

PAPER • OPEN ACCESS

Preparing ground and excited states using adiabatic CoVaR

To cite this article: Wooseop Hwang and Bálint Koczor 2025 *New J. Phys.* **27** 023025

View the [article online](#) for updates and enhancements.

You may also like

- [Exploring *ab initio* machine synthesis of quantum circuits](#)
Richard Meister, Cica Gustiani and Simon C Benjamin
- [Application of Machine Learning and CoVaR Model on Intelligent Decision Method](#)
Zheyu Yang and Anna Dai
- [Risk spillovers between China and other BRICS countries during COVID-19 pandemic: A CoVaR-copula approach](#)
Yangnan Cheng, Jianxu Liu and Songsak Sriboonchitta



PAPER

Preparing ground and excited states using adiabatic CoVaR

OPEN ACCESS

RECEIVED
5 November 2024REVISED
25 January 2025ACCEPTED FOR PUBLICATION
3 February 2025PUBLISHED
17 February 2025

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.

Wooseop Hwang*  and Bálint Koczor 

Mathematical Institute, University of Oxford, Woodstock Road, Oxford OX2 6GG, United Kingdom

* Author to whom any correspondence should be addressed.

E-mail: useop02@gmail.com and koczor@maths.ox.ac.uk**Keywords:** quantum simulation, quantum information, variational algorithms, quantum algorithms for chemical simulations

Abstract

CoVariance Root finding with classical shadows (CoVaR) was recently introduced as a new paradigm for training variational quantum circuits. Common approaches, such as variants of the variational quantum eigensolver, aim to optimise a non-linear classical cost function and thus suffer from, e.g. poor local minima, high shot requirements and barren plateaus. In contrast, CoVaR fully exploits powerful classical shadows and finds joint roots of a very large number of covariances using only a logarithmic number of shots and linearly scaling classical computing resources. As a result, CoVaR has been demonstrated to be particularly robust against local traps, however, its main limitation has been that it requires a sufficiently good initial state. We address this limitation by introducing an adiabatic morphing of the target Hamiltonian and demonstrate in a broad range of application examples that CoVaR can successfully prepare eigenstates of the target Hamiltonian when no initial warm start is known. CoVaR succeeds even when Hamiltonian energy gaps are very small—this is in stark contrast to adiabatic evolution and phase estimation algorithms where circuit depths scale inversely with the Hamiltonian energy gaps. On the other hand, when the energy gaps are relatively small then adiabatic CoVaR may converge to higher excited states as opposed to a targeted specific low-lying state. Nevertheless, we exploit this feature of adiabatic CoVaR and demonstrate that it can be used to map out the low lying spectrum of a Hamiltonian which can be useful in practical applications, such as estimating thermal properties or in high-energy physics.

1. Introduction

One of the most natural applications of quantum computers is the simulation of quantum systems which, for certain tasks, promises exponential speedups over classical computers. However, it is anticipated that early generations of quantum computers (NISQ and early fault tolerant) will be limited in the numbers of qubits and in the circuit depths that can be implemented [1]. This motivated the field to consider hybrid algorithms whereby a quantum computer is used to run shallow quantum circuits whose outputs are post processed in classical computers. Such variational quantum algorithms [2–4] have received significant attention and have been extended to a broad range of application areas such as simulating time evolution [5] or finding ground and excited states [6–10].

On the other hand, variational quantum algorithms suffer from significant performance limitations. First, estimating the cost function and its gradient to high accuracy may require a prohibitively large number of shots [3, 4, 11]. Second, the optimisation landscape is particularly complex and training is generally NP-hard due to the presence of a large number of local optima [12, 13]. Furthermore, if no good initial parameters are known then the optimisation may require exponential efforts to estimate gradients due to barren plateaus [14].

Recently, [15] introduced the CoVariance Root finding (CoVaR) approach, which fundamentally redefines the problem of training circuit parameters by posing the problem as finding joint roots of a large number of covariance functions rather than the minimisation of a classical cost function. The approach is powerful as it can fully exploit classical shadows [16] to estimate a large number of covariance functions

using only a logarithmic number of shots and using linearly scaling classical compute resources—while the use of high performance computers (HPCs) offers further improvements by enabling very large operator pools. The approach has the significant benefit that it can avoid local traps and is robust against imperfections, such as gate noise [15].

However, the main limitation of the CoVaR approach is its sensitivity to good initial parameters, i.e. [15] numerically observed that CoVaR—similarly to phase estimation algorithms—randomly converges to one of the nearby eigenstates with a probability approximately proportional to the overlap in the initial state. While CoVaR has been a promising approach for training circuits when good initial parameters are available, in the present work we significantly extend its capabilities to the practically more important case when no good initial states are known. As such, we demonstrate in a broad range of application examples that CoVaR can successfully prepare eigenstates of systems by initialising in an exact eigenstate of an initially solvable Hamiltonian which is then gradually morphed into the desired problem Hamiltonian. The only downside, however, is that in general one does not have full control as to which eigenstate the approach converges to—albeit we can guarantee convergence to specific eigenstates when the ansatz is sufficiently deep and the relevant energy gaps in the Hamiltonian are sufficiently large.

The present approach, called adiabatic CoVaR, uses a time-dependent Hamiltonian $\mathcal{H}(t)$ and builds on the concept of adiabatic computing, such that one initialises in an eigenstate of a trivial Hamiltonian at $\mathcal{H}(t=0)$ and then gradually interpolates in steps of Δt until the final problem Hamiltonian is approached at $\mathcal{H}(t=1)$. At each step CoVaR is used to prepare an instantaneous eigenstate of the formally time-dependent Hamiltonian and given a sufficiently small Δt , the initial state supplied to CoVaR is always guaranteed to have a high overlap with a targeted instantaneous eigenstate.

We numerically demonstrate the effectiveness of this approach in a broad range of application examples, including finding ground and excited states of spin problems, finding solutions to combinatorial optimisation problems and solving Hamiltonian eigenstates in a high-energy physics example. While variational quantum eigensolver (VQE) could, in principle, be used as part of this approach to find instantaneous eigenstates, in all numerical simulations of VQE in the present work got trapped in local optima. In contrast, in all numerical demonstrations in the present work CoVaR successfully converged to an eigenstate of the problem Hamiltonian providing strong evidence for its robustness against local traps. However, the main limitation of the approach is that small energy gaps and shallow ansatz circuit depths force CoVaR to converge to higher excited states even when it was initialised in a low-lying state of the initial model. Nevertheless, this may be a valuable feature of the present approach as adiabatic CoVaR allows to map out the low lying spectrum of a Hamiltonian—as we demonstrate this can benefit an example application in high-energy physics.

Crucially, while the complexity of conventional adiabatic evolution and phase estimation protocols depend inversely on Hamiltonian energy gaps, we numerically observe that adiabatic CoVaR has a fundamentally improved robustness against small energy gaps and empirically observe that the step size Δt has a logarithmic dependence on the smallest Hamiltonian energy gap within the finite regime investigated. Furthermore, the parameter Δt does not affect the precision of the final eigenstate preparation, nor the circuit depths required, and we demonstrate that CoVaR can prepare eigenstates to relatively high precision even when using relatively large Δt —this is in stark contrast to phase estimation and adiabatic evolution algorithms whereby the circuit depth needs to scale inversely with the energy gaps otherwise the approaches cannot output an eigenstate. Of course, as we demonstrate, adiabatic CoVaR potentially outputs higher excited states when Δt is too large.

