

ARTICLE OPEN



Variational quantum state eigensolver

M. Cerezo^{1,2,✉}, Kunal Sharma^{1,3}, Andrew Arrasmith¹ and Patrick J. Coles¹

Extracting eigenvalues and eigenvectors of exponentially large matrices will be an important application of near-term quantum computers. The variational quantum eigensolver (VQE) treats the case when the matrix is a Hamiltonian. Here, we address the case when the matrix is a density matrix ρ . We introduce the variational quantum state eigensolver (VQSE), which is analogous to VQE in that it variationally learns the largest eigenvalues of ρ as well as a gate sequence V that prepares the corresponding eigenvectors. VQSE exploits the connection between diagonalization and majorization to define a cost function $C = \text{Tr}(\tilde{\rho}H)$ where H is a non-degenerate Hamiltonian. Due to Schur-concavity, C is minimized when $\tilde{\rho} = V\rho V^\dagger$ is diagonal in the eigenbasis of H . VQSE only requires a single copy of ρ (only n qubits) per iteration of the VQSE algorithm, making it amenable for near-term implementation. We heuristically demonstrate two applications of VQSE: (1) Principal component analysis, and (2) Error mitigation.

npj Quantum Information (2022)8:113; <https://doi.org/10.1038/s41534-022-00611-6>

INTRODUCTION

Near-term quantum computers hold great promise but also pose great challenges. Low qubit counts place constraints on problem sizes that can be implemented. Decoherence and gate infidelity place constraints on the circuit depth that can be implemented. These constraints are captured in the (now widely used) term noisy intermediate-scale quantum (NISQ)¹.

To address the circuit depth constraint, variational quantum algorithms (VQAs) have been proposed for many applications^{2–27}. VQAs employ a quantum-classical optimization loop to train the parameters θ of a quantum circuit $V(\theta)$. Leveraging classical optimizers allows the quantum circuit depth to remain shallow. This makes VQAs powerful tools for error mitigation on NISQ devices.

A particularly important application of NISQ computers will be extracting the spectra, eigenvalues, and eigenvectors, of very large matrices. Indeed the most famous VQA, known as the variational quantum eigensolver (VQE), aims to variationally determine the energies and state-preparation circuits for the ground state and low-lying excited states of a given Hamiltonian, i.e., a Hermitian matrix. VQE promises to revolutionize the field of quantum chemistry^{28,29}, and perhaps even nuclear³⁰ and condensed matter^{31,32} physics.

If one instead considers a positive-semidefinite matrix, then extracting the spectrum has direct application as a machine-learning primitive known as principal component analysis (PCA). Along these lines, Lloyd et al.³³ introduced a quantum algorithm called quantum PCA (qPCA) to deterministically extract the spectrum of an n -qubit density matrix ρ . qPCA employs quantum phase estimation and density matrix exponentiation as subroutines and hence requires a large number of quantum gates and copies of ρ . In an effort to reduce circuit depth in the NISQ era, LaRose et al.⁶ developed a VQA for this application called variational quantum state diagonalization (VQSD). VQSD requires two copies of ρ , hence $2n$ qubits, and trains the parameters θ of a gate sequence $V(\theta)$ so that $\tilde{\rho} = V(\theta)\rho V^\dagger(\theta)$ is approximately diagonal. A different variational approach, called quantum singular value decomposition (QSVD), was introduced by Bravo-Prieto et al.²⁷. QSVD takes a purification $|\psi\rangle$ of ρ as its input and

hence requires however many qubits it takes to purify ρ (possibly $2n$ qubits).

In this work, we introduce a variational algorithm for PCA that only requires a single copy of ρ and hence only n qubits per iteration of the algorithm. Our approach, called the variational quantum state eigensolver (VQSE), exploits the mathematical connection between diagonalization and majorization. Namely, it is well known that the eigenvalues of a density matrix ρ majorize the diagonal elements in any basis. Hence, by choosing a cost function C that is a Schur concave function of the diagonal elements of ρ , one can ensure that the cost function is minimized when ρ is diagonalized. Specifically, we write the cost as $C = \text{Tr}(\tilde{\rho}H)$, where H is some Hamiltonian with a non-degenerate spectrum, which ensures the Schur concavity property. Note that evaluating C simply involves measuring the expectation value of H on $\tilde{\rho}$, and hence one can see why only n qubits are required.

To learn the optimal θ parameters, we introduce a new training approach, not previously used in other VQAs. Specifically, we employ a time-dependent Hamiltonian H that we adapt based on information gained from measurements performed throughout the optimization. The aim of this adaptive approach is: (1) to mitigate barren plateaus in training landscapes, and (2) to get out of local minima. With our numerics, we find that using an adaptive Hamiltonian is better than simply fixing the Hamiltonian throughout the optimization. Here, we further provide a rigorous analysis of the measurement shot requirements of VQSE where we show that the relative error induced from statistical sampling error is, with high probability, smaller than δ , if one measures the system response with a number of shots that scales only as $\Omega(\log(1/\delta)/\lambda_m^2)$, with λ_m being the smallest eigenvalue one wishes to estimate.

Finally, we illustrate two important applications of VQSE with our numerical implementations. First, we use VQSE for error mitigation of the W -state preparation circuit. Namely, by projecting the state onto the eigenvector with the largest eigenvalue, we re-purify the state, mitigating the effects of incoherent errors. Second, we use VQSE to perform entanglement spectroscopy (which is essentially PCA on the reduced state of a bipartition) on

¹Theoretical Division, MS B213, Los Alamos National Laboratory, Los Alamos, NM 87545, USA. ²Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, NM, USA. ³Hearne Institute for Theoretical Physics and Department of Physics and Astronomy, Louisiana State University, Baton Rouge, LA, USA. ✉email: cerezo@lanl.gov

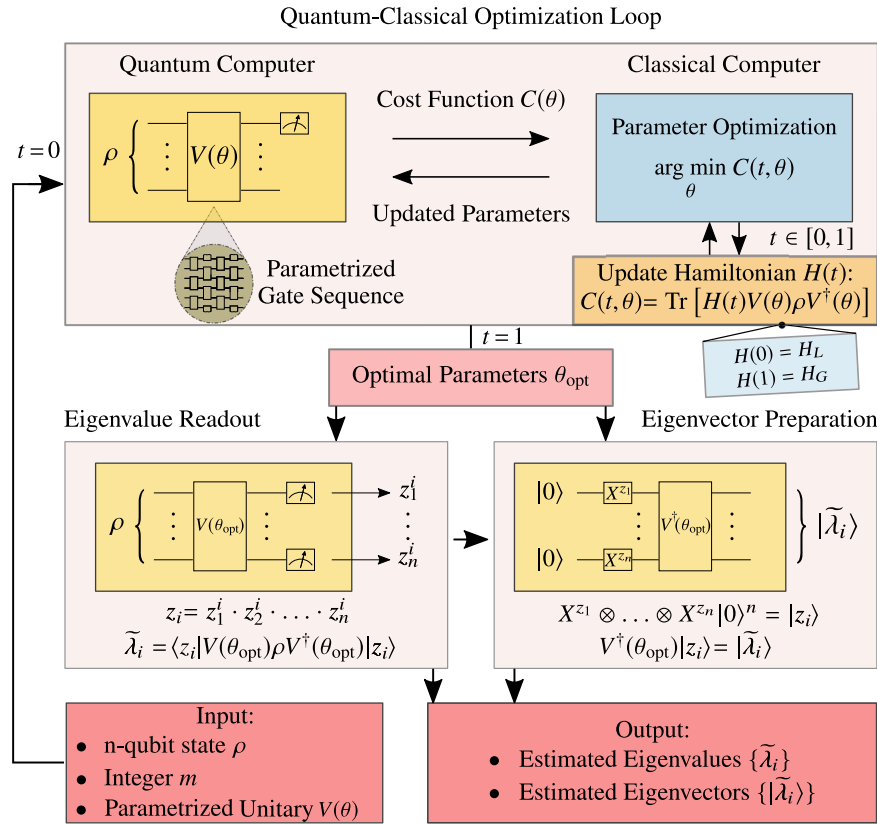


Fig. 1 Schematic diagram of VQSE. VQSE takes as inputs an n -qubit state ρ , an integer m , and a parametrized unitary $V(\theta)$. It then outputs estimates of the m -largest eigenvalues of ρ , and their associated eigenvectors. The first step of the algorithm is a hybrid quantum-classical optimization loop to train the parameters θ , and minimize the cost function defined in (4) as the expectation value of a Hamiltonian $H(t)$ over the state $\tilde{\rho} = V(\theta)\rho V^\dagger(\theta)$. To facilitate this optimization, we adaptively update $H(t)$ using information obtained via measurements on $\tilde{\rho}$. When this optimization terminates, at which point we say $\theta = \theta_{\text{opt}}$, one reads off the eigenvalues. Namely, by preparing $V(\theta_{\text{opt}})\rho V^\dagger(\theta_{\text{opt}})$ and measuring in the standard basis, one obtains bitstrings \mathbf{z} whose associated frequencies are estimates of the eigenvalues of ρ . Finally, one prepares the estimated eigenvectors by preparing the states $|\mathbf{z}\rangle$ and acting on them with $V^\dagger(\theta_{\text{opt}})$.

the ground state of an XY-model spin chain. This allows us to identify quantum critical points in this model.

RESULTS

Theoretical basis of VQSE

Consider an n -qubit quantum state ρ with (unknown) spectral decomposition $\rho = \sum_k \lambda_k |\lambda_k\rangle \langle \lambda_k|$, such that the eigenvalues are ordered in decreasing order (i.e., $\lambda_k \geq \lambda_{k+1}$ for $k = 1, \dots, \text{rank}(\rho)$), while $\lambda_k = 0$ for $k \geq \text{rank}(\rho)$). The goal of VQSE is to estimate the m -largest eigenvalues of ρ , where $m \ll 2^n$, and furthermore to return a gate sequence $V(\theta)$ that approximately prepares their associated eigenvectors from standard basis elements.

At first sight, this looks like a matrix diagonalization problem. Indeed, this is the perspective taken in the literature, e.g., by the VQSD algorithm⁶ which employs a cost function that quantifies how far $\tilde{\rho} = V(\theta)\rho V^\dagger(\theta)$ is from a diagonal matrix. However, our VQSE algorithm takes a conceptually different approach, focusing on majorization instead of diagonalization.

