf{F}}_Q^{-1}] \leq \eta^{-1}(\nu r_g + \eta)$ hold in general. Here r_g is the largest Pauli rank of the ansatz gates from Appendix A1.

Proof. It immediately follows that

$$\sigma_{\max}[(\mathbf{F}_Q + \eta \text{Id}_\nu)^{-1}] = [\sigma_{\min}(\mathbf{F}_Q + \eta \text{Id}_\nu)]^{-1} \leq \eta^{-1}$$

via $\sigma_{\min}(\mathbf{F}_Q + \eta \text{Id}_\nu) \geq \eta$. Now we use the boundedness of the matrix elements as $|[\mathbf{F}_Q^{-1}]_{kl}| \leq r_g^2$ from Lemma 1 that establishes the matrix norm $\|\mathbf{F}_Q\|_{\max} := \max_{k,l} |[\mathbf{F}_Q]_{kl}| \leq r_g^2$. This bounds the largest singular value of \mathbf{F}_Q as

$$r_g^2 \geq \|\mathbf{F}_Q\|_{\max} \geq \|\mathbf{F}_Q\|_{\infty}/\nu := \sigma_{\max}(\mathbf{F}_Q)/\nu.$$

The smallest singular value of the inverse is therefore bounded as

$$\begin{aligned} \sigma_{\min}[(\mathbf{F}_Q + \eta \text{Id}_\nu)^{-1}] &= [\sigma_{\max}(\mathbf{F}_Q + \eta \text{Id}_\nu)]^{-1} \\ &\geq (\nu r_g^2 + \eta)^{-1}. \end{aligned}$$

We can now establish the bound

$$\begin{aligned} (\nu r_g^2 + \eta)^{-2} &\leq \sigma_{\min}^2(\tilde{\mathbf{F}}_Q^{-1}) \leq \text{Spc}[\tilde{\mathbf{F}}_Q^{-1}] \\ &\leq \sigma_{\max}^2(\tilde{\mathbf{F}}_Q^{-1}) \leq \eta^{-2}. \end{aligned} \quad (\text{B13})$$

We can therefore bound the growth rate of the quantity $\text{Spc}[\tilde{\mathbf{F}}_Q^{-1}]$ as $\text{Spc}[\tilde{\mathbf{F}}_Q^{-1}] = \mathcal{O}(\nu^s)$ with $-2 \leq s \leq 0$. ■

APPENDIX C: OPTIMAL MEASUREMENTS

Lemma 4. *Measurements are distributed optimally when the number of samples for determining individual elements of the matrix and gradient are given by*

$$[\mathbf{N}_F]_{kl} = \epsilon^{-2} \Sigma \sqrt{a_{kl} \text{Var}\{[\mathbf{F}_Q]_{kl}\}}, \quad (\text{C1})$$

$$[\mathbf{N}_g]_k = \epsilon^{-2} \Sigma \sqrt{b_k \text{Var}[g_k]}, \quad (\text{C2})$$

respectively. Here $\text{Var}[\cdot]$ is the variance of a single measurement of the corresponding element and we explicitly define Σ via the coefficients a_{kl} and b_k from Appendix B as

$$\Sigma := \sum_{k,l=1}^v \sqrt{a_{kl} \text{Var}\{[\mathbf{F}_Q]_{kl}\}} + \sum_{k=1}^v \sqrt{b_k \text{Var}[g_k]}. \quad (\text{C3})$$

Proof. From Lemma 2 we write the error measure as

$$\epsilon^2 = \sum_{k,l=1}^v a_{kl} \text{Var}\{[\mathbf{F}_Q]_{kl}\} + \sum_{k=1}^v b_k \text{Var}[g_k], \quad (\text{C4})$$

where $\text{Var}[\cdot]$ denotes the variance in the statistical average over many measurements. Now we allow $\text{Var}[\cdot]$ to denote variance in a single measurement, while $[\mathbf{N}_F]_{kl}$ and $[\mathbf{N}_g]_k$ are the number of measurements assigned each element $[\mathbf{F}_Q]_{kl}$ and g_k , respectively, so that the error measure becomes

$$\epsilon^2 = \sum_{k,l=1}^v \frac{a_{kl} \text{Var}\{[\mathbf{F}_Q]_{kl}\}}{[\mathbf{N}_F]_{kl}} + \sum_{k=1}^v \frac{b_k \text{Var}[g_k]}{[\mathbf{N}_g]_k}. \quad (\text{C5})$$