This manuscript is organised as follows. In the rest of this introduction we briefly review basics of preparing eigenstates using variational circuits and using adiabatic evolution. Then in section 2 we present details of our adiabatic CoVaR approach and then demonstrate its effectiveness in a range of practically important application examples in section 3. Finally, we demonstrate the robustness of the present approach against imperfections and against different hyper parameter choices and then conclude.

1.1. Preliminaries: training variational circuits

Variational quantum circuits are constructed as a sequence of L gate operations that are often chosen to be native, parametrised gate operations of the hardware platform [4] as $U(\boldsymbol{\theta}) = U_L(\theta_L) \dots U_2(\theta_2)U_1(\theta_1)$. These circuits are then applied to an initial state, such as $|0 \dots 0\rangle$, to prepare a parametrised family of variational states as

$$|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})|0 \dots 0\rangle. \quad (1)$$

The particular construction of an ‘ansatz’ circuit $U(\boldsymbol{\theta})$ is often motivated by physical considerations, e.g. in the case of the Hamiltonian Variational Ansatz the individual gates in the ansatz circuit correspond to Pauli terms in the problem Hamiltonian [4].

A common task in quantum simulation is to find the ground state of a problem Hamiltonian $\mathcal{H} = \sum_a w_a P_a$ that is specified as a linear combination of Pauli strings $P_a \in \{\mathbb{1}, X, Y, Z\}^{\otimes N}$. In VQEs, the expected value of this Hamiltonian $C(\boldsymbol{\theta}) = \langle \psi(\boldsymbol{\theta}) | \mathcal{H} | \psi(\boldsymbol{\theta}) \rangle$ is minimised as a cost function according to the variational principle such that the global minimum over the parameters $\boldsymbol{\theta}$ approximates the ground state [2, 9, 11].

While the approach may be effective at moderate system sizes, as one scales up, the classical optimization procedure may require exponential training efforts as in general the VQE optimization problem is NP-hard [12]. In particular, if no good initial guess of the optimal parameters is known then training may become prohibitively expensive due to barren plateaus [14, 17]. However, even if good initial states are known, it has been shown that the optimization landscape of VQE contains exponentially many local traps [13]. This motivates alternative techniques for training variational circuits, and as we demonstrate in the present work, adiabatic CoVaR mitigates both issues: the parameters are guided by an adiabatic evolution and thus the system is always ‘warm started’ while getting trapped in local optima is mitigated through the use of CoVaR.

CoVaR Finding was introduced in [15]; The approach finds eigenstates through a root finding problem over a very large number of covariance surfaces rather than through the minimisation of a single energy surface as in VQE. In particular, given any two Hermitian operators A and B one can define the covariance between them that depends on a specific input quantum state $|\psi\rangle$ (which is assumed to be pure for ease of notation [15]) as

$$\langle A, B \rangle_\psi := \langle \psi | AB | \psi \rangle - \langle \psi | A | \psi \rangle \langle \psi | B | \psi \rangle. \quad (2)$$

We choose these operators from a predefined operator pool $\mathcal{P} := \{O_k\}_{k=1}^p$ which in the present work we assume to be local Pauli strings $O_k \in \{\mathbb{1}, X, Y, Z\}^{\otimes N}$, but any other orthonormal operator basis can also be used [15].

The approach builds on the main observation that given our problem Hamiltonian $\mathcal{H} = \sum_a h_a \mathcal{H}_a$, such that we denote the Hamiltonian decomposition operators as $\mathcal{H}_a \in \mathcal{Q}$ with $\mathcal{Q} \subseteq \mathcal{P}$, then a quantum state $|\psi\rangle$ is an eigenstate of \mathcal{H} if the following conditions are met as

$$\begin{aligned} \text{Sufficient condition: } & \langle \mathcal{H}_a, \mathcal{H} \rangle_\psi = 0, \quad \forall \mathcal{H}_a \in \mathcal{Q}, \\ \text{Necessary condition: } & \langle O_k, \mathcal{H} \rangle_\psi = 0, \quad O_k \in \mathcal{P}. \end{aligned}$$

The sufficient condition already ensures that the eigenvalue equation $\mathcal{H}|\psi\rangle = \langle \mathcal{H} \rangle |\psi\rangle$ holds, however, further necessary conditions are introduced by any other operator O_k . Then, given a parametrised variational quantum state $|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})|0\rangle^{\otimes N}$, where U is a variational quantum circuit (an example is illustrated in figure 6), the parameterised covariance is defined as [15]:

$$f_k(\boldsymbol{\theta}) := \langle O_k, \mathcal{H} \rangle_{\psi(\boldsymbol{\theta})}. \quad (3)$$

Indeed, similarly to VQE cost functions [18], the covariance function $f_k(\boldsymbol{\theta})$ is an infinitely differentiable function of the circuit parameters $\boldsymbol{\theta}$ for any Hermitian operator O_k and Hamiltonian \mathcal{H} .

This converts the problem of finding eigenstates into the problem of finding roots of covariances over the parameters, i.e, one searches for parameter values $\boldsymbol{\theta}$ such that covariances in the state $|\psi(\boldsymbol{\theta})\rangle$ are simultaneously all zero for all operators in the operator pool. There exist numerous methods for finding simultaneous roots of such vector-valued functions as variants of Newton’s method [19]; One iterates the parameter values according to the update rule as

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \mathbf{J}^{-1} \mathbf{f}. \quad (4)$$

Here both the vector of covariances \mathbf{f} and the Jacobian \mathbf{J} , which is a matrix of all partial derivatives of \mathbf{f} , are estimated using a quantum computer and they are defined as

$$\mathbf{f}_k(\boldsymbol{\theta}) = \langle O_k, \mathcal{H} \rangle_{\psi(\boldsymbol{\theta})}, \quad [\mathbf{J}]_{kl} = \partial_k f_l(\boldsymbol{\theta}). \quad (5)$$

For ease of notation, we focus on operator pools of local Pauli strings and corresponding efficient Pauli shadows, however, as we detail below a range of more advanced classical shadow techniques are immediately applicable.

The significant advantage of CoVaR is that we can use a very large operator pool as we can simultaneously estimate a large number of covariances using classical shadows [16]. Reference [15] analysed in detail how an increasing number of constraints improves the convergence rate of the CoVaR approach, and therefore, one would ideally use as many constraints as possible – but certainly more than the absolute

minimum as the number of circuit parameters. However, in the specific case of Pauli shadows, [16] proved that the shot noise cost of classical shadows scales exponentially as $O(3^w)$ with the weight of the Pauli string. This, fortunately, poses no strict limitation in practice, e.g. [15] demonstrated that even 2-local Pauli strings with $w = 2$ can provide rich information, while [20] used a similar approach in a 100-qubit system reconstructing millions of Pauli strings from classical shadows.

Classical shadows based on random Pauli basis measurements involve the following steps. First, we prepare the variational state using our ansatz circuit $|\psi(\boldsymbol{\theta})\rangle$ and then measure each qubit randomly in either the X , Y , or Z bases (by applying random single-qubit rotations) which results in an N -bit classical measurement outcome as the bitstring $|b_i\rangle \in \{0, 1\}^N$. The collection of N_s measurement bases and N_s outcome bitstrings forms a classical dataset which is the classical shadow of the quantum state $|\psi(\boldsymbol{\theta})\rangle$ [16]. All covariances can then be estimated efficiently by post-processing this classical shadow dataset to a precision ϵ that scales with the square root as $O(N_s^{-1/2})$. Specifically, the required Pauli operator expected values can be estimated from classical shadows to a guaranteed shot noise precision ϵ using the median of means estimator detailed in [16] and a number of shots that scales with the logarithm of N_c . Furthermore, [15] detailed that N_c covariances can be reconstructed from classical shadows, through processing Pauli operator expected values, using only linear $O(N_c)$ classical post processing resources. As we will detail below, this yields the guaranteed shot noise cost of a single iteration of CoVaR in equation (17).