We write the VQSE cost function as an energy, or the expectation value of a Hamiltonian:

$$C(\theta) \equiv \langle H \rangle = \text{Tr}[HV(\theta)\rho V^\dagger(\theta)]. \quad (1)$$

Here, H is a simple n -qubit Hamiltonian that is diagonal in the standard basis and whose eigenenergies and associated eigenstates are known and respectively given by $\{E_k\}$ and $\{|\mathbf{e}_k\rangle\}$ (where $\mathbf{e}_k = e_k^1 \dots e_k^n$ for $k = 1, \dots, 2^n$ are bitstrings of length n). Moreover, we henceforth assume that the eigenenergies are

non-negative and ordered in increasing order, i.e., $E_k \leq E_{k+1}$. We have

$$C(\theta) = \sum_{k=1}^{2^n} E_k p_k = \mathbf{E} \cdot \mathbf{p}, \quad p_k = \langle \mathbf{e}_k | \tilde{\rho} | \mathbf{e}_k \rangle, \quad (2)$$

where we defined the vectors $\mathbf{E} = (E_1, E_2, \dots)$ and $\mathbf{p} = (p_1, p_2, \dots)$. Similarly, let us define the vector of eigenvalues of ρ as $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \dots)$. Then, since the eigenvalues of a positive semidefinite matrix majorize its diagonal elements $\boldsymbol{\lambda} \succ \mathbf{p}$, and since the dot product with an increasingly ordered vector is a Schur concave function^{34,35}, we have

$$C(\theta) = \mathbf{E} \cdot \mathbf{p} \geq \mathbf{E} \cdot \boldsymbol{\lambda} = \sum_k E_k \lambda_k, \quad (3)$$

where we have used the fact that ρ and $\tilde{\rho}$ have the same eigenvalues. Hence, one can see that $C(\theta)$ is minimized when $V(\theta)$ maps the eigenbasis of ρ to the eigenbasis of H , with appropriate ordering. Since the latter is chosen to be the standard basis, this corresponds to diagonalizing ρ . Thus, even though it may not be obvious at first sight, minimizing $C(\theta)$ corresponds to diagonalizing ρ .

The VQSE algorithm

Figure 1 shows a schematic diagram of the VQSE algorithm. The three inputs to VQSE are: (1) a n -qubit quantum state ρ , (2) an integer m , and (3) a parameterized gate sequence or ansatz $V(\theta)$. The outputs of VQSE are: (1) estimates $\{\tilde{\lambda}_i\}_{i=1}^m$ of the m -largest eigenvalues $\{\lambda_i\}_{i=1}^m$ of ρ , and (2) a gate sequence $V(\theta_{\text{opt}})$ that

prepares approximate versions $\{\tilde{\lambda}_i\}_{i=1}^m$ of the associated m eigenvectors $\{|\lambda_i\rangle\}_{i=1}^m$. While in principle m can be as large as 2^n , we assume that one is interested in a number m of eigenvalues and eigenvectors that grows at worse as $\mathcal{O}(\text{poly}(n))$.

After taking in the inputs, VQSE enters a hybrid quantum-classical optimization loop to train the parameters θ in the ansatz $V(\theta)$. This loop employs a quantum computer to evaluate the VQSE cost function, denoted

$$C(t, \theta) \equiv \langle H(t) \rangle = \text{Tr}[H(t)\tilde{\rho}], \quad \tilde{\rho} = V(\theta)\rho V^\dagger(\theta). \quad (4)$$

Here, $H(t)$ is a Hamiltonian that could, in general, depend on the time t , where $t \in [0, 1]$ is a parameter that indicates the optimization loop run-time such that the loop starts at $t=0$ and ends at $t=1$. For all t , we assume that $H(t)$ can be efficiently measured on a quantum computer and that it is diagonal in the standard basis, with its lowest m eigenenergies being non-degenerate and non-negative. We further elaborate on how to choose $H(t)$ in section 'Cost functions'. Note that the quantum circuit to evaluate the cost $C(t, \theta)$, as depicted in Fig. 1, simply involves applying $V(\theta)$ to the state ρ and then measuring the Hamiltonian $H(t)$.

The quantum computer then feeds the value of the cost (or the gradient of the cost for gradient-based optimization) to a classical computer, which adjusts the parameters θ for the next round of the loop. The ultimate goal is to find the global minimum of the cost landscape at $t=1$, i.e., to solve the problem:

$$\theta_{\text{opt}} \equiv \arg \min_{\theta} C(1, \theta). \quad (5)$$

In reality, one will need to impose some termination condition on the optimization loop and hence the final parameters obtained (which we still denote as θ_{opt}) will only approximately satisfy Eq. (5). Nevertheless, we provide a verification procedure below in section 'Verification of solution quality' that allows one to quantify the quality of the solution even when (5) is not exactly satisfied.

As shown in Fig. 1, the next step of VQSE is the eigenvalue readout. From the parameters θ_{opt} one can estimate the eigenvalues of ρ by acting with the gate sequence $V(\theta_{\text{opt}})$ and then measuring in the standard basis $\{|z_k\rangle\}$. Let $\text{Pr}(z_k)$ be the probability of the z_k outcome. Then by taking the m largest of these probabilities we define $\mathcal{L} \equiv \{\tilde{\lambda}_i\}_{i=1}^m$ as the ordered set of estimates of the m -largest eigenvalues of ρ , and we define \mathcal{Z} as the set of bitstrings $\{z_i\}_{i=1}^m$ associated with the elements of \mathcal{L} :

$$\tilde{\lambda}_i = \text{Pr}(z_i) = \langle z_i | \tilde{\rho} | z_i \rangle, \quad \text{such that } \tilde{\lambda}_i \geq \tilde{\lambda}_{i+1}. \quad (6)$$

Note that $\tilde{\lambda}_i$ in (6) correspond to diagonal elements of $\tilde{\rho}$ in the standard basis, and not to its eigenvalues.

In practice, when estimating the eigenvalues one measures $\tilde{\rho}$ in the standard basis a finite number of times N_{runs} . Hence, if a bitstring $z_i \in \mathcal{Z}$ has frequency f_i for N_{runs} total runs, then we can estimate $\tilde{\lambda}_i$ as

$$\tilde{\lambda}_i^{\text{est}} = \frac{f_i}{N_{\text{runs}}}. \quad (7)$$

One can think of this as a Bernoulli trial. Let Λ_i be a random variable that takes value 1 if we get outcome z_i (with probability $\tilde{\lambda}_i$), and takes value 0 otherwise (with probability $1 - \tilde{\lambda}_i$). After repeating the experiment N_{runs} times we are interested in bounding the probability that the relative error $\varepsilon_i \equiv \tilde{\lambda}_i^{\text{est}} - \tilde{\lambda}_i / \tilde{\lambda}_i$ is larger than a certain value $c \geq 0$. From Hoeffding's inequality, we find

$$\text{Pr}(\varepsilon_i \geq c) \leq e^{-2N_{\text{runs}}c^2\tilde{\lambda}_i}, \quad \forall c > 0. \quad (8)$$

For fixed N_{runs} , Eq. (8) shows that the smaller the inferred eigenvalue $\tilde{\lambda}_i$, the larger the probability of having a given relative error. Equation (8) also implies that increasing N_{runs} reduces the probability of large relative errors. Hence, we can always choose

N_{runs} such that the probability of error is smaller than a given δ for all m eigenvalues via

$$\forall i \in [1, m], \quad \text{Pr}(\varepsilon_i \geq c) \leq \delta \rightarrow N_{\text{runs}} \geq \frac{\log(1/\delta)}{2c^2\lambda_m^2}, \quad (9)$$

where λ_m is the smallest eigenvalue of interest. Analogously, from (9) we have that all eigenvalues larger than $\sqrt{\frac{\log(1/\delta)}{2c^2N_{\text{runs}}}}$ have a probability of error smaller than δ .

The last step of VQSE is to prepare the inferred eigenvectors of ρ . Given a bitstring $z_i \in \mathcal{Z}$, one can prepare the associated inferred eigenvector by taking the state $|\mathbf{0}\rangle = |0\rangle^{\otimes n}$, acting on it with the gate $X^{z_1} \otimes X^{z_2} \otimes \dots \otimes X^{z_n}$, and then applying the gate sequence $V(\theta_{\text{opt}})^\dagger$:

$$|\tilde{\lambda}_i\rangle = V^\dagger(\theta_{\text{opt}})|z_i\rangle, \quad |z_i\rangle = X^{z_1} \otimes \dots \otimes X^{z_n}|\mathbf{0}\rangle. \quad (10)$$

Note that while the inferred eigenvalues can be stored classically, the eigenvectors are prepared on a quantum computer, and hence one needs to perform measurements to extract information about these eigenvectors.

Cost functions

Consider the Hamiltonian $H(t)$ that defines the VQSE cost function in (4). Recall that we choose $H(t)$ so that: (1) it is diagonal in the standard basis, (2) its lowest m eigenvalues are non-negative and non-degenerate, and (3) it can be efficiently measured on a quantum computer. Let us now discuss possible choices for $H(t)$.

Fixed Hamiltonians. When the Hamiltonian is fixed (i.e., time-independent), we write $H(t) \equiv H$, and $C(t, \theta) \equiv C(\theta)$. In this case, a simple, intuitive cost function is given by

$$C_G(\theta) = \text{Tr}[H_G\tilde{\rho}], \quad H_G = \mathbb{1} - \sum_{i=1}^m q_i |\mathbf{e}_i\rangle \langle \mathbf{e}_i|, \quad (11)$$

with $q_i > 0$ (such that $q_i > q_{i+1}$), and where the $|\mathbf{e}_i\rangle$ are orthogonal states in the standard basis. The spectrum of H_G is composed of m non-degenerate eigenenergies, and a $(2^n - m)$ -fold degenerate eigenenergy.

On the one hand, this large degeneracy makes it easier to find a global minimum as the solution space is large. That is, denoting as V_{opt} an optimal unitary that minimizes (11), then there is a large set of such optimal unitaries $\mathcal{S}_{\text{opt}} = \{V_{\text{opt}}\}$, which are not related by global phases. This is due to the fact that one is only interested in the m rows and the m columns of $V(\theta)$ that diagonalize $\tilde{\rho}$ in the subspace spanned by $\{|\mathbf{e}_i\rangle\}_{i=1}^m$. Specifically, any optimal unitary must satisfy $\langle z_i | V_{\text{opt}} |\lambda_i\rangle = \langle \lambda_i | V_{\text{opt}} |z_i\rangle = \delta_{z_i, \mathbf{e}_i}$ for $i = 1, \dots, m$ (and with $z_i \in \mathcal{Z}$), while the $(2^n - m) \times (2^n - m)$ unitary principal submatrix of V_{opt} with matrix elements $\langle z_i | V_{\text{opt}} |z_i'\rangle$, where $z_i, z_i' \notin \mathcal{Z}$, remains completely arbitrary.

On the other hand, it has been shown that when employing hardware-efficient ansatzes³⁶ for $V(\theta)$, global cost functions like $C_G(\theta)$ are untrainable for large problem sizes as they exhibit exponentially vanishing gradients (i.e., barren plateaus³⁷) even when the ansatz is short depth³⁸. Such barren plateaus can be avoided by employing a different type of cost function known as a local cost^{38,39}, where C is defined such that one compares states or operators with respect to each individual qubit rather than comparing them in a global sense.