By minimizing the error measure, in this form, subject to the constraint of a fixed total number of measurements, so that

$$N_{\text{opt}} = \sum_{k,l=1}^v [\mathbf{N}_F]_{kl} + \sum_{k=1}^v [\mathbf{N}_g]_k, \quad (\text{C6})$$

we find that the optimal fraction of measurement to be assigned to each element is

$$\frac{[\mathbf{N}_F]_{kl}}{N_{\text{opt}}} = \frac{\sqrt{a_{kl} \text{Var}\{[\mathbf{F}_Q]_{kl}\}}}{\Sigma}, \quad \frac{[\mathbf{N}_g]_k}{N_{\text{opt}}} = \frac{\sqrt{b_k \text{Var}[g_k]}}{\Sigma}, \quad (\text{C7})$$

$$\text{where } \Sigma := \sum_{k,l=1}^v \sqrt{a_{kl} \text{Var}\{[\mathbf{F}_Q]_{kl}\}} + \sum_{k=1}^v \sqrt{b_k \text{Var}[g_k]}. \quad (\text{C8})$$

By substituting this results into the error measure we can remove the dependence on the total number of measurements N_{opt} to yield the required result. ■

1. Fisher matrix symmetry

Lemma 5. *The symmetry of the Fisher matrix can be accounted for by replacing the elements a_{kl} with a'_{kl} , where*

$$a'_{kl} := \begin{cases} 0, & k < l, \\ a_{kk}, & k = l, \\ a_{kl} + a_{lk}, & k > l. \end{cases} \quad (\text{C9})$$

Proof. As the Fisher matrix is symmetric, measurements of the $[\mathbf{F}_Q]_{kl}$ element also constitute measurements of the $[\mathbf{F}_Q]_{lk}$, so $\text{Var}\{[\mathbf{F}_Q]_{kl}\} = \text{Var}\{[\mathbf{F}_Q]_{lk}\}$. Thus, the error measure can be written as

$$\begin{aligned} \epsilon^2 &= \sum_{k=1}^v a_{kk} \text{Var}\{[\mathbf{F}_Q]_{kk}\} + \sum_{k>l}^v 2a_{kl} \text{Var}\{[\mathbf{F}_Q]_{kl}\} \\ &+ \sum_{k=1}^v b_k \text{Var}[g_k]. \end{aligned} \quad (\text{C10})$$

It is possible to force this back into the original form of the error measure if we define

$$a'_{kl} := \begin{cases} 0, & k < l, \\ a_{kk}, & k = l, \\ a_{kl} + a_{lk}, & k > l, \end{cases} \quad (\text{C11})$$

so that the error measure error measure can be written as

$$\epsilon^2 = \sum_{k,l=1}^v a'_{kl} \text{Var}\{[\mathbf{F}_Q]_{kl}\} + \sum_{k=1}^v b_k \text{Var}[g_k]. \quad (\text{C12})$$

Using the error measure written in this form as a starting point for the derivation in the proof of Lemma 4, we trivially obtain the same results with the elements a_{kl} replaced with a'_{kl} . ■

APPENDIX D: APPLICATIONS BEYOND NATURAL GRADIENT

Let us now comment on how the main results of this work can be applied to other quantum algorithms beyond natural gradient optimization. For this reason, we now consider three categories of algorithms and review how the results in the main text can be tailored to these algorithms. We note that all algorithms considered in the following use a parameter update rule whereby an inverse matrix A^{-1} is applied to a vector v .

Metric-aware optimization algorithms. We have covered metric-aware optimization algorithms in the main text that include quantum natural gradient descent and imaginary time evolution [37,39,42–44]. In this case the matrix object A is the quantum Fisher information, which depends

only on the ansatz circuit, while the vector object is the gradient vector that depends on both the ansatz circuit and on the Hamiltonian.