A remarkable and unique feature of CoVaR is that it is robust against local traps in the energy surface. Specifically, ref [15] demonstrated that VQE can get trapped in local optima, however, when initialising CoVaR in those traps, the approach can ‘jump out’ and converge to the desired eigenstate.

In fact, CoVaR can be compared with phase estimation algorithms, as the success rate of converging to the desired eigenstate was found proportional to the initial state overlap in numerical simulations. On the other hand, the main limitation of CoVaR was found that it fails to make progress when initialised away from an eigenstate, i.e. when no good initial parameters are known. In the present work we demonstrate adiabatic CoVaR in which the CoVaR optimiser is used iteratively in a way to guarantee that at each step the inputted initial state remains close to an eigenstate of the system.

1.2. Adiabatic evolution

In adiabatic quantum computation one prepares the ground state of a trivial Hamiltonian and slowly morphs it into the true problem Hamiltonian such that the system remains in the instantaneous ground state of the time-dependent Hamiltonian [21, 22].

For example, one assumes a Hamiltonian H_0 is available whose ground state is known and the aim is to find the ground state of the Hamiltonian \mathcal{H} . The adiabatic approach is then typically applied to the linear combination as

$$\mathcal{H}(s) = (1 - s)\mathcal{H}_0 + s\mathcal{H}. \quad (6)$$

In general, given a one-parameter family of Hamiltonians $\mathcal{H}(s)$ for $0 \leq s \leq 1$ we define its instantaneous ground-state via the eigenvalue equation as

$$\mathcal{H}(s)|\psi_0(s)\rangle = E_0(s)|\psi_0(s)\rangle. \quad (7)$$

Then, choosing the initial state $|\psi_0(0)\rangle$ as the ground state of the initial trivial Hamiltonian $H(0)$, the adiabatic theorem guarantees that for a sufficiently long time evolution, the final state will be arbitrarily close to $|\psi_0(s=1)\rangle$ as long as the system is gapped as $E_1(s) - E_0(s) > 0$ at all times $0 \leq s \leq 1$.

Setting $s = t/T$, where t is a physical time variable, the total evolution time T needs to be chosen according to the condition as

$$T \gg \frac{\epsilon}{g_{\min}^2}, \quad (8)$$

which guarantees that the time-evolved state stays close to the instantaneous ground state $|\psi_0(s)\rangle$ throughout the evolution in a Landau–Zener problem setting [23]. Here the minimum energy gap is defined as

$$g_{\min} = \min_{0 \leq s \leq 1} [E_1(s) - E_0(s)], \quad (9)$$

while ϵ depends on the transition amplitudes as

$$\epsilon = \max_{0 \leq s \leq 1} |\langle \psi_1(s) | \frac{d\mathcal{H}}{ds} | \psi_0(s) \rangle|. \quad (10)$$

The time evolution under the parametrised Hamiltonian can be simulated using a range of techniques, including trotterisation, variational quantum simulation [24], quantum walk-based methods [25], Taylor series expansions [25], and qubitization [26]. While adiabatic computing is indeed a general and powerful approach, its implementation may be challenging due to deep quantum circuits required for simulating sufficiently long time evolutions. While our approach is indeed inspired by adiabatic evolution, we do not use time evolution but rather use CoVaR to train a parametrised short-depth circuit to follow instantaneous eigenstates of the problem Hamiltonian.

2. Adiabatic CoVaR

We now detail our approach which combines an adiabatic evolution, i.e. gradual morphing of a Hamiltonian, and CoVaR which is used to train a shallow-depth parametrised circuit in order to closely follow instantaneous eigenstates of the adiabatic Hamiltonian. The significant advantage of this approach is that no time evolution is used but rather parameters of a shallow ansatz circuit are slowly varied using CoVaR. While CoVaR has been observed in numerical simulations to rapidly and robustly converge to eigenstates, its main limitation has been that it needs an initial state that has sufficient overlap with the target eigenstate. The present adiabatic morphing guarantees that CoVaR is always supplied an initial state that has high overlap with an instantaneous eigenstate.

2.1. Hamiltonian morphing

In the present work we explore two different Hamiltonian morphing techniques both of which interpolate between an initial and the final Hamiltonians by gradually increasing a time parameter $t \in \{0, \Delta t, 2\Delta t, \dots, 1\}$ in discrete time steps of Δt that we will refer to as the adiabatic time-step.

First, we consider the usual adiabatic evolution Hamiltonian from equation (6) by setting $T \equiv 1$ as

$$\mathcal{H}(t) = (1-t)\mathcal{H}_0 + t\mathcal{H}, \quad (11)$$

where \mathcal{H}_0 is an initial Hamiltonian whose eigenstates can be solved analytically and can be prepared with a shallow circuit, and \mathcal{H} is the problem Hamiltonian of interest. We will refer to this as the ‘mixing approach’.

Second, we consider a special case that can be applied when the Hamiltonian decomposes into a—potentially dominant—contribution \mathcal{H}_0 whose eigenstates can be solved analytically and can be prepared with shallow circuits, and into additional terms as $\mathcal{H} = \mathcal{H}_0 + \mathcal{H}_1$ [27]. We then consider the adiabatic Hamiltonian as

$$\mathcal{H}(t) = \mathcal{H}_0 + t\mathcal{H}_1. \quad (12)$$

This variant is particularly relevant to physical systems where non-trivial perturbative interactions \mathcal{H}_1 are added to an otherwise solvable system \mathcal{H}_0 , and thus we will refer to this as the ‘perturbative approach’.