One can construct a local cost where the Hamiltonian is a weighted sum of local z -Pauli operators:

$$C_L \equiv \langle H_L \rangle, \quad H_L = \mathbb{1} - \sum_{j=1}^n r_j Z_j, \quad (12)$$

where $r_j \in \mathbb{R}$ and Z_j is the z -Pauli operator acting on qubit j . Care must be taken when choosing the coefficients $\{r_j\}_{j=1}^n$ to ensure that the lowest m -eigenenergies of H_L are non-degenerate.

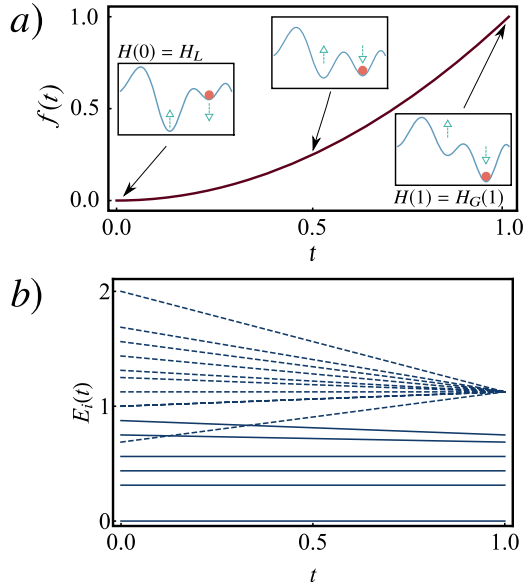


Fig. 2 Adaptive cost function. **a** Schematic representation of the function $f(t)$ and the cost landscape of $C(t, \theta)$ versus t . We choose $f(t)$ as a slowly growing function with t . When the optimization starts at $t=0$, the cost function does not exhibit a barren plateau as the Hamiltonian is local $H(0) = H_L$. As t increases $H(t)$ becomes a linear combination of H_L and a global Hamiltonian $H_G(t)$ which is adaptively updated using information gained from measurements on $V(\theta)\rho V^\dagger(\theta)$. As shown in the insets, this procedure allows for local minima to become global minima. Finally, when the algorithm ends at $t=1$ the Hamiltonian is global $H(1) = H_G(1)$, and as shown in panel (b) of Fig. 2, the spectrum of H becomes highly degenerate and the dimension of the solution space increases. In section ‘Algorithm for the adaptive cost function’ of the ‘Methods’, we present an algorithm to illustrate how one can update $H(t)$. **b** Schematic representation of the eigenenergies of $H(t)$ versus t . For small t the Hamiltonian is local and hence its spectrum contains non-degeneracies that reduce the space of solutions. At $t=1$, $H(t)$ becomes a global Hamiltonian and the spectrum has m non-degenerate levels and a $(2^n - m)$ -degenerate level.

For instance, when targeting the largest eigenvalue of ρ ($m=1$), the simple choice $r_j=1, \forall j$ achieves this goal. On the other hand, if one is interested in $m=n+1$ eigenvalues, then one can choose $r_j=r_1+(j-1)\delta$ with $r_1 \gg \delta$, which will ensure that the m -lowest energy levels, $\{E_1, E_1+r_1, E_1+r_1+\delta, \dots, E_1+r_1+(m-1)\delta\}$, are non-degenerate. Henceforth, we will assume that one has chosen $\{r_j\}_{j=1}^n$ such that the m -lowest energy levels are non-degenerate.

While fixed local cost functions do not exhibit barren plateaus for shallow depth, they still have several trainability issues. First, having less degeneracy in H_L leads to a more difficult optimization problem. Since degeneracy allows for additional freedom in the solution space, non-degeneracy constrains the possible solutions. Therefore, there is a tradeoff between engineering non-degeneracy (which allows one to distinguish more eigenvalues of ρ) versus keeping degeneracy (which allows for more solutions). Second, we expect both C_L and C_G to have a high density of local minima, especially for large m . This is because there will be partial solutions to the problem where one correctly assigns some eigenvalues of ρ to the right energy levels of the Hamiltonian, while incorrectly assigning other eigenvalues. This local minima issue is what motivates the following adaptive approach.

Adaptive Hamiltonian. Let us now we introduce an approach to adaptively update the VQSE Hamiltonian (and hence the cost function) based on information obtained via measurements during the optimization loop. This method allows us to mitigate the issues discussed in the previous section that arise for cost functions with fixed local or global Hamiltonians. Namely, the adaptive cost function solves the following three problems: (1) barren plateaus for shallow depth³⁸, (2) high density of local minima, (3) smaller solution space arising from non-degeneracies.

Consider a time-dependent Hamiltonian of the form

$$H(t) \equiv (1 - f(t))H_L + f(t)H_G(t), \quad (13)$$

where $f(t)$ is a real-valued function such that $f(0) = 0$, $f(1) = 1$, and H_L is a local Hamiltonian as in (12). We recall here that $t \in [0, 1]$ is a parameter that indicates the optimization loop run time. Moreover, we define the time-dependent global Hamiltonian

$$H_G(t) \equiv \mathbb{1} - \sum_{i=1}^m q_i |\mathbf{z}_i(t)\rangle \langle \mathbf{z}_i(t)|, \quad (14)$$

where the coefficients q_i are real and positive, and chosen in the same way as in (11). In addition, the states $|\mathbf{z}_i(t)\rangle$ are adaptively chosen throughout the optimization loop by preparing $\hat{\rho}$, measuring in the standard basis to obtain the sets \mathcal{L} and \mathcal{Z} , and updating $H_G(t)$ so that $\mathbf{z}_i(t) \in \mathcal{Z}$.

As schematically shown in Fig. 2a, in order to mitigate the barren plateau phenomena it is important to choose a function $f(t)$ which is not rapidly growing with t . Hence, for small t , $H(t) \sim H_L$ and the cost function will be trainable as it will not present a barren plateau. Then, as t increases, one can deal with the issue of local minima by updating $H_G(t)$. As depicted in the insets of Fig. 2a, adaptively changing $H_G(t)$ transforms local minima in the cost landscape into global minima. Then, by the end of the algorithm, we have $H(1) = H_G(1)$, and as shown in panel (b) of Fig. 2, the spectrum of H becomes highly degenerate and the dimension of the solution space increases. In section ‘Algorithm for the adaptive cost function’ of the ‘Methods’, we present an algorithm to illustrate how one can update $H(t)$.

We remark that ref. 40 proposed a method called adiabatically assisted VQE (AAVQE), which dynamically updates the VQE cost function by driving between a simple Hamiltonian to the non-trivial problem Hamiltonian. Note that the goals of AAVQE and our adaptive training method are different. Furthermore, in our method, one adaptively updates the cost function based on information obtained through measurements, while AAVQE does not use information gained during the optimization.

Operational meaning of the cost function. Here we discuss the operational meaning of the VQSE cost function, showing that small cost values imply small eigenvalue and eigenvector errors. Let $\{\tilde{|\lambda_i\rangle}\}_{i=1}^m$ be the set of the inferred eigenvector associated with every $\tilde{\lambda}_i$ in \mathcal{L} , and let $|\delta_i\rangle = \rho|\tilde{\lambda}_i\rangle - \tilde{\lambda}_i|\tilde{\lambda}_i\rangle$. We then define eigenvalue and eigenvector errors as follows:

$$\varepsilon_\lambda \equiv \sum_{i=1}^m (\lambda_i - \tilde{\lambda}_i)^2, \quad \varepsilon_v \equiv \sum_{i=1}^m \langle \delta_i | \delta_i \rangle. \quad (15)$$

Here $\langle \delta_i | \delta_i \rangle$ quantifies the component of $\rho|\tilde{\lambda}_i\rangle$ that is orthogonal to $|\tilde{\lambda}_i\rangle$, which follows from the following identity: $|\delta_i\rangle = (\mathbb{1} - |\tilde{\lambda}_i\rangle \langle \tilde{\lambda}_i|)\rho|\tilde{\lambda}_i\rangle$.

Then by using the Cauchy–Schwarz inequality, majorization conditions, and Schur convexity, we establish the following upper bound on eigenvalue and eigenvector errors (see section ‘Operational meaning of the cost function’ for more details):

$$\varepsilon_\lambda, \varepsilon_v \leq \text{Tr}[\rho^2] - \frac{(E_{m+1} - C(\theta))^2}{\sum_{i=1}^m (E_{m+1} - E_i)^2}, \quad (16)$$

where (E_1, \dots, E_m) are the m -smallest eigenenergies of H , and where for simplicity we have omitted the t dependence. Thus Eq. (16) provides an operational meaning to our cost function, as small values of the cost function lead to small eigenvalue and eigenvector errors.

Verification of solution quality

Let us show how to verify the results obtained from the VQSE algorithm. We remark that this verification step is optional,




Fig. 3 Ansatz diagram. **a** Layered hardware-efficient ansatz for $V(\theta)$. A single layer of the ansatz is composed of two-qubit gates $B_{\mu}(\theta_{\mu})$ acting on neighboring qubits. Shown is the case of two layers. **b** While there are many choices for each block $B_{\mu}(\theta_{\mu})$, in our numerics we employed two different parameterizations. Top: Each gate is composed of a controlled-Z gate preceded and followed by single-qubit rotations about the y -axis $R_y(\theta) = e^{i\theta\sigma_y}$. Bottom: Each gate is composed of a CNOT gate preceded and followed by a single-qubit rotation $G(\theta_1, \theta_2, \theta_3) = e^{i\theta_3\sigma_z/2}e^{i\theta_2\sigma_y/2}e^{i\theta_1\sigma_x/2}$. The number of parameters in θ increases linearly with the number of layers and the number of qubits n .

particularly because it requires $2n$ qubits, whereas the rest of VQSE only requires n qubits.

In section ‘Verification of solution quality’ of “Methods”, we prove the following useful bound on eigenvalue and eigenvector error:

$$\varepsilon_{\lambda}, \varepsilon_{\nu} \leq \text{Tr}[\rho^2] - \left(\sum_{i=1}^{\hat{m}} \tilde{\lambda}_i^2 + \frac{(1 - \sum_{i=1}^{\hat{m}} \tilde{\lambda}_i)^2}{2^n - \hat{m}} \right), \quad (17)$$

where one can take \hat{m} as any integer between m and 2^n . One can efficiently estimate the right-hand-side of (17) as follows. Given two copies of ρ , $\text{Tr}[\rho^2]$ can be estimated by a depth-two quantum circuit with classical post-processing that scales linearly with n^{41} . Moreover, since $\text{Tr}[\rho^2]$ is independent of $V(\theta)$, one only needs to compute it once (outside of the optimization loop). Estimating the $\tilde{\lambda}_i$ for $i = 1, \dots, \hat{m}$ essentially comes for free as part of the eigenvalue readout step of VQSE, where we note that taking $\hat{m} > m$ simply involves keeping track of the frequencies of more bitstrings (more than the m -largest) during this readout step. Finally, we remark that while Eq. (16) can also be used for verification, in section ‘Verification of solution quality’ we show that (17) provides a tighter bound, particularly as one increases \hat{m} .