Variational quantum simulation. The time evolution of a quantum system under a Hamiltonian can be simulated using techniques described in Refs. [42,43]. In such a scenario the matrix object (imaginary part of the quantum geometric tensor) still depends only on the ansatz circuit and the vector object is related to the gradient vector and therefore all our results apply, except for Result 1. The reason is the following. In Result 1 we assumed that, when increasing the number of iterations, the norm of the gradient vector v vanishes due to convergence. However, in the case of variational simulation, the norm of the vector object does not necessarily decrease.

Hessian optimization. Analogously to metric-aware optimizations here the inverse of the Hessian matrix is applied to the gradient vector. Thus, our error propagation formula in Eq. (2) and our optimal measurement distribution scheme in Result 3 immediately apply to this scenario too. The main difference to the previously discussed scenarios is that the Hessian matrix now depends on both the ansatz circuit and on the Hamiltonian. For this reason, Theorems 1 and 2 need to be modified such that the dependence on the Hamiltonian is taken into account. As such, the main conclusion of Result 1 will still hold: as the optimization converges, the vanishingly small gradient becomes increasingly more expensive to determine to a sufficient precision. However, via Result 2, it is expected that, when we increase the number of qubits, the Hessian matrix becomes increasingly more expensive to estimate due to its dependence on the Hamiltonian.

-
- [1] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](#) (2014).
 - [2] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, 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).
 - [3] Y. Wang, F. Dolde, J. Biamonte, R. Babbush, V. Bergholm, S. Yang, I. Jakobi, P. Neumann, A. Aspuru-Guzik, and J. D. Whitfield *et al.*, Quantum simulation of helium hydride cation in a solid-state spin register, *ACS Nano* **9**, 7769 (2015).
 - [4] P. J. J. O’Malley *et al.*, Scalable Quantum Simulation of Molecular Energies, *Phys. Rev. X* **6**, 031007 (2016).
 - [5] Y. Shen, X. Zhang, S. Zhang, J.-N. Zhang, M.-H. Yung, and K. Kim, Quantum implementation of the unitary coupled cluster for simulating molecular electronic structure, *Phys. Rev. A* **95**, 020501 (2017).
 - [6] 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).
 - [7] S. Paesani, A. A. Gentile, R. Santagati, J. Wang, N. Wiebe, D. P. Tew, J. L. O’Brien, and M. G. Thompson, Experimental Bayesian Quantum Phase Estimation on a Silicon Photonic Chip, *Phys. Rev. Lett.* **118**, 100503 (2017).
 - [8] Y. Li and S. C. Benjamin, Efficient Variational Quantum Simulator Incorporating Active Error Minimization, *Phys. Rev. X* **7**, 021050 (2017).
 - [9] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, Computation of Molecular Spectra on a Quantum Processor with an Error-Resilient Algorithm, *Phys. Rev. X* **8**, 011021 (2018).
 - [10] R. Santagati, J. Wang, A. A. Gentile, S. Paesani, N. Wiebe, J. R. McClean, S. Morley-Short, P. J. Shadbolt, D. Bonneau, J. W. Silverstone, D. P. Tew, X. Zhou, J. L. O’Brien, and M. G. Thompson, Witnessing eigenstates for quantum simulation of hamiltonian spectra, *Sci. Adv.* **4**, eaap9646 (2018).
 - [11] 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* **549**, 242 (2017).
 - [12] A. Kandala, K. Temme, A. D. Corcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, Error mitigation extends the computational reach of a noisy quantum processor, *Nature* **567**, 491 (2019).
 - [13] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, and C. F. Roos, Quantum Chemistry Calculations on a Trapped-Ion Quantum Simulator, *Phys. Rev. X* **8**, 031022 (2018).
 - [14] J. Romero, R. Babbush, J. R. McClean, C. Hempel, P. Love, and A. Aspuru-Guzik, Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz, [arXiv:1701.02691](#) (2017).
 - [15] O. Higgott, D. Wang, and S. Brierley, Variational Quantum Computation of Excited States, [arXiv:1805.08138](#) (2018).
 - [16] T. Jones, S. Endo, S. McArdle, X. Yuan, and S. C. Benjamin, Variational quantum algorithms for discovering hamiltonian spectra, *Phys. Rev. A* **99**, 062304 (2019).
 - [17] J. R. McClean, M. E. Kimchi-Schwartz, J. Carter, and W. A. de Jong, Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states, *Phys. Rev. A* **95**, 042308 (2017).
 - [18] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, Robust determination of molecular spectra on a quantum processor, [arXiv:1707.06408](#) (2017).
 - [19] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, and C. F. Roos *et al.*, Self-Verifying Variational Quantum Simulation of the Lattice Schwinger Model, [arXiv:1810.03421](#) (2018).
 - [20] K. Sharma, S. Khatri, M. Cerezo, and P. J. Coles, Noise resilience of variational quantum compiling, *New J. Phys.* **22**, 043006 (2020).
 - [21] M. Cerezo, K. Sharma, A. Arrasmith, and P. J. Coles, Variational Quantum State Eigensolver, [arXiv:2004.01372](#) (2020).
 - [22] S. Endo, J. Sun, Y. Li, S. C. Benjamin, and X. Yuan, Variational Quantum Simulation of General Processes, *Phys. Rev. Lett.* **125**, 010501 (2020).
 - [23] B. Koczor and S. C. Benjamin, Quantum analytic descent, [arXiv:2008.13774](#) (2020).