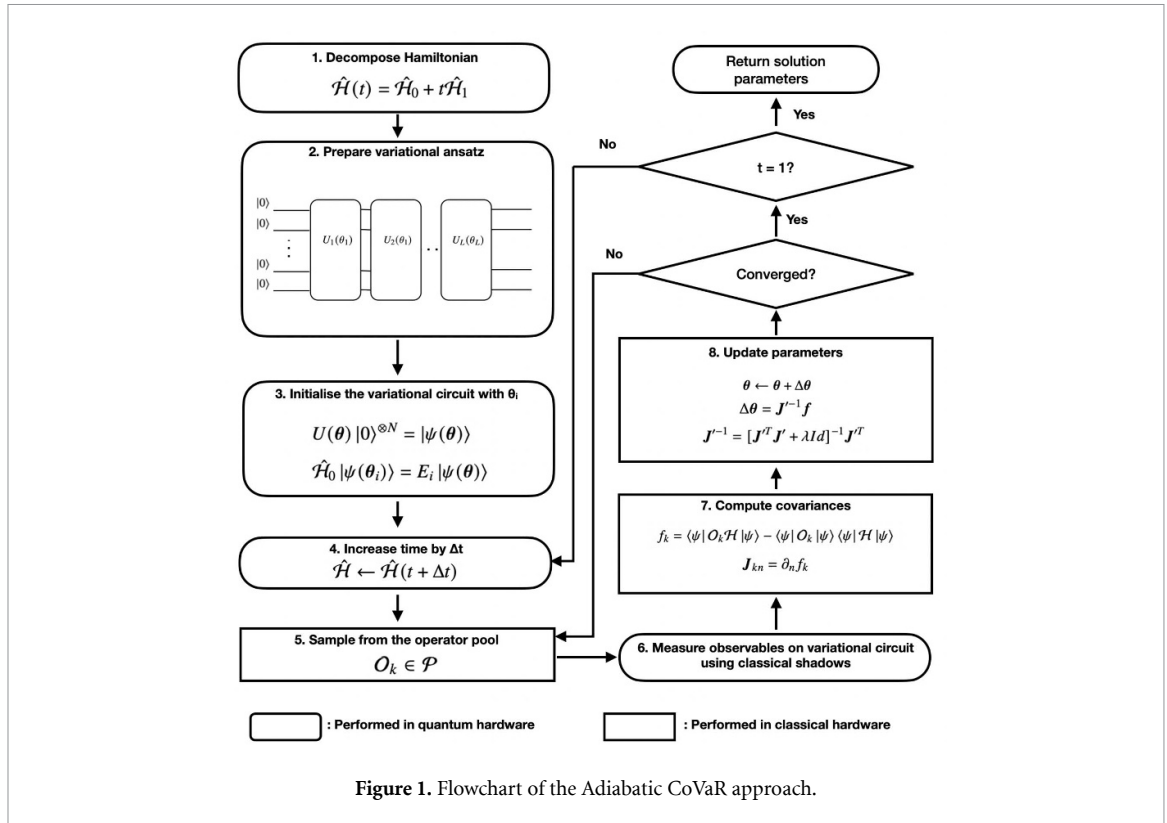
For example, we will later consider the spin ring Hamiltonian in equation (18) for which we can define $\mathcal{H}_0 = \sum_{i=1}^N c_i Z_i$ and $\mathcal{H}_1 = J \sum_{i=1}^N \sigma_i \cdot \sigma_{i+1}$, where J is a coupling constant, σ_i are vectors of single-qubit Pauli matrices acting on the i th qubit, and $-1 \leq c_i \leq 1$ are uniformly randomly generated onsite interactions. This naturally leads to the parametrisation of the Hamiltonian as

$$\mathcal{H}(t) = \sum_{i=1}^N c_i Z_i + tJ \sum_{i=1}^N \sigma_i \cdot \sigma_{i+1}. \quad (13)$$

The above decomposition is convenient as one can not only diagonalize the initial trivial model \mathcal{H}_0 analytically but also prepare the eigenstates as computational basis states using simple bitflip circuits. However, for systems in which such separation is not possible, the mixing approach should be employed. In cases when J is small, the perturbative approach equation (12) may be preferable over the mixing approach equation (11) as the initial state is merely perturbed by the additional non-trivial terms in \mathcal{H}_1 , whereas the mixing approach transforms between Hamiltonians of entirely different physical systems. An added benefit of the perturbative approach is that it prepares eigenstates for the entire path $0 \leq J \leq 1$ not just for the final, target Hamiltonian with, e.g. $J = 1$. We also note that the efficacy of both the mixing and perturbative approaches can be improved by including counterdiabatic terms [28].

2.2. Adiabatic CoVaR

Our approach first initialises the variational circuit in an eigenstate of the initial Hamiltonian $\mathcal{H}(t=0)$. The parameter value t is then increased to Δt and CoVaR is used to train the circuit parameters such that the



variational circuit prepares an eigenstate of the the Hamiltonian $\mathcal{H}(\Delta t)$. This is then repeated iteratively until $t = 1$ is reached.

More specifically, we use the variational ansatz $|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})|0\rangle^{\otimes N}$, in which the parameters $\boldsymbol{\theta} := \boldsymbol{\theta}(t, \kappa)$ depend on t and κ . As illustrated in figure 1, t is increased in steps $t \in \{0, \Delta t, 2\Delta t, \dots, 1\}$ and for each fixed value of t , a CoVaR optimisation is run whose iteration count is denoted by κ and is capped at a maximum allowed CoVaR steps K .

The initial parameters $\boldsymbol{\theta}(t=0, \kappa=0)$ by construction prepare an energy eigenstate of the initial Hamiltonian \mathcal{H}_0 as

$$\mathcal{H}_0|\psi(\boldsymbol{\theta}(0,0))\rangle = E_i|\psi(\boldsymbol{\theta}(0,0))\rangle. \quad (14)$$

The parameters are then updated through the following update rules as

$$\boldsymbol{\theta}(t + \Delta t, 0) = \boldsymbol{\theta}(t, K) \quad (15)$$

$$\boldsymbol{\theta}(t, \kappa + 1) = \boldsymbol{\theta}(t, \kappa) - \mathbf{J}^{-1}\mathbf{f}(\boldsymbol{\theta}_i(t, \kappa)). \quad (16)$$

Here, $\mathbf{J} := \mathbf{J}(\boldsymbol{\theta}(t, \kappa))$ is the Jacobian and $\mathbf{f} := \mathbf{f}(\boldsymbol{\theta}(t, \kappa))$ is the covariance vector as defined in equation (5).

In contrast to algorithms that only find energies of the eigenstates, a significant advantage of the present approach is that it allows the direct preparation of both the ground state and excited states using a short-depth variational ansatz—albeit one does not have full control over which eigenstate CoVaR converges to as we demonstrate below. This is in contrast to VQD [6, 29], which we summarise in appendix A.1 and which requires the computation of the below $n - 1$ energy eigenstates for the preparation of n th excited state.

The other significant advantage of adiabatic CoVaR is that a relatively large Δt can be used compared to conventional adiabatic evolution. As detailed in equation (8), the time complexity of adiabatic evolution is $O(g_{\min}^{-2})$, where g_{\min} is the minimum energy difference between the consecutive eigenstates throughout the adiabatic evolution [30]. We will later demonstrate that indeed Δt scales favourably with the energy gap, and that Δt does not significantly affect the precision of eigenstate preparation. The present approach performs in total $K/(\Delta t)$ iterations of CoVaR. Boyd and Koczor [15] bounded the number of samples N_s , i.e. number of circuit repetitions, required for a single CoVaR iteration: using classical shadows for estimating the Jacobian $\mathbf{J} \in \mathbb{C}^{N_c \times \nu}$ to an error ϵ requires only a logarithmic overhead in the number N_c of covariances as

$$N_s = \mathcal{O}[\nu \log(N_c) \epsilon^{-2}], \quad (17)$$

which allows us to estimate a very large number of covariances. The cost also depends proportionally on the number ν of circuit parameters.

One of the other advantages we inherit from CoVaR is its robustness to gate noise, shot noise and robustness to getting trapped in local optima as we are adopting a stochastic Levenberg–Marquadt evolution [31, 32], whereby the operator pool \mathcal{P} , and thus the covariances are randomly chosen at each iteration as in [15].

As we demonstrate later, when the energy gap g_{\min} is too small, CoVaR may converge to a different eigenstate located close in parameter space. A possible approach to allow CoVaR to get back to the desired eigenstate could be by adding small random fluctuations in the parameters $\theta_i(t, 0)$ at every timestep Δt —or by occasionally performing VQE energy minimisation throughout the evolution to force CoVaR to converge to low-lying eigenstates.

3. Application examples

In this section, we present numerical simulations in a range of practically relevant application examples as (a) finding ground and excited states of a spin ring Hamiltonian, (b) solving binary optimisation problems by finding the ground state of a Hamiltonian that encodes the solution to a max-cut problem and (c) finding excited states of the Schwinger model which is a toy model in high-energy physics.