Ansatz

While there are many possible choices for the ansatz $V(\theta)$, we are here restricted to state-agnostic ansatzes which do not require any a priori information about ρ . One such ansatz is the Layered Hardware Efficient Ansatz³⁶ shown in Fig. 3a. Here, $V(\theta)$ consists of a fixed number L of layers of two-qubit gates $B_{\mu}(\theta_{\mu})$ acting on alternating pairs of neighboring qubits. Figure 3b illustrates possible choices for $B_{\mu}(\theta_{\mu})$. Note that with this structure, the number of parameters in θ grows linearly with the n and L .

Let us remark that the Layered Hardware Efficient Ansatz can lead to trainability issues as the system size increases^{37,38}. Hence, different strategies have been proposed to mitigate such difficulties, such as learning to initialize parameters⁴², layer-by-layer training⁴³, and correlating the parameters⁴⁴. In addition, these methods can be combined with more sophisticated ansatzes, such as a variable-structure ansatz^{6,41} where the structure of the ansatz is not fixed, and where the gate placement becomes an optimizable hyper-parameter. This variable-structure approach has already been shown to improve performance in the context of extracting the eigensystem of a quantum state⁶.

Finally, since VQSE optimization corresponds to an energy minimization problem, a natural ansatz that can also be used to mitigate trainability issues is the quantum alternating operator ansatz (QAOA)^{3,45}. Specifically, one could employ $H(t)$ as the

problem Hamiltonian in the QAOA and use a standard mixing Hamiltonian. While we do not employ this ansatz in our heuristics, it is nevertheless of interest for future work.

Optimization

Regarding the optimization of the parameters θ , while gradient-free methods are an option^{46,47}, there has been recent evidence that gradient-based methods can perform better^{48–51}. Moreover, as shown in refs. ^{52,53}, for cost functions like (1), gradients can be analytically determined (see section ‘Gradient of the cost function’ in the “Methods” for an explicit derivation of the gradient formula). Therefore, in our heuristics, we employ gradient-based optimization.

Numerical Implementations

Here we present the numerical results obtained from implementing VQSE. We first employ VQSE to estimate the spectrum of quantum states of different dimensions and compare the performance of cost functions based on the global, local, and adaptive Hamiltonians discussed in section ‘Cost functions’. Then we use VQSE for error mitigation of the W -state preparation circuit. Finally, we implement VQSE for entanglement spectroscopy on the ground-state of an XY -spin chain, which allows us to detect the presence of quantum critical points.

VQSE for quantum principal component analysis. Figure 4 presents the results of implementing VQSE to estimate the six largest eigenvalues ($m = 6$) of quantum states with $n = 6, 8$, and 10 qubits. In all cases, we have $\text{rank}(\rho) = 16$, as the states were prepared by randomly entangling the system qubits with four ancillary qubits, which were later traced out. Moreover, we chose ρ to be real and not sparse in the standard basis.

In our heuristics, we used the layered hardware efficient ansatz of Fig. 3(b, top), and we employed the fixed-local, fixed-global, and adaptive cost functions of section ‘Cost functions’. The termination condition was stated in terms of the maximum number of iterations in the optimization loop. Hardware noise and finite sampling were not included in these heuristics. (The next subsection shows heuristics with noise.) For the fixed local cost function, we chose the $\{r_j\}_{j=1}^n$ in (12) so that the first six energy eigenvalues of H_L were non-degenerate. Moreover, we defined the fixed global Hamiltonian such that the first six energy levels (i.e., associated eigenvectors and spectral gaps) coincided with those of H_L . Finally, the adaptive Hamiltonian was constructed according to the procedure described above, and more specifically, in Algorithm 1 in the “Methods” section.

Since for these examples, we can calculate the exact eigenvalues λ_i , we compute and plot the following quantities which we use as figures of merit for the performance of the VQSE




Fig. 4 Relative and absolute error versus the number of iterations. We implemented VQSE for states of: **a** $n = 6$, **b** $n = 8$, and **c** $n = 10$ qubits. In all cases, the ansatz for $V(\theta)$ was given by three layers of the Layered Hardware Efficient Ansatz of Fig. 3(b, top). Each curve represents the absolute or relative error (denoted Abs error or Rel error, respectively) of (18) obtained by training $V(\theta)$ when employing an adaptive, fixed-local, or fixed-global Hamiltonian. The number of iterations was 330 for (a) and 360 for (b) and (c). For the adaptive runs, we employed Algorithm 1, with the Hamiltonian being updated every 30 iterations. In each case, the adaptive approach performs the best as it achieves the smallest errors.

algorithm:

$$\varepsilon_\lambda \equiv \sum_{i=1}^6 (\lambda_i - \tilde{\lambda}_i)^2, \quad \varepsilon_r \equiv \sum_{i=1}^6 (\lambda_i - \tilde{\lambda}_i)^2 / \lambda_i^2. \quad (18)$$

Here ε_λ and ε_r respectively quantify absolute error and relative error in estimating the exact eigenvalues. We remark that these two quantities provide different information: The absolute error is biased towards the error in estimating the large eigenvalues of ρ , while on the other hand, the relative error is more sensitive to errors in estimating the small eigenvalues of ρ .

Figure 4 plots the relative and absolute errors versus number of iterations (with the total number of iterations fixed). While we performed many runs, these plots show only the run that achieved the lowest absolute error. For all system sizes considered ($n = 6, 8$ and 10), VQSE achieves smaller relative and absolute errors when employing the adaptive Hamiltonian approach than using a fixed Hamiltonian. For $n = 6$, the errors obtained by adaptively updating $H(t)$ are two orders of magnitude smaller than those obtained with fixed Hamiltonians, while for $n = 8$ they are one order of magnitude smaller. As shown in Fig. 4c for $n = 10$, the adaptive Hamiltonian approach achieves error of the order: $\sim 10^{-5}$ for the relative error, and $\sim 10^{-7}$ for the absolute error, and again outperforms the fixed local Hamiltonian approaches. We here finally remark that we can use Eq. (9) to determine the number of shots needed to guarantee that with a probability larger than 99% the relative error induced by finite sampling is smaller than 0.001. Namely, we find that one needs a number of shots larger than 50.2K, which is well within the order of magnitude of shots regularly used.

It is natural to ask whether the runs shown in Fig. 4 are representative of the algorithm performance. To provide an analysis of the average VQSE performance, we plot in Fig. 5 the runs-per-success versus $1/\varepsilon_\lambda$ for each of the aforementioned examples. Here, runs-per-success is defined as the total number of runs divided by the number of runs with an absolute error smaller than a target ε_λ .

From all three panels in Fig. 5 we see that for large $1/\varepsilon_\lambda$, the adaptive Hamiltonian always has the best performance as it requires less runs-per-success to achieve smaller errors.

Finally, it is interesting to note from the insets of Fig. 5 that there is a regime where the run time has a linear dependence on $\log(1/\varepsilon_\lambda)$ when employing an adaptive approach. This suggests that VQSE may perform quite efficiently for large ε_λ . However, the linear dependence breaks down for small ε_λ , where the number of runs-per-success seems to grow exponential with $1/\varepsilon_\lambda$. Despite, such growth, for up to 8 qubits we only need 100 repetitions to achieve an error of order 10^{-6} . We leave for future work a more detailed study of the dependence of runs-per-success for small error. Finally, the results presented in Figs. 4 and 5 suggest that for a sufficiently large value of number of iterations, the adaptive Hamiltonian approach outperforms the fixed Hamiltonian approaches as it requires the least number of iterations to converge to very small values of ε_λ .

Error mitigation. Here we discuss an important application of the VQSE algorithm for error mitigation. Quantum state preparation circuits (gate sequences U which prepare a target state $|\psi\rangle$) are used as subroutines in many quantum algorithms. However, since current quantum computers are noisy, all state preparation circuits produce mixed states ρ . If there is little enough incoherent noise, we can expect that the largest eigenvalue of ρ is associated with $|\psi\rangle$. Here we show that VQSE can be implemented to re-purify ρ and estimate $|\psi\rangle$. Naturally, when running the VQSE eigenvector preparation circuit, noise will also produce a mixed state σ . However, if the depth of $V(\theta)$ is shorter than the depth of U , one can obtain a higher fidelity between σ and $|\psi\rangle$ in comparison to the fidelity between ρ and $|\psi\rangle$. In this case one can mitigate errors by replacing the state preparation circuit by the VQSE eigenvector preparation circuit.

Let us now consider the three qubit W -state preparation circuit from ref. ⁵⁴, section 2.2 (see also ref. ⁵⁵). By employing a noisy quantum computer simulator with the noise profile of IBM's Melbourne processor⁵⁶, we find that the fidelity between ρ and the exact W state $|\psi\rangle$ is $F(\rho, |\psi\rangle) \approx 0.785$. We then train 10 instances of VQSE with two layers of the ansatz in Fig. 3(b, bottom) and with a termination condition of 50 iterations. Moreover, we employ the adaptive Hamiltonian, where we update $H(t)$ every 10 iterations according to Algorithm 1. Figure 6 shows the average cost function value and average fidelity between $|\psi\rangle$ and the state σ obtained by running the VQSE eigenvector preparation circuit. As the number of iterations

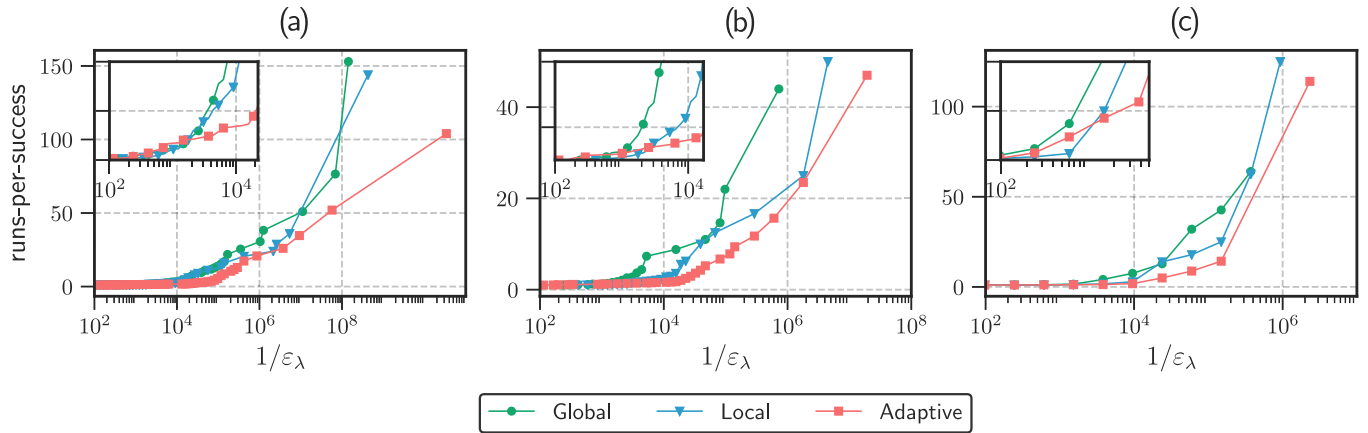


Fig. 5 Runs-per-success versus inverse absolute error $1/\varepsilon_\lambda$. We implemented VQSE for the states of: **a** $n = 6$, **b** $n = 8$, and **c** $n = 10$ qubits corresponding to Fig. 4. The insets depict the same data in the small $1/\varepsilon_\lambda$ regime. Runs-per-success is defined as the total number of runs divided by the number of runs with a relative error smaller than a target ε_λ . For all three cases, we can see that as $1/\varepsilon_\lambda$ increases, the adaptive Hamiltonian has the lowest number of runs-per-success, and hence the best performance. In all cases, the x -axis is plotted on a log scale.