- [24] B. Koczor, Exponential error suppression for near-term quantum devices, *Phys. Rev. X* (to be published), [arXiv:2011.05942](#) (2020).
- [25] J. Preskill, Quantum Computing in the NISQ era and beyond, [arXiv:1801.00862](#) (2018).
- [26] I. Kassal, J. D. Whitfield, A. Perdomo-Ortiz, M.-H. Yung, and A. Aspuru-Guzik, Simulating chemistry using quantum computers, *Annu. Rev. Phys. Chem.* **62**, 185 (2011).
- [27] D. Lu, B. Xu, N. Xu, Z. Li, H. Chen, X. Peng, R. Xu, and J. Du, Quantum chemistry simulation on quantum computers: Theories and experiments, *Phys. Chem. Chem. Phys.* **14**, 9411 (2012).
- [28] K. B. Whaley, A. R. Dinner, and S. A. Rice, *Quantum Information and Computation for Chemistry* (John Wiley & Sons, Hoboken, New Jersey, 2014).
- [29] S. McArdle, S. Endo, A. Aspuru-Guzik, S. Benjamin, and X. Yuan, Quantum computational chemistry, *Rev. Mod. Phys.* **92**, 15003 (2020).
- [30] G. A. Quantum *et al.*, Hartree-fock on a superconducting qubit quantum computer, *Science* **369**, 1084 (2020).
- [31] R. Sweke, F. Wilde, J. Meyer, M. Schuld, P. K. Fährmann, B. Meynard-Piganeau, and J. Eisert, Stochastic gradient descent for hybrid quantum-classical optimization, [arXiv:1910.01155](#) (2019).
- [32] J. M. Kübler, A. Arrasmith, L. Cincio, and P. J. Coles, An adaptive optimizer for measurement-frugal variational algorithms, [arXiv:1909.09083](#) (2019).
- [33] L. Gentini, A. Cuccoli, S. Pirandola, P. Verrucchi, and L. Bianchi, Noise-Assisted Variational Hybrid Quantum-Classical Optimization, [arXiv:1912.06744](#) (2019).
- [34] O. Crawford, B. van Straaten, D. Wang, T. Parks, E. Campbell, and S. Brierley, Efficient quantum measurement of Pauli operators, [arXiv:1908.06942](#) (2019).
- [35] A. Arrasmith, L. Cincio, R. D. Somma, and P. J. Coles, Operator Sampling for Shot-frugal Optimization in Variational Algorithms, [arXiv:2004.06252](#) (2020).
- [36] C. Hadfield, S. Bravyi, R. Raymond, and A. Mezzacapo, Measurements of Quantum Hamiltonians with Locally-Biased Classical Shadows, [arXiv:2006.15788](#) (2020).
- [37] B. Koczor and S. C. Benjamin, Quantum natural gradient generalised to non-unitary circuits, [arXiv:1912.08660](#) (2019).
- [38] R. R. Ferguson, L. Dellantonio, A. A. Balushi, K. Jansen, W. Dür, and C. A. Muschik, Measurement-Based Variational Quantum Eigensolver, *Phys. Rev. Lett.* **126**, 220501 (2021).
- [39] S. McArdle, T. Jones, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan, Variational ansatz-based quantum simulation of imaginary time evolution, *Npj Quantum Inf.* **5**, 1 (2019).
- [40] P. Rebentrost, M. Schuld, L. Wossnig, F. Petruccione, and S. Lloyd, Quantum gradient descent and newton's method for constrained polynomial optimization, *New J. Phys.* **21**, 073023 (2019).
- [41] D. Wierichs, C. Gogolin, and M. Kastoryano, Avoiding local minima in variational quantum eigensolvers with the natural gradient optimizer, *Phys. Rev. Research* **2**, 043246 (2020).