3.1. Eigenstate discovery of spin problems

Spin models are important in a broad range of applications including condensed matter physics and material science [3, 15, 33–35]. Furthermore, certain hard optimisation problems can be mapped to spin models [36] and indeed certain spin models may be hard to simulate classically for large numbers of qubits [37, 38]. Conversely, finding ground states of certain spin problems for large N may require exponentially increasing circuit depths [39].

Here we consider the perturbative approach introduced in section 2.1 to find eigenstates of the random field Heisenberg model as

$$\mathcal{H} = \sum_{i=1}^N [c_i Z_i + J \sigma_i \cdot \sigma_{i+1}]. \quad (18)$$

We fix the coupling strength $J=1.0$ to match the size of the onsite interactions in which case the eigenstates are most complex and we choose a system size of $N=10$ qubits. In the present perturbative morphing approach we initialise the variational circuit to prepare the ground state of the initial Hamiltonian $\mathcal{H}(0)$ as a computational basis state and then gradually increase the coupling strength J of the model as in equation (13).

Figure 2(left, blue) shows the energy obtained with adiabatic CoVaR throughout the evolution and confirms that, indeed, the present approach can closely follow the instantaneous ground states as the blue curves rapidly converge to the instantaneous ground state energies of the problem Hamiltonian (black dashed lines). In fact, each time the blue line in figure 2(left) converges to one of the dashed lines represents the successful preparation of the ground state of the Hamiltonian at intermediate values of J —this nicely illustrates that the present approach prepares an entire family of ground states for increasing values of J as opposed to preparing only the ground state of the final Hamiltonian at $J=1$.

For comparison, in figure 2(left, orange) we apply a VQE evolution to the same morphed Hamiltonian (effectively replacing the CoVaR subroutine with VQE). In the present demonstration all 5 runs (initialised randomly around the same initial parameters) got trapped in local optima significantly above the ground state energy—indeed it is a known limitation of VQE that it is vulnerable to getting trapped in local optima [13]. In contrast, the present adiabatic CoVaR approach clearly converges to the true ground state and achieves an energy difference $\Delta E = 0.055$ which is almost 2 orders of magnitude smaller. We note that the energy difference of the CoVaR solution is limited by the depth of the ansatz circuit—we use a coupling constant $J=1.0$ which requires a relatively expressive variational circuit to achieve a good representation of the ground state and indeed CoVaR achieves significantly lower energy differences for the intermediate ground states at $J < 1$, for example, ΔE is an order of magnitude smaller as $\Delta E \leq 10^{-3}$ for $J \leq 0.4$. Furthermore, we note that adiabatic CoVaR converges quickly to the instantaneous ground states at each time step as shown figure 2(left, blue), and thus likely fewer total iterations could be used than 1050—the relatively high iteration count was used to ensure a fair comparison with VQE.

In figure 2(left, green) we also compare to conventional VQE and plot the evolution of the energy of 10 different rounds that were all initialised randomly in the vicinity of the ground state of the trivial model. While VQE can initially rapidly lower the energy, all runs get trapped in local optima, far away from the true

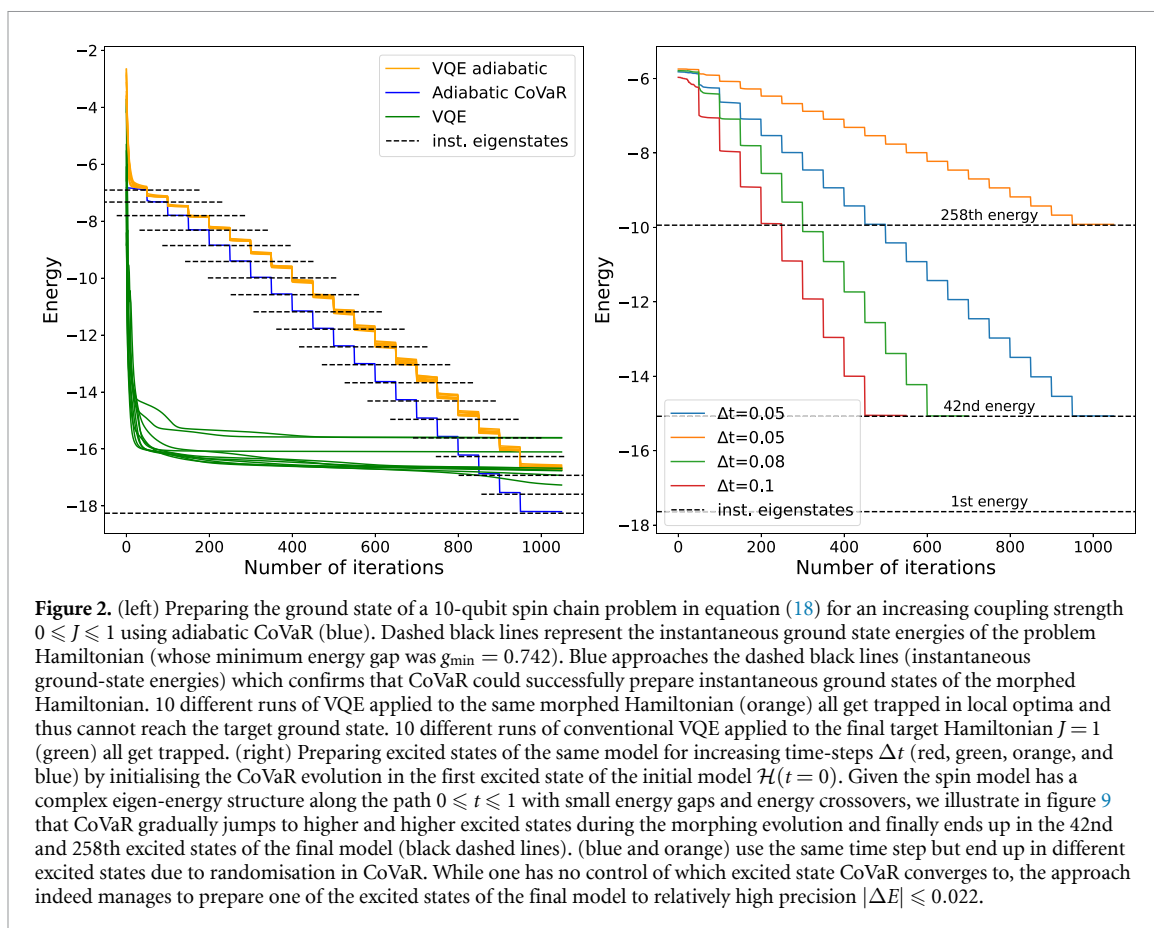


Figure 2. (left) Preparing the ground state of a 10-qubit spin chain problem in equation (18) for an increasing coupling strength $0 \leq J \leq 1$ using adiabatic CoVaR (blue). Dashed black lines represent the instantaneous ground state energies of the problem Hamiltonian (whose minimum energy gap was $g_{\min} = 0.742$). Blue approaches the dashed black lines (instantaneous ground-state energies) which confirms that CoVaR could successfully prepare instantaneous ground states of the morphed Hamiltonian. 10 different runs of VQE applied to the same morphed Hamiltonian (orange) all get trapped in local optima and thus cannot reach the target ground state. 10 different runs of conventional VQE applied to the final target Hamiltonian $J = 1$ (green) all get trapped. (right) Preparing excited states of the same model for increasing time-steps Δt (red, green, orange, and blue) by initialising the CoVaR evolution in the first excited state of the initial model $\mathcal{H}(t=0)$. Given the spin model has a complex eigen-energy structure along the path $0 \leq t \leq 1$ with small energy gaps and energy crossovers, we illustrate in figure 9 that CoVaR gradually jumps to higher and higher excited states during the morphing evolution and finally ends up in the 42nd and 258th excited states of the final model (black dashed lines). (blue and orange) use the same time step but end up in different excited states due to randomisation in CoVaR. While one has no control of which excited state CoVaR converges to, the approach indeed manages to prepare one of the excited states of the final model to relatively high precision $|\Delta E| \leq 0.022$.

ground state as consistent with theoretical expectations [13]. Specifically, [13] proved that in a large class of problems, only a super-polynomially vanishing fraction of all local minima are found in a fixed region of the ground state energy.