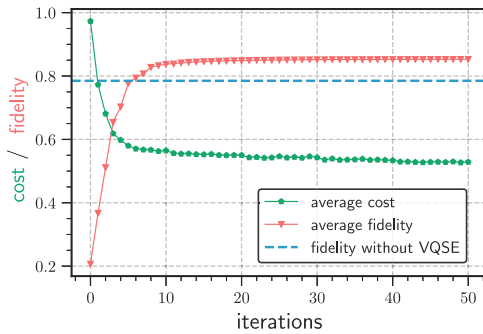


Fig. 6 Cost function value and fidelity versus number of iterations. We implement VQSE for error mitigation of the three qubit W -state preparation circuit. The input state ρ corresponds to the mixed state obtained by running the W -state preparation circuit on a noisy simulator. The dashed line corresponds to the fidelity $F(\rho, |\psi\rangle)$ between ρ and the exact W state $|\psi\rangle$. For each iteration step, we compute the fidelity $F(\sigma, |\psi\rangle)$, where the mixed state σ is obtained by running the VQSE eigenvector preparation circuit on the noisy simulator. Curves depict the average of 10 instances of the algorithm. As the number of iterations increases the cost function value decreases, which implies that we are able to train $V(\theta)$ in the presence of noise. After a few iterations of the VQSE optimization loop, we find $F(\sigma, |\psi\rangle) > F(\rho, |\psi\rangle)$.

increases, the cost value tends to decrease, showing that we are able to train in the presence of noise. Moreover, we also see that $F(\sigma, |\psi\rangle)$ increases and saturates at a value larger than $F(\rho, |\psi\rangle)$, namely at 0.853, hence showing that we are in fact mitigating the effect of noise. This can be explained by the fact that we reduced the circuit depth, as our ansatz contains two CNOTs, while the textbook circuit contains three CNOTs.

Entanglement spectroscopy. We now discuss the possibility of employing VQSE to compute the entanglement spectrum of a state ρ which is obtained as the reduced state of a bipartite quantum system $|\psi_{AB}\rangle$, i.e., $\rho = \text{Tr}_B |\psi_{AB}\rangle \langle \psi_{AB}|$. Let d denote the dimension of ρ . The entanglement spectrum⁵⁷ refers to the collection $\{\lambda_k\}_{k=1}^d$ of eigenvalues of ρ , and as discussed in ref. 58, entanglement spectroscopy is a useful tool to analyze states $|\psi_{AB}\rangle$ prepared by simulating many-body systems on a quantum computer. Specifically, the entanglement spectrum is useful to study the bipartite entanglement, as it contains more universal signatures than the von Neumann entropy alone⁵⁷, and it can detect the presence of quantum critical points^{59,60}.

Let us now consider an $N = 8$ spin-1/2 cyclic chain interacting through uniform XY first-neighbor Heisenberg coupling in the presence of a non-transverse magnetic field. The Hamiltonian of the system is

$$H = - \sum_j (h_x S_j^x + h_z S_j^z + J_x S_j^x S_{j+1}^x + J_y S_j^y S_{j+1}^y), \quad (19)$$

where j labels the site in the chain, S_j^μ the spin operator (with $\mu = x, y, z$), J_μ the coupling strength, and h_μ the magnetic fields. Here, $J_\mu > 0$ leads to ferromagnetic (FM) coupling, while $J_\mu < 0$ to antiferromagnetic (AFM) coupling. As shown in refs. 60,61, for specific values of the fields h_μ (known as *factorizing fields*) the Hamiltonian in (19) presents quantum critical points known as "factorization" points. At the non-transverse factorizing field, the ground-state of H becomes a separable non-degenerate state such that one of its eigenvalues is exactly equal to one, while the rest are exactly zero.

In Fig. 7a, c, we show results of implementing VQSE with an adaptive Hamiltonian to compute the three largest eigenvalues of the state ρ defined as the reduced state of 4 neighboring spins obtained from the ground state of (19). For simplicity, we have parametrized the fields as $(h_z, h_x) = h(\cos(\gamma), \sin(\gamma))$ with γ fixed. Specifically, in Fig. 7a, c we plot the estimated eigenvalues versus the field magnitude h for a system with FM and AFM couplings, respectively. Moreover, dashed lines indicate the exact eigenvalues. For each field value, we run 8 instances of VQSE, and even for such a small number of runs, the estimated eigenvalues give good approximations as we get relative errors which in general are of the order of $\sim 10^{-2}$.

In Fig. 7b, d, we show the same data as in (a) and (c) but the y axis is plotted on a logarithmic scale, and where instead of plotting the largest eigenvalue λ_1 , we plot $1 - \lambda_1$. For the FM (AFM) case, there is a factorization point at $h/J_x \approx 0.76$ ($h/J_x \approx 1.43$). As depicted in these panels, around critical points we correctly find $\lambda_1 \approx 1$, and $\lambda_2, \lambda_3 \approx 0$. These results show that VQSE can detect quantum critical factorization points.

DISCUSSION

In the NISQ era, every qubit and every gate counts. Wasteful usage of qubits or gates will ultimately limit the problem size that an algorithm can solve. In this work, we presented an algorithm for extracting the eigensystem of a quantum state ρ that is as frugal as we could imagine, with respect to qubit count.

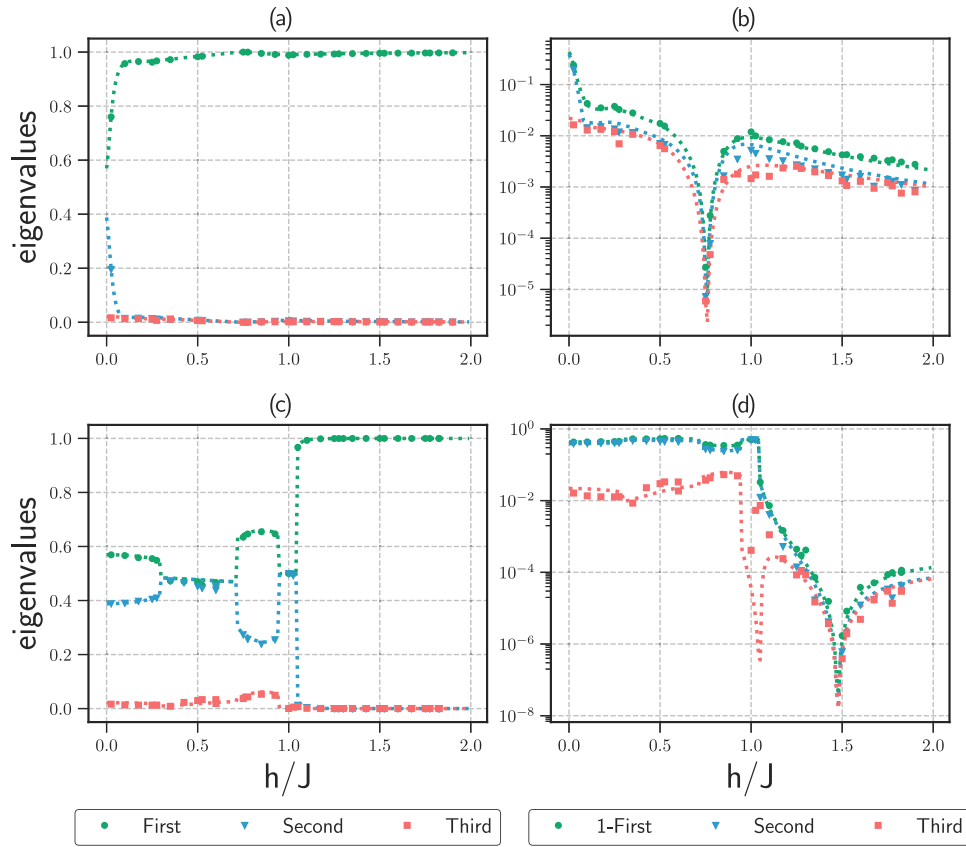


Fig. 7 Exact and estimated eigenvalues versus field value, for the VQSE entanglement spectroscopy implementations. The input state ρ is given as the reduced state of 4 neighboring qubits from the ground state of (19). Top and bottom rows correspond to ferromagnetic and antiferromagnetic couplings, respectively. Dashed curves represent the exact three largest eigenvalues of ρ , while plot markers indicate the VQSE estimated eigenvalues. In **a** and **c** we see that VQSE can accurately estimate the eigenvalues. In **b** and **d** we plot $1 - \lambda_1$, and the y axis is on a log scale. Here we see that the quantum critical factorization points are detected at $h/J_x \approx 0.76$ and $h/J_x \approx 1.43$ in **(b)** and **(d)**, respectively, since at those points we have $\bar{\lambda}_1 \approx 1$, and $\bar{\lambda}_2, \bar{\lambda}_3 \approx 0$.

We introduced the VQSE, which estimates the m -largest eigenvalues and associated eigenvectors of ρ , using only a single copy of ρ , and hence only n qubits per iteration of the VQSE. VQSE exploits the mathematical connection between diagonalization and majorization to define an efficiently computable cost function as the expectation value of a Hamiltonian. We derived an operational meaning of this cost function as a bound on eigensystem error. Furthermore, we introduced a training method that involved adaptively updating the VQSE cost function based on the information gained from measurements performed throughout the optimization. This was aimed at addressing both barren plateaus and local minima in the cost landscape.