- [42] Y. Li and S. C. Benjamin, Efficient Variational Quantum Simulator Incorporating Active Error Minimization, *Phys. Rev. X* **7**, 021050 (2017).
- [43] X. Yuan, S. Endo, Q. Zhao, Y. Li, and S. C. Benjamin, Theory of variational quantum simulation, *Quantum* **3**, 191 (2019).
- [44] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, Quantum natural gradient, [arXiv:1909.02108](#) (2019).
- [45] Where the ansatz parameter θ corresponds to a global Z rotation of all the qubits and therefore $r_g = N$.
- [46] L. Pezzè, A. Smerzi, M. K. Oberthaler, R. Schmied, and P. Treutlein, Quantum metrology with nonclassical states of atomic ensembles, *Rev. Mod. Phys.* **90**, 035005 (2018).
- [47] V. Giovannetti, S. Lloyd, and L. Maccone, Advances in quantum metrology, *Nat. Phot.* **5**, 222 (2011).
- [48] B. Koczor, S. Endo, T. Jones, Y. Matsuzaki, and S. C. Benjamin, Variational-state quantum metrology, *New J. Phys.* **22**, 083038 (2020).
- [49] In Fig. 1 in the uniformly distributed scenario the same number of measurements are used to determine every entry of the metric tensor (gradient vector) and only the overall number N_F (N_g) of measurements to determine the metric tensor (gradient vector) is chosen optimally. In the naive scheme both N_F and N_g are additionally fixed.
- [50] T.-C. Yen, V. Verteletskyi, and A. F. Izmaylov, Measuring all compatible operators in one series of single-qubit measurements using unitary transformations, *J. Chem. Theory Comput.* **16**, 2400 (2020).
- [51] A. Jena, S. Genin, and M. Mosca, Pauli partitioning with respect to gate sets, [arXiv:1907.07859](#) (2019).
- [52] P. Gokhale, O. Angiuli, Y. Ding, K. Gui, T. Tomesh, M. Suchara, M. Martonosi, and F. T. Chong, $O(N^3)$ measurement cost for variational quantum eigensolver on molecular hamiltonians, *IEEE Trans. Quantum Eng.* **1**, 1 (2020).
- [53] 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](#) (2019).
- [54] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 1 (2018).
- [55] E. Grant, L. Wossnig, M. Ostaszewski, and M. Benedetti, An initialization strategy for addressing barren plateaus in parametrized quantum circuits, *Quantum* **3**, 214 (2019).
- [56] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost-Function-Dependent Barren Plateaus in Shallow Quantum Neural Networks, [arXiv:2001.00550](#) (2020).
- [57] T. Jones, A. Brown, I. Bush, and S. C. Benjamin, QuEST and high performance simulation of quantum computers, *Sci. Rep.* **9**, 10736 (2019).
- [58] T. Jones and S. C. Benjamin, QuESTlink—Mathematica embiggened by a hardware-optimised quantum emulator, *Quantum Sci. Technol.* **5**, 034012 (2020).
- [59] H. Ku, Notes on the use of propagation of error formulas, *J. Res. Natl. Bur. Stand. Sec. C: Eng. Instrum.* **70C**, 263 (1966).
- [60] M. Lefebvre, R. K. Keeler, R. Sobie, and J. White, Propagation of errors for matrix inversion, *Nucl. Instrum. Methods Phys. Res. Sec. A: Accelerators, Spectrometers, Detectors Associated Equipment* **451**, 520 (2000), [arXiv:9909031 \[hep-ex\]](#).