Finally, we perform simulations for preparing excited states of the present spin model, focusing only on assessing the performance of CoVaR rather than comparing to VQE-based deflation methods. Figure 2(right) shows adiabatic CoVaR evolutions for different time-steps $\Delta t \in \{0.01, 0.05, 0.08, 0.1\}$ when the evolution is initialised in the first excited state of the initial trivial Hamiltonian \mathcal{H}_0 . If Δt was very small and there were no energy-level crossovers then the approach could prepare the first excited state of the final model. However, we will illustrate below in more detail that when Δt is relatively large and small gaps are present, CoVaR may converge to other energy eigenstates that are close in parameter space, potentially resulting in excitations. In figure 9 we plot how adiabatic CoVaR gradually jumps to higher and higher excited states and, indeed figure 2(right) demonstrates that following multiple consecutive jumps, adiabatic CoVaR ends up preparing the 42nd and 258th excited states of the final model ($J = 1$). In appendix C we also detail that both Figure 2(right, orange and blue) use exactly the same hyperparameters, however, due to randomisation in CoVaR, we end up in a significantly higher excited state in figure 2(right, orange).

Still, it appears to be a general feature of CoVaR that even though one cannot fully control which excited state to prepare, at least CoVaR prepares a very good approximation to one of the eigenstates as $\Delta E = 0.007$ (blue), $\Delta E = 0.022$ (orange), $\Delta E = 0.007$ (green), and $\Delta E = 0.021$ (red). While we only feature excited states of the final model ($J = 1$) in figure 2(right, dashed lines), CoVaR indeed prepares excited states of the parametric model along the entire path $0 < J \leq 1$. We can also infer this from figure 9, which confirms that the present approach prepares a range of low and intermediate excited states for the entire family of Hamiltonians for increasing J . A potential application as estimating thermal properties could be enabled by running CoVaR on multiple excited initial states in order to map out the low-energy eigenspace of a Hamiltonian. This process could be aided by existing techniques, such as VQE or complementary techniques that can find approximations to eigenstates through resonant transitions [40, 41].

Finally, we perform further numerical analysis in appendices B.1 and B.2, whereby we systematically explore how different hyperparameters affect the performance of the present approach. These results indicate that the number of variational layers can significantly affect the number of iterations required to reach a desired precision. The reason is that choosing a low number of layers results in a subspace spanned by the

variational ansatz that is not sufficiently expressive to contain the relevant eigenstates of all instantaneous Hamiltonians. In appendix D we also track the evolution of the imprecision ΔE throughout the adiabatic path. While we find that in the particular system simulated, ground-state preparation was rather stable, the imprecision monotonically increased throughout the evolution path suggesting an increased complexity of preparing excited states. Specifically, while the ansatz circuit can successfully prepare the initial states, it is not sufficiently expressive to generate the evolving eigenstates of the morphed Hamiltonian as the circuit depth is kept constant and shallow.

3.2. Eigenstate discovery of the lattice Schwinger model

Here we consider a simple toy model in High-Energy Physics following [42] whereby the lattice Schwinger model was mapped to a spin Hamiltonian by enforcing Gauss' law and following a Jordan–Wigner transformation as $\mathcal{H}_{\text{sch}} = \mathcal{H}_{ZZ} + \mathcal{H}_{\pm} + \mathcal{H}_Z$ where the individual Hamiltonian terms are defined as

$$\begin{aligned}\mathcal{H}_{ZZ} &= \frac{J}{2} \sum_{n=2}^{N-1} \sum_{1 \leq k \leq l \leq n} Z_k Z_l, \\ \mathcal{H}_{\pm} &= \frac{J}{2} \sum_{n=1}^{N-1} \left[w - (-1)^n \frac{m}{2} \sin(\theta) \right] [X_n X_{n+1} + Y_n Y_{n+1}], \\ \mathcal{H}_Z &= \frac{m \cos \theta}{2} \sum_{n=1}^N (-1)^n Z_n - \frac{J}{2} \sum_{n=1}^{N-1} (n \bmod 2) \sum_{l=1}^n Z_l.\end{aligned}$$

Chakraborty *et al* [42] then used this Hamiltonian to prepare massless and massive eigenstates of the lattice Schwinger model using digital quantum algorithms. Here we consider an adiabatic mixing approach via the following time dependent Hamiltonian as

$$\mathcal{H}(t) = (1-t) \sum_i^N X_i + t \mathcal{H}_{\text{sch}}. \quad (19)$$

In figure 3 we perform adiabatic CoVaR on a 5-qubit (left) and on a 10-qubit (right) instance of this Hamiltonian. First, we initialise our approach in the ground state of the trivial mixing Hamiltonian and plot the evolution of the energies obtained via CoVaR in figure 3(blue). As the model has a significant ground-state energy gap for a smaller system (left), our approach manages to very closely follow the true ground state, however, for the larger system size (right) CoVaR eventually converges to the first excited state.

Second, we demonstrate how our approach extends ground-state preparations, as relevant in [42], to excited state preparations by initialising our approach in low lying eigenstates of the initial model at $t = \Delta t$ (given the energy levels are degenerate at $t = 0$). While, again, the CoVaR evolutions in figure 3(orange, purple, red) undergo energy excitations around small intermediate energy gaps, the approach successfully prepares one of the eigenstates of the final model to high precision (all final energies are within a distance $\Delta E \leq 10^{-3}$ from nearest energy eigenstates).

We also note that a clear difference from conventional adiabatic evolution can be observed in figure 3(left, red, $0.3 \leq t \leq 0.4$), whereby excitations to higher energy eigenstates happen in regions where the energy gaps are relatively large. This is likely caused by closeness of those eigenstates in parameter space or high overlaps (fidelity) with the excited energy eigenstates. Furthermore, CoVaR converged to the excited states with high precision in figure 3(right, purple, $0.6 \leq t \leq 0.8$) despite a large step size $\Delta t = 0.15$ and a number of nearby eigenstates with vanishingly small gaps—in contrast phase estimation and conventional adiabatic evolution would require long evolution times that scale inversely with the energy gaps g_{min} .

3.3. Solving combinatorial optimisation problems

We now demonstrate that adiabatic CoVaR can naturally be applied to approximating solutions to combinatorial optimisation problems. We consider a weighted max-cut problem which can be encoded [36, 43] as the ground state of the following Hamiltonian as

$$\mathcal{H} = \sum_i w_i Z_i + \sum_{i < j} w_{ij} Z_i Z_j, \quad (20)$$

where we uniformly randomly generate the weights $0 \leq w_i \leq 1$ and $0 \leq w_{ij} \leq 1$. Given all terms in the Hamiltonian commute, and their common eigenstates are computational basis states, one cannot apply the




Figure 3. Scaling up adiabatic CoVaR: energy evolution in a 5-qubit (left) and 10-qubit (right) Lattice Schwinger Hamiltonian. Dashed lines are instantaneous eigenstates of the time-dependent Hamiltonian in equation (19) using parameters $J = 1.0$, $m = 0.1$, $\theta = 0$, and $w = 0.1$. Solid lines indicate the evolution of the energy using adiabatic CoVaR with a time step $\Delta t = 0.15$. Blue represents ground-state preparation by initialising in the ground state of the initial model. Low-lying excited states of the final model were prepared by initialising in the first (orange) and second (purple) excited states of the initial model.