We have numerically implemented VQSE for several applications. We showed that VQSE can be employed for PCA by implementing the VQSE algorithm on states of $n = 6, 8, \text{ and } 10$ qubits to estimate the six largest eigenvalues. Our numerical results (Figs. 4 and 5) indicate that our adaptive cost function approach leads to smaller errors than the ones obtained by training a fixed cost function. We also showed (Fig. 7) that one can detect quantum critical points by performing entanglement spectroscopy with the eigenvalues obtained via VQSE. Finally, we employed VQSE to mitigate errors that occur during the W -state preparation circuit. This involved running VQSE on a noisy simulator to re-purify the state, i.e., find the circuit that prepares the eigenvector with the largest eigenvalue. We found (Fig. 6) that the re-purified state obtained by VQSE improved the fidelity with the target W state, and hence reduced the effects of noise.

Comparison to literature

Since VQSE only requires n qubits, it is as qubit frugal as it can possible be when compared to other algorithms for the same task, such as quantum PCA (qPCA)³³, VQSD⁶, and quantum state singular value decomposition (QSVD)²⁷. The quantum phase estimation and density matrix exponentiation primitives in qPCA make it difficult to implement in the near term⁶², and this is supported by attempted implementation in ref. 6 that resulted in poor performance. On the other hand, VQSD and QSVD are variational algorithms and hence have the possibility of lower-depth requirements. But they still need to employ a larger number of qubits than VQSE. Specifically, VQSD needs to perform the so-called Diagonalized Inner Product Test⁵ which requires two copies of ρ , i.e., requires twice as many qubits as VQSE. In addition, it is also worth noting that VQSD is vulnerable to noise, since any asymmetry between the noise acting of each copy of ρ will affect the result of the algorithm. Finally, in QSVD, one needs to either compute or have access to a purification $|\psi\rangle$ of ρ . Hence QSVD requires a number of qubits between n and $2n$. Moreover, we expect that noise will be a bigger issue for QSVD than for VQSE, since in practice the assumption that one has a pure state in QSVD can often be violated due to incoherent noise during state preparation.

Quantum-inspired classical algorithms⁶³ for PCA can perform well in practice, provided that the matrix has a very large dimension, low rank, and low condition number⁶⁴. We note that VQSE does not have such limitations, except the fact that VQSE yields results with high accuracy for low-rank states. Here is also

paramount to recall that recent results have shown that an exponential advantage is still possible for PCA^{65,66}, even for near-term algorithms, as the quantum-inspired classical algorithms are artificially given too much power via access to quantum state amplitudes. Thus, in view of these recent results, VQSE can be useful in the quest for achieving a quantum speedup with quantum PCA, particularly for analysis of quantum data. For the case of classical data, the success of VQSE for PCA relies on the efficiency of preparing a quantum state corresponding to the covariance matrix of the classical data⁶⁷. In addition, VQSE has applications not only for PCA but also for other tasks such as entanglement spectroscopy and error mitigation on NISQ devices. For error mitigation, we leave for future work combining our approach with other error mitigation techniques such as Virtual Distillation^{68,69}, which also seeks to re-purify noisy quantum states.

Future directions

Due to the rapid rise of VQE², much research has gone into how to prepare ground and excited states on NISQ devices. However, more research is needed on how to characterize these states, once prepared. This is where VQSE comes in, as VQSE can extract the entanglement spectra of these states and hence characterize important properties like topological order⁵⁷. Hence it is worth exploring in the future the idea of pairing up the VQE and VQSE algorithms, where VQSE is implemented immediately after VQE.

Furthermore, VQSE has immediate application for estimating the fidelity of two quantum states with reduced resource requirements. This is because an algorithm was previously introduced⁸ to estimate fidelity by using state diagonalization as a subroutine, and hence VQSE can provide a more efficient version of this subroutine.

A crucial technical idea in this work was our adaptively-updated cost function, which improved optimization performance. It is worth investigating whether this adaptive method can improve the performance of other VQAs²⁻²⁷.

Another direction to explore is whether VQSE exhibits noise resilience¹⁷. We suspect this to be true given the similar structure of VQSE and the variational quantum compiling algorithms investigated in ref. ¹⁷.

This is important as we are proposing that VQSE will be a useful tool for error mitigation. Namely, we envision that VQSE could be used as a subroutine to improve the accuracy of several quantum algorithms. For example, one could use VQSE to re-purify the noisy quantum state obtained as the outcome of the VQE algorithm. Alternatively, one could periodically perform VQSE whilst running a dynamical quantum simulation on a NISQ device, which would re-purify the state as it is evolving in time. This could allow one to simulate long-time dynamics, i.e., times significantly beyond the coherence time of a NISQ device.

METHODS

Operational meaning of the cost function

In this section, we provide a derivation for Eq. (16). First, we rewrite the eigenvalue error in Eq. (15) as follows:

$$\varepsilon_\lambda = \boldsymbol{\lambda}^m \cdot \boldsymbol{\lambda}^m + \tilde{\boldsymbol{\lambda}}^m \cdot \tilde{\boldsymbol{\lambda}}^m - 2\boldsymbol{\lambda}^m \cdot \tilde{\boldsymbol{\lambda}}^m, \quad (20)$$

where $\boldsymbol{\lambda}^m \equiv (\lambda_1, \dots, \lambda_m)$ and $\tilde{\boldsymbol{\lambda}}^m \equiv (\tilde{\lambda}_1, \dots, \tilde{\lambda}_m)$. Since the eigenvalues of a positive semidefinite operator majorize its diagonal elements, we have that $\boldsymbol{\lambda}^m \succ \tilde{\boldsymbol{\lambda}}^m$. Moreover, from the Schur convexity property of the dot product with an ordered vector, it follows that $\boldsymbol{\lambda}^m \cdot \tilde{\boldsymbol{\lambda}}^m \geq \tilde{\boldsymbol{\lambda}}^m \cdot \tilde{\boldsymbol{\lambda}}^m$, which further implies the following inequality:

$$\varepsilon_\lambda \leq \boldsymbol{\lambda}^m \cdot \boldsymbol{\lambda}^m - \tilde{\boldsymbol{\lambda}}^m \cdot \tilde{\boldsymbol{\lambda}}^m. \quad (21)$$

Similarly, from Eq. (15) we get

$$\varepsilon_\nu \leq \boldsymbol{\lambda}^m \cdot \boldsymbol{\lambda}^m - \tilde{\boldsymbol{\lambda}}^m \cdot \tilde{\boldsymbol{\lambda}}^m, \quad (22)$$

where we again used the fact that the eigenvalues of a positive semidefinite operator majorize its diagonal elements, and hence $\boldsymbol{\lambda}^m \cdot \boldsymbol{\lambda}^m \geq \sum_{i=1}^m (\lambda_i |\rho^2| \tilde{\lambda}_i)$.

We recall from Eq. (3) that the VQSE cost function can be expressed as $C = \sum_{i=1}^d E_i p_i$, where we omit the $\boldsymbol{\theta}$, and t dependence of C . Therefore, the following chain of inequalities hold:

$$\begin{aligned} C &\geq \sum_{i=1}^m E_i p_i + E_{m+1} \sum_{i>m} p_i \\ &= E_{m+1} - \left(\sum_{i=1}^m p_i (E_{m+1} - E_i) \right) \\ &\geq E_{m+1} - \sqrt{\left(\sum_{i=1}^m p_i^2 \right) \left(\sum_{i=1}^m (E_{m+1} - E_i)^2 \right)}, \end{aligned} \quad (23)$$

where $d = 2^n$. The first inequality follows the fact that $E_i \geq E_{m+1}, \forall i \geq m+1$ and $\sum_{i>m} p_i = 1 - \sum_{i=1}^m p_i$. The second inequality follows from the Cauchy-Schwarz inequality for the dot product of two vectors $|\mathbf{u} \cdot \mathbf{v}| \leq |\mathbf{u}| |\mathbf{v}|$. By combining Eq. (23) with the fact that $\sum_{i=1}^m \tilde{\lambda}_i^2 \geq \sum_{i=1}^m p_i^2$ (since $\tilde{\lambda}_i \in \mathcal{L}$ are the largest diagonal elements of $\tilde{\rho}$), we find that

$$\sqrt{\sum_{i=1}^m \tilde{\lambda}_i^2} \geq \frac{E_{m+1} - C}{\sqrt{\sum_{i=1}^m (E_{m+1} - E_i)^2}}. \quad (24)$$

Using the fact that $\boldsymbol{\lambda}^m \cdot \boldsymbol{\lambda}^m \leq \boldsymbol{\lambda} \cdot \boldsymbol{\lambda} = \text{Tr}[\rho^2]$, we obtain the following equality from (24)

$$\boldsymbol{\lambda}^m \cdot \boldsymbol{\lambda}^m - \tilde{\boldsymbol{\lambda}}^m \cdot \tilde{\boldsymbol{\lambda}}^m \leq \text{Tr}[\rho^2] - \frac{(E_{m+1} - C)^2}{\sum_{i=1}^m (E_{m+1} - E_i)^2}.$$

Combining this with (21) and (22) leads to (16).

Verification of solution quality

Here we provide a proof of Eq. (17), and we show that this bound is tighter than the bound in (16). From the definition of the eigenvalue and eigenvector error in (15), it is straightforward to see that $\varepsilon_\lambda \leq \sum_{i=1}^d (\lambda_i - \tilde{\lambda}_i)^2$, and $\varepsilon_\nu \leq \sum_{i=1}^d \langle \delta_i | \delta_i \rangle$, where $d = 2^n$. By following a procedure similar to the one employed in deriving (22), we find

$$\varepsilon_\lambda \leq \boldsymbol{\lambda} \cdot \boldsymbol{\lambda} - \tilde{\boldsymbol{\lambda}} \cdot \tilde{\boldsymbol{\lambda}}, \quad (25)$$

where we recall that $\boldsymbol{\lambda}$ and $\tilde{\boldsymbol{\lambda}}$ denote d -dimensional vectors of ordered exact and estimated eigenvalues of ρ , respectively. Moreover, from $|\delta_i\rangle = (1 - |\tilde{\lambda}_i\rangle \langle \tilde{\lambda}_i|) \rho |\tilde{\lambda}_i\rangle$, it is straightforward to get

$$\varepsilon_\nu \leq \sum_{i=1}^d \langle \tilde{\lambda}_i | \rho^2 | \tilde{\lambda}_i \rangle - \tilde{\boldsymbol{\lambda}} \cdot \tilde{\boldsymbol{\lambda}} = \boldsymbol{\lambda} \cdot \boldsymbol{\lambda} - \tilde{\boldsymbol{\lambda}} \cdot \tilde{\boldsymbol{\lambda}}, \quad (26)$$

where we used the fact that $\sum_{i=1}^d \langle \tilde{\lambda}_i | \rho^2 | \tilde{\lambda}_i \rangle = \text{Tr}[\rho^2]$, which follows from the invariance of trace under a basis transformation.

Let $\hat{\boldsymbol{\lambda}} = (\hat{\lambda}_1, \dots, \hat{\lambda}_m, \frac{1 - \sum_{i=1}^m \hat{\lambda}_i}{2^n - m}, \dots, \frac{1 - \sum_{i=1}^m \hat{\lambda}_i}{2^n - m})$, with $\hat{m} > m$, be a vector majorized by $\tilde{\boldsymbol{\lambda}}$, i.e., $\tilde{\boldsymbol{\lambda}} \succ \hat{\boldsymbol{\lambda}}$. Since the dot product with an ordered vector is a Schur convex function, we have $\hat{\boldsymbol{\lambda}} \cdot \tilde{\boldsymbol{\lambda}} \leq \tilde{\boldsymbol{\lambda}} \cdot \tilde{\boldsymbol{\lambda}} \leq \tilde{\boldsymbol{\lambda}} \cdot \hat{\boldsymbol{\lambda}}$, which further implies the following inequality:

$$\boldsymbol{\lambda} \cdot \boldsymbol{\lambda} - \tilde{\boldsymbol{\lambda}} \cdot \tilde{\boldsymbol{\lambda}} \leq \boldsymbol{\lambda} \cdot \boldsymbol{\lambda} - \left(\sum_{i=1}^{\hat{m}} \hat{\lambda}_i^2 + \frac{(1 - \sum_{i=1}^{\hat{m}} \hat{\lambda}_i)^2}{2^n - \hat{m}} \right). \quad (27)$$

This inequality can be combined with (25) and (26) to obtain the bound in (17).

We now show that (17) is tighter than (16). Specifically, we prove that the negative term in the right-hand side of (17) is larger than the one in (16). Consider the following chain of inequalities:

$$\begin{aligned} \left(\sum_{i=1}^{\hat{m}} \hat{\lambda}_i^2 + \frac{(1 - \sum_{i=1}^{\hat{m}} \hat{\lambda}_i)^2}{2^n - \hat{m}} \right) &\geq \sum_{i=1}^m \tilde{\lambda}_i^2 \\ &\geq \frac{(E_{m+1} - C)^2}{\sum_{i=1}^m (E_{m+1} - E_i)^2}, \end{aligned}$$

where we used $\hat{m} > m$, and where the last inequality follows from (24).

Gradient of the cost function

Here we show that the partial derivative of (4) with respect to an angle θ_v is given by

$$\frac{\partial C(t, \boldsymbol{\theta})}{\partial \theta_v} = \frac{1}{2} (\text{Tr}[H(t)V(\boldsymbol{\theta}_+) \rho V^\dagger(\boldsymbol{\theta}_+)] - \text{Tr}[H(t)V(\boldsymbol{\theta}_-) \rho V^\dagger(\boldsymbol{\theta}_-)]). \quad (28)$$

Writing $\boldsymbol{\theta} = (\theta_1, \dots, \theta_v, \dots)$, then $\boldsymbol{\theta}_\pm$ are simply given by $\boldsymbol{\theta}_\pm = (\theta_1, \dots, \theta_v \pm \pi/2, \dots)$, which shows that the gradient values are efficiently accessible by shifting the parameters in $\boldsymbol{\theta}$ and measuring the expectation value $\langle H(t) \rangle$.

Let us consider the layered hardware efficient ansatz of Fig. 3a. Here $V(\boldsymbol{\theta})$ consists of a fixed number L of layers of 2-qubit gates $B_\mu(\boldsymbol{\theta}_\mu)$ acting on alternating pairs of neighboring qubits. Moreover, $B_\mu(\boldsymbol{\theta}_\mu)$ can always be expressed as a product of η_μ gates from a given alphabet $\mathcal{A} = \{U_k(\theta_k)\}$ as $B_\mu(\boldsymbol{\theta}_\mu) = U_{\eta_\mu}(\theta_{\eta_\mu}^\mu) \dots U_\nu(\theta_\nu^\mu) \dots U_1(\theta_1^\mu)$. (29)

Here θ_μ^ν are continuous parameters, and we can always write without loss of generality $U_k(\theta) = R_k(\theta)T_k$ where $R_k(\theta) = e^{i\theta\sigma_k/2}$ is a single qubit rotation and T_k is an unparametrized gate.

We can then compute $\partial_\nu B_\mu(\boldsymbol{\theta}_\mu) \equiv \partial B_\mu(\boldsymbol{\theta}_\mu)/\partial \theta_\nu^\mu$ as

$$\partial_\nu B_\mu(\boldsymbol{\theta}_\mu) = \frac{i}{2} U_{\eta_\mu}(\theta_{\eta_\mu}^\mu) \dots \sigma_\nu U_\nu(\theta_\nu^\mu) \dots U_1(\theta_1^\mu). \quad (30)$$

Then, without loss of generality let us write $V(\boldsymbol{\theta}) = V_L(\boldsymbol{\theta}_L)B_\mu(\boldsymbol{\theta}_\mu)V_R(\boldsymbol{\theta}_R)$, where $V_L(\boldsymbol{\theta}_L)$ and $V_R(\boldsymbol{\theta}_R)$ contain all gates in $V(\boldsymbol{\theta})$ except for $B_\mu(\boldsymbol{\theta}_\mu)$. By noting that $\partial_\nu V(\boldsymbol{\theta}) = V_L(\boldsymbol{\theta}_L)\partial_\nu B_\mu(\boldsymbol{\theta}_\mu)V_R(\boldsymbol{\theta}_R)$, we have

$$\begin{aligned} \partial_\nu C &= \text{Tr}[HV_L\partial_\nu B_\mu V_R \rho V_R^\dagger B_\mu^\dagger V_L^\dagger] \\ &+ \text{Tr}[HV_L B_\mu V_R \rho V_R^\dagger \partial_\nu B_\mu^\dagger V_L^\dagger], \end{aligned}$$

where we omitted the parameter dependence for simplicity. Then, from Eq. (30) and using the following identity (which is valid for any matrix A)

$$i[\sigma_\nu A] = R_\nu\left(-\frac{\pi}{2}\right)AR_\nu^\dagger\left(-\frac{\pi}{2}\right) - R_\nu\left(\frac{\pi}{2}\right)AR_\nu^\dagger\left(\frac{\pi}{2}\right), \quad (31)$$

where $R_k(\theta) = e^{i\theta\sigma_k/2}$, we obtain

$$\frac{\partial C(t, \boldsymbol{\theta})}{\partial \theta_\nu} = \frac{1}{2} (\text{Tr}[H(t)V(\boldsymbol{\theta}_+) \rho V^\dagger(\boldsymbol{\theta}_+)] - \text{Tr}[H(t)V(\boldsymbol{\theta}_-) \rho V^\dagger(\boldsymbol{\theta}_-)]). \quad (32)$$

Algorithm for the adaptive cost function

Algorithm 1 shows a simple adaptive strategy that illustrates how one can update $H(t)$. Specifically, we consider the case when $f(t)$ is a stepwise function. In addition, we define the VQSE optimization loop termination condition in terms of the maximum number of iterations allowed N_{\max} . We also define an updating parameter s (with N_{\max}/s being an integer) such that we update $H_G(t)$ every s steps. Finally, here we use the term optimizer, denoted as opt , as a function that takes as inputs a set of parameters $\boldsymbol{\theta}$ and a cost function $C(t, \boldsymbol{\theta})$ (or the gradient of the cost for gradient-based optimization) and returns an updated set of parameters that attempts to solve the minimization problem of Eq. (5).

Algorithm 1. Adaptive cost function with stepwise schedule $f(t)$

Input: state ρ ; trainable unitary $V(\boldsymbol{\theta})$; integer m ; timestep $\delta t = 1/N_{\max}$; adapting stepsize $t_s = 1/s$; local time-independent Hamiltonian H_L ; a set of constant parameters $\{q_i\}_{i=1}^m$; classical optimizer opt .

Output: parameters $\boldsymbol{\theta}_{\text{opt}}$ which minimize the cost function, i.e., $\boldsymbol{\theta}_{\text{opt}} = \arg \min_{\boldsymbol{\theta}} C(\boldsymbol{\theta})$.

Init: randomly choose a set of initial parameters $\boldsymbol{\theta}$; $H(t) \leftarrow H_L$; $t \leftarrow \delta t$

```

1: while  $t \leq 1$  do
2:   if  $t$  if divisible by  $t_s$ , then
3:     measure  $V(\boldsymbol{\theta})\rho V^\dagger(\boldsymbol{\theta})$  in the standard basis.
       define the sets  $\mathcal{L}$  and  $\mathcal{Z}$ 
4:      $H_G(t) \leftarrow \mathbb{1} - \sum_{i=1}^m q_i |\mathbf{z}_i\rangle\langle \mathbf{z}_i|$  with  $\mathbf{z}_i \in \mathcal{Z}$ 
5:      $H(t) \leftarrow (1-t)H_L + tH_G(t)$ 
6:     run  $\text{opt}$  with  $C$  and  $\boldsymbol{\theta}$  as input, and  $\boldsymbol{\theta}_{\min}$  as output
7:      $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}_{\min}$ 
8:      $t \leftarrow t + \delta t$ 
9:   if  $t = 1$  then
10:     $\boldsymbol{\theta}_{\text{opt}} \leftarrow \boldsymbol{\theta}$ 
Return:  $\boldsymbol{\theta}_{\text{opt}}$ 

```

DATA AVAILABILITY

Data generated and analyzed during the current study are available from the corresponding author upon reasonable request.