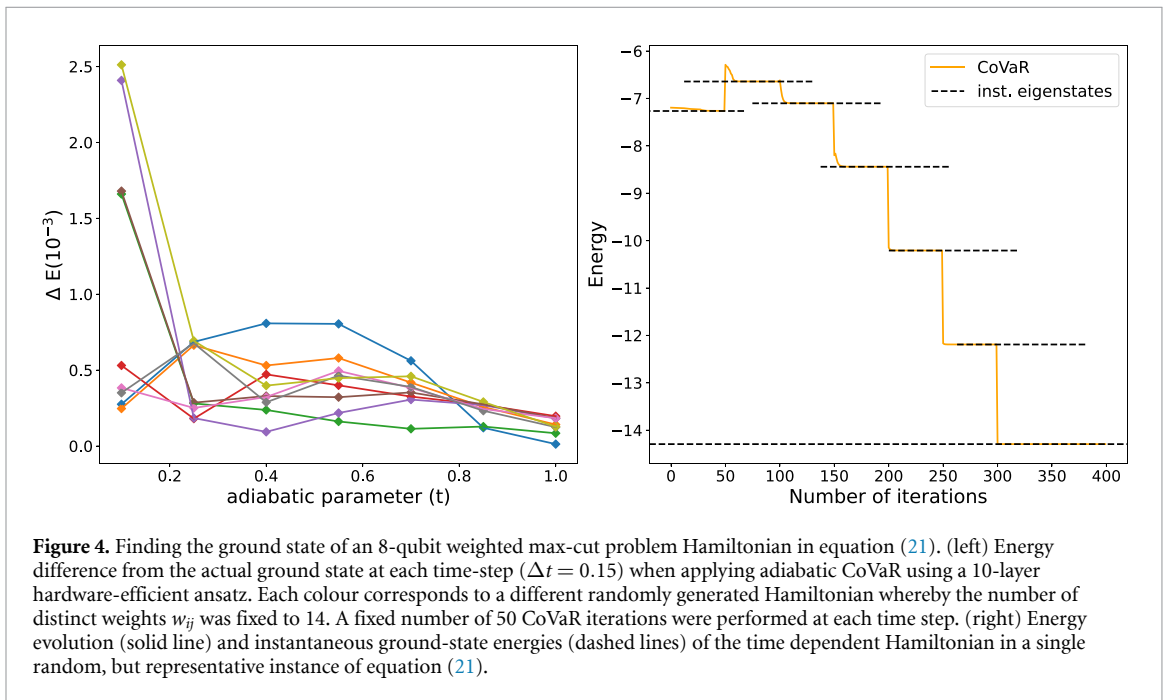


Figure 4. Finding the ground state of an 8-qubit weighted max-cut problem Hamiltonian in equation (21). (left) Energy difference from the actual ground state at each time-step ($\Delta t = 0.15$) when applying adiabatic CoVaR using a 10-layer hardware-efficient ansatz. Each colour corresponds to a different randomly generated Hamiltonian whereby the number of distinct weights w_{ij} was fixed to 14. A fixed number of 50 CoVaR iterations were performed at each time step. (right) Energy evolution (solid line) and instantaneous ground-state energies (dashed lines) of the time dependent Hamiltonian in a single random, but representative instance of equation (21).

perturbative approach. Instead, most commonly one uses a mixer Hamiltonian $\sum_i X_i$ to construct the time-dependent Hamiltonian as

$$\mathcal{H}(t) = (1-t) \sum_i X_i + t \left\{ \sum_i w_i Z_i + \sum_{i < j} w_{ij} Z_i Z_j \right\}. \quad (21)$$

Figure 4(left) shows the difference from the instantaneous ground states for 10 different, randomly generated Hamiltonians, and demonstrates that CoVaR indeed closely follows instantaneous ground states, with an energy difference $\Delta E \leq 2.5 \times 10^{-3}$. It is worth noting that these simulations did not include shot noise in order to test the ultimate performance of the approach and indeed we analyse the effect of shot noise separately in section 4.2. We also note that the final precision achieved here ($\Delta E \leq 5 \times 10^{-4}$) is significantly




Figure 5. (left) Largest time-step Δt sufficient to achieve a final state preparation $\Delta E = \pm 0.0015$ vs minimum energy gap g_{\min} of the Hamiltonian. Error bars indicate variance across multiple randomly generated, 7-qubit spin Hamiltonians from equation (18). For illustrative purposes, we plot (orange dotted line) the quadratic scaling of Δt in conventional adiabatic evolution. This demonstrates that adiabatic CoVaR appears to have a fundamentally different dependence on the energy gap. (right) Final energy precision ΔE with (orange) and without (blue) shot noise confirming robustness of CoVaR against shot noise as long as the level of shot noise is slightly smaller than the sought final precision. The dashed line at $\Delta E = 10^{-3}$ indicates a target precision that is sought in typical applications, e.g. chemical accuracy.

higher than in case of the spin problem in figure 2 given the ground state of the maxcut Hamiltonian is a computational basis state which can be prepared exactly using our shallow ansatz.

To illustrate the effectiveness of CoVaR, in figure 4(right) we plot the evolution of a single, representative random instance and demonstrate that CoVaR (orange solid line) rapidly converges to the instantaneous ground states (black dashed lines). Figure 4(right) also features sudden jumps in the energy evolution curve; We note here that while VQE is guaranteed to monotonically and smoothly decrease the energy, CoVaR makes no reference to the energy, i.e. it finds joint roots of covariances. Therefore, we do not expect CoVaR to yield a smooth, monotonic evolution of the energy.

4. Further numerical analysis

4.1. Relationship between energy gap and step size

In the conventional adiabatic time evolution approach a sufficiently small timestep $\Delta t \sim g_{\min}^2$ needs to be chosen that is proportional to the square of the minimal energy gap in order to prevent excitation to higher eigenstates [44]. In figure 5(left) we numerically estimate the largest feasible Δt required to prepare the ground state of a series of randomly generated (by randomly sampling c_i) 7-qubit spin problems (equation (18)) with different energy gaps g_{\min} .

The value of Δt was found through a line search whereby the largest Δt was accepted that achieved a final state preparation error $\epsilon \leq 0.0015$. The error bars indicate variance across the different random Hamiltonians. Fitting a non-linear curve in figure 5(left) suggests a logarithmic scaling $(\Delta t)^{-1} \sim \ln(g_{\min})$ within the finite regime of g_{\min} explored. This is in stark contrast to the conventional adiabatic evolution approach whose scaling is quadratic $\Delta t \propto g_{\min}^2$, which we illustrate in figure 5(left, dashed orange line).

4.2. Robustness to shot noise

Finally, we confirm the robustness of our approach to shot noise. We use the approach of [15, 45] to simulate the effect of shot noise by adding Gaussian noise of standard deviation $N_s^{-1/2}$ to individual covariances and assume a number of shots $N_s = 10^6$. We again consider the ground-state preparation of the 7-qubit spin Hamiltonian in equation (18) using an ansatz of 16 layers.