Received: 3 June 2020; Accepted: 4 August 2022;

Published online: 21 September 2022

REFERENCES

- Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
- Peruzzo, A. et al. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **5**, 4213 (2014).
- Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm. Preprint at <https://arxiv.org/abs/1411.4028> (2014).
- Johnson, P. D., Romero, J., Olson, J., Cao, Y. & Aspuru-Guzik, A. QVECTOR: An algorithm for device-tailored quantum error correction. Preprint at <https://arxiv.org/abs/1711.02249> (2017).
- Romero, J., Olson, J. P. & Aspuru-Guzik, A. Quantum autoencoders for efficient compression of quantum data. *Quantum Sci. Technol.* **2**, 045001 (2017).
- LaRose, R., Tikku, A., O'Neel-Judy, Étude, Cincio, L. & Coles, P. J. Variational quantum state diagonalization. *Npj Quantum Inf.* **5**, 57 (2019).
- Arrasmith, A., Cincio, L., Sornborger, A. T., Zurek, W. H. & Coles, P. J. Variational consistent histories as a hybrid algorithm for quantum foundations. *Nat. Commun.* **10**, 3438 (2019).
- Cerezo, M., Poremba, A., Cincio, L. & Coles, P. J. Variational quantum fidelity estimation. *Quantum* **4**, 248 (2020).
- Jones, T., Endo, S., McArdle, S., Yuan, X. & Benjamin, S. C. Variational quantum algorithms for discovering hamiltonian spectra. *Phys. Rev. A* **99**, 062304 (2019).
- Yuan, X., Endo, S., Zhao, Q., Li, Y. & Benjamin, S. C. Theory of variational quantum simulation. *Quantum* **3**, 191 (2019).
- Li, Y. & Benjamin, S. C. Efficient variational quantum simulator incorporating active error minimization. *Phys. Rev. X* **7**, 021050 (2017).
- Kokail, C. et al. Self-verifying variational quantum simulation of lattice models. *Nature* **569**, 355 (2019).
- Khatris, S. et al. Quantum-assisted quantum compiling. *Quantum* **3**, 140 (2019).
- Jones, T. & Benjamin, S. C. Robust quantum compilation and circuit optimisation via energy minimisation. *Quantum* **6**, 628 (2022).
- Heya, K., Suzuki, Y., Nakamura, Y. & Fujii, K. Variational quantum gate optimization. Preprint at <https://arxiv.org/abs/1810.12745> (2018).
- Endo, S., Sun, J., Li, Y., Benjamin, S. C. & Yuan, X. Variational quantum simulation of general processes. *Phys. Rev. Lett.* **125**, 010501 (2020).
- Sharma, K., Khatris, S., Cerezo, M. & Coles, P. Noise resilience of variational quantum compiling. *New J. Phys.* **22**, 043006 (2020).
- Carolan, J. et al. Variational quantum unsampling on a quantum photonic processor. *Nat. Phys.* **16**, 322–327 (2020).
- Yoshioka, N., Nakagawa, Y. O., Mitarai, K. & Fujii, K. Variational quantum algorithm for non-equilibrium steady states. *Phys. Rev. Res.* **2**, 043289 (2020).
- Bravo-Prieto, C. et al. Variational quantum linear solver: A hybrid algorithm for linear systems. Preprint at <https://arxiv.org/abs/1909.05820> (2019).
- Xu, X. et al. Variational algorithms for linear algebra. *Sci. Bull.* **66**, 2181–2188 (2021).
- McArdle, S. et al. Variational ansatz-based quantum simulation of imaginary time evolution. *Npj Quantum Inf.* **5**, 1–6 (2019).
- Cirstoiu, C. et al. Variational fast forwarding for quantum simulation beyond the coherence time. *Npj Quantum Inf.* **6**, 1–10 (2020).
- Otten, M., Cortes, C. L. & Gray, S. K. Noise-resilient quantum dynamics using symmetry-preserving ansatzes. Preprint at <https://arxiv.org/abs/1910.06284> (2019).
- Lubasch, M., Joo, J., Moinier, P., Kiffner, M. & Jaksch, D. Variational quantum algorithms for nonlinear problems. *Phys. Rev. A* **101**, 010301 (2020).
- Verdon, G., Marks, J., Nanda, S., Leichenauer, S. & Hidary, J. Quantum hamiltonian-based models and the variational quantum thermalizer algorithm. Preprint at <https://arxiv.org/abs/1910.02071> (2019).
- Bravo-Prieto, C., Garcia-Martin, D. & Latorre, José. Quantum singular value decomposer. *Phys. Rev. A* **101**, 062310 (2020).
- Cao, Y. et al. Quantum chemistry in the age of quantum computing. *Chem. Rev.* **119**, 10856–10915 (2019).
- McArdle, S., Endo, S., Aspuru-Guzik, A., Benjamin, S. C. & Yuan, X. Quantum computational chemistry. *Rev. Mod. Phys.* **92**, 015003 (2020).
- Dumitrescu, E. F. et al. Cloud quantum computing of an atomic nucleus. *Phys. Rev. Lett.* **120**, 210501 (2018).
- Wecker, D., Hastings, M. B. & Troyer, M. Progress towards practical quantum variational algorithms. *Phys. Rev. A* **92**, 042303 (2015).
- Bauer, B., Wecker, D., Millis, A. J., Hastings, M. B. & Troyer, M. Hybrid quantum-classical approach to correlated materials. *Phys. Rev. X* **6**, 031045 (2016).

33. Lloyd, S., Mohseni, M. & Rebentrost, P. Quantum principal component analysis. *Nat. Phys.* **10**, 631 (2014).
34. Horn, R. A. & Johnson, C. R. *Matrix Analysis* (Cambridge University Press, 1990).
35. Bhatia, R. *Matrix Analysis* Vol. 169 (Springer Science & Business Media, 2013).
36. Kandala, A. et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242 (2017).
37. McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R. & Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **9**, 4812 (2018).
38. Cerezo, M., Sone, A., Volkoff, T., Cincio, L. & Coles, P. J. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nat. Commun.* **12**, 1–12 (2021).
39. Sharma, K., Cerezo, M., Cincio, L. & Coles, P. J. Trainability of dissipative perceptron-based quantum neural networks. *Phys. Rev. Lett.* **128**, 180505 (2022).
40. Garcia-Saez, A. & Latorre, J. I. Addressing hard classical problems with adiabatically assisted variational quantum eigensolvers. Preprint at <https://arxiv.org/abs/1806.02287> (2018).
41. Cincio, L., Subaşı, Y., Sornborger, A. T. & Coles, P. J. Learning the quantum algorithm for state overlap. *New J. Phys.* **20**, 113022 (2018).
42. Verdon, G. et al. Learning to learn with quantum neural networks via classical neural networks. Preprint at <https://arxiv.org/abs/1907.05415> (2019).
43. Grant, E., Wossnig, L., Ostaszewski, M. & Benedetti, M. An initialization strategy for addressing barren plateaus in parametrized quantum circuits. *Quantum* **3**, 214 (2019).
44. Volkoff, T. & Coles, P. J. Large gradients via correlation in random parameterized quantum circuits. *Quantum Sci. Technol.* **6**, 025008 (2021).
45. Hadfield, S. et al. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms* **12**, 34 (2019).
46. Nakanishi, K. M., Fujii, K. & Todo, S. Sequential minimal optimization for quantum-classical hybrid algorithms. *Phys. Rev. Res.* **2**, 043158 (2020).
47. Parrish, R. M., Josue, J. T., Ozaeta, A. & McMahon, P. L. A Jacobi diagonalization and Anderson acceleration algorithm for variational quantum algorithm parameter optimization. Preprint at <https://arxiv.org/abs/1904.03206> (2019).
48. Harrow, A. W. & Napp, J. C. Low-depth gradient measurements can improve convergence in variational hybrid quantum-classical algorithms. *Phys. Rev. Lett.* **126**, 140502 (2021).
49. Kübler, J. M., Arrasmith, A., Cincio, L. & Coles, P. J. An adaptive optimizer for measurement-frugal variational algorithms. *Quantum* **4**, 263 (2020).
50. Sweke, R. et al. Stochastic gradient descent for hybrid quantum-classical optimization. *Quantum* **4**, 314 (2020).
51. Arrasmith, A., Cincio, L., Somma, R. D. & Coles, P. J. Operator sampling for shot-frugal optimization in variational algorithms. Preprint at <https://arxiv.org/abs/2004.06252> (2020).
52. Mitarai, K., Negoro, M., Kitagawa, M. & Fujii, K. Quantum circuit learning. *Phys. Rev. A* **98**, 032309 (2018).
53. Schuld, M., Bergholm, V., Gogolin, C., Izaac, J. & Killoran, N. Evaluating analytic gradients on quantum hardware. *Phys. Rev. A* **99**, 032331 (2019).
54. Cruz, D. et al. Efficient quantum algorithms for GHZ and W states, and implementation on the ibm quantum computer. *Adv. Quantum Technol.* **2**, 1900015 (2019).
55. Bärttschi, A. & Eidenbenz, S. *International Symposium on Fundamentals of Computation Theory* 126–139 (Springer, 2019).
56. Abdo, B. et al. IBM Q 16 Melbourne backend specification, <https://github.com/Qiskit/ibmq-device-information/tree/master/backends/melbourne/V1> (2018).
57. Li, H. & Haldane, F. D. M. Entanglement spectrum as a generalization of entanglement entropy: Identification of topological order in non-abelian fractional quantum hall effect states. *Phys. Rev. Lett.* **101**, 010504 (2008).
58. Subaşı, Y., Cincio, L. & Coles, P. J. Entanglement spectroscopy with a depth-two quantum circuit. *J. Phys. A Math. Theor.* **52**, 044001 (2019).
59. Giampaolo, S. M., Montangero, S., Dell'Anno, F., De Siena, S. & Illuminati, F. Universal aspects in the behavior of the entanglement spectrum in one dimension: Scaling transition at the factorization point and ordered entangled structures. *Phys. Rev. B* **88**, 125142 (2013).
60. Cerezo, M., Rossignoli, R. & Canosa, N. Nontransverse factorizing fields and entanglement in finite spin systems. *Phys. Rev. B* **92**, 224422 (2015).
61. Cerezo, M., Rossignoli, R. & Canosa, N. Factorization in spin systems under general fields and separable ground-state engineering. *Phys. Rev. A* **94**, 042335 (2016).
62. Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge University Press, 2010).
63. Tang, E. A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* 217–228 (2019).
64. Arrazola, J. M., Delgado, A., Bardhan, B. R. & Lloyd, S. Quantum-inspired algorithms in practice. *Quantum* **4**, 307 (2020).
65. Cotler, J., Huang, H.-Y. & McClean, J. R. Revisiting dequantization and quantum advantage in learning tasks. Preprint at <https://arxiv.org/abs/2112.00811> (2021).
66. Huang, H.-Y. et al. Quantum advantage in learning from experiments. *Science* **376**, 1182–1186 (2022).
67. Aaronson, S. Read the fine print. *Nat. Phys.* **11**, 291–293 (2015).
68. Koczor, B. Exponential error suppression for near-term quantum devices. *Phys. Rev. X* **11**, 031057 (2021).
69. Huggins, W. J. et al. Virtual distillation for quantum error mitigation. *Phys. Rev. X* **11**, 041036 (2021).

ACKNOWLEDGEMENTS

We thank Lukasz Cincio for helpful conversations. All authors acknowledge support from LANL's Laboratory Directed Research and Development (LDRD) program. M.C. was also supported by the Center for Nonlinear Studies at LANL. P.J.C. also acknowledges support from the LANL ASC Beyond Moore's Law Project. This work was also supported by the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research, under the Quantum Computing Application Teams program, and by the U.S. DOE, Office of Science, Basic Energy Sciences, Materials Sciences and Engineering Division, Condensed Matter Theory Program.

AUTHOR CONTRIBUTIONS

The project was conceived by P.J.C. The manuscript was written by M.C., K.S., A.A., and P.J.C. The theoretical results were derived by M.C., K.S., and P.J.C. M.C. and K.S. performed the numerical simulations in Figs. 4, 5, and 7. A.A. developed the noisy simulator and performed the numerical simulations for error mitigation of Fig. 6. M.C. and K.S. contributed equally to this work and are considered as co-first authors.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Correspondence and requests for materials should be addressed to M. Cerezo.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022