Figure 5(right) shows the final achieved energy error ΔE both with (orange) and without (blue) shot noise for increasing values of $0.01 \leq \Delta t \leq 0.1$. Comparing the blue and orange curves it is clear that the present level of shot noise does not significantly influence the performance of the approach.

These results nicely confirm observations of [15] whereby a similar robustness of CoVaR against shot noise was observed and whereby the primary limitation due to shot noise was found that the final precision reached by CoVaR is limited as $\Delta E \propto N_s^{-1/2}$, and here $\Delta E \approx 10^{-3}$.

5. Discussion and conclusion

We consider CoVaR [15] which allows one to find eigenstates of a Hamiltonian; Parameters of an ansatz circuit are guided through a root-finding approach to ultimately find joint roots of a large number of covariance functions. Joint roots guarantee that the quantum state prepared by the parametrised circuit is an eigenstate of the problem Hamiltonian. A large number of covariances ($\propto 10^8$ – 10^{10} is reasonable using HPC resources) can be extracted efficiently for large systems using classical shadows, guaranteeing remarkable robustness and convergence speed of the approach. However, CoVaR is oblivious as to which eigenstate it converges to and thus its vanilla implementation required an initial state that has high overlap with the desired target eigenstate—an assumption most typical in quantum simulation.

In the present work we attempt to address this limitation and combine the powerful CoVaR approach with a discretised morphing of an initial trivial Hamiltonian into the final desired problem Hamiltonian. By design, the initial trivial Hamiltonian's eigenstates can be solved analytically and can be prepared with a shallow circuit; the approach is then efficiently initialised in one of these eigenstates and, after slightly adjusting the Hamiltonian as in adiabatic computing, CoVaR is employed to prepare the instantaneous eigenstates of the time-dependent Hamiltonian. This way the initial state supplied to the CoVaR subroutine has a high overlap with the relevant instantaneous eigenstates.

We numerically simulated the present approach in a broad range of practically relevant applications. First, we considered a perturbative approach whereby an initial trivial Hamiltonian is morphed by gradually turning on a strong coupling term. We compared adiabatic CoVaR in a 10-qubit spin chain system to two different implementations of VQE. (a) By performing standard VQE using the same initial parameters that prepare the ground state of the initial trivial model. (b) We performed an adiabatic VQE by effectively using VQE instead of CoVaR in the present morphing approach. All instances of these VQE optimisations failed and got stuck in local minima whereas the present approach converged to the true ground state in all instances with relatively high precision.

Second, a broad application area where adiabatic techniques are very naturally applicable are combinatorial optimisation problems [43]. We apply adiabatic CoVaR to solve weighted max-cut problems and demonstrate its effectiveness by successfully preparing exact ground states in all instances.

A significant strength of adiabatic CoVaR is that it naturally allows for the direct preparation of excited states similarly to phase estimation—this is in contrast to variants of VQE, such as VQD, that prepare excited states one-by-one. We demonstrated excited-state preparation in two practically relevant application examples: in preparing highly excited states in a random Heisenberg model and in the Schwinger model which is a toy model relevant in High-Energy-Physics applications. While our approach is demonstrably robust against local traps and did in all numerical experiments converge to an eigenstate, one does not have full control as to which eigenstate we prepare. In particular, depending on the choice of Δt , small energy gaps typically led to convergence to nearby eigenstates in parameter space, effectively resulting in an excitation or de-excitation of the variational state.

However, the most remarkable feature of adiabatic CoVaR is that it appears to tolerate small or vanishing energy gaps in the following sense. We numerically demonstrated that, while one does not have full control as to which eigenstate CoVaR converges to, the present approach successfully prepares an eigenstate even when the energy gaps are quite small and the time step is very large as $\Delta t = 0.15$. In one example we empirically found that the adiabatic time step Δt scaled logarithmically with the minimum energy gap g_{\min} of the problem Hamiltonian. This is a significant and unique advantage of the present approach as even phase estimation—when used to prepare instantaneous eigenstates of the time-dependent Hamiltonian—requires deep quantum circuits due to long evolution times $T \propto g_{\min}^{-1}$ when energy gaps are too small. Similarly, this dependence is quadratic as $T \propto g_{\min}^{-2}$ in conventional adiabatic evolution.

Finally, we analysed the robustness of the present approach against imperfections and against variability in hyperparameter settings. Our numerical results confirm prior observations that CoVaR is robust against shot noise. Furthermore, [15] confirmed that CoVaR is robust against gate noise given the perturbations to the ideal Jacobian matrix due to gate imperfections are averaged out due to the large dimensionality of the Jacobian matrix – and also because the update rule in (16) is invariant under averaged out global depolarizing noise. Still, we note that a broad range of error mitigation techniques [46–48] are immediately applicable to CoVaR through error mitigated classical shadows [49].

Finally, we note that we focused on conventional ‘Pauli shadows’ [16] which are straightforward to implement. However, a broad range of more advanced shadow measurement techniques are available that are immediately applicable to extend the range of the present approach or to make it more shot efficient including shallow shadows [50] or fermionic shadows [51, 52]. While the latter enables to study ground and excited states of molecular and materials systems [7, 53–55], it may require increased circuit depths for implementing matchgate scrambling circuits. Implementations on early fault-tolerant architectures are

therefore an exciting and promising future direction which may be enabled by techniques in [56, 57] for overcoming technical challenges associated with continuous angle rotations in our ansatz circuits [58].

Data availability statement

The core algorithm presented in this work is fully implemented and demonstrated in the repository openly available at https://github.com/seop02/adiabatic_covar/tree/main. All data that support the findings of this study are included within the article (and any supplementary files).

Acknowledgments

The authors thank Bipasha Chakraborty for helpful discussions. BK thanks the University of Oxford for a Glasstone Research Fellowship and Lady Margaret Hall, Oxford for a Research Fellowship. All simulations were performed using the open-source tools Quantum Exact Simulation Toolkit (QuEST) [59] and its Mathematica-based interface QuEST link [60]. BK thanks UKRI for the Future Leaders Fellowship Theory to Enable Practical Quantum Advantage (MR/Y015843/1). The authors also acknowledge funding from the EPSRC projects Robust and Reliable Quantum Computing (RoQRQ, EP/W032635/1) and Software Enabling Early Quantum Advantage (SEEQA, EP/Y004655/1). This research was funded in whole, or in part, by UKRI (MR/Y015843/1). For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

Appendix A. Comparison to prior techniques

A.1. Variational quantum deflation (VQD)

VQD is a variant of VQE for preparing excited states [6, 61]. The approach augments the VQE cost function with additional terms that account for the overlap with known energy eigenstates, i.e. the VQD cost function for finding the k th excited state is

$$C_k(\boldsymbol{\theta}_k) = \langle \psi(\boldsymbol{\theta}_k) | \mathcal{H} | \psi(\boldsymbol{\theta}_k) \rangle + \sum_{i=0}^{k-1} \beta_i |\langle \psi(\boldsymbol{\theta}_k) | \psi(\boldsymbol{\theta}_i) \rangle|^2.$$

Here β_i is an arbitrary hyperparameter and $\boldsymbol{\theta}_i$ are parameter settings that approximate the i th excited energy eigenstate $|\psi_i\rangle$.

This ensures that we minimize the energy with the constraint that the variational state be orthogonal to the previously found eigenstates [29]. The downside of the approach is that one requires explicit knowledge of the parameter settings $\boldsymbol{\theta}_0, \dots, \boldsymbol{\theta}_{k-1}$ up to the $(k-1)$ th excited state in order to find the k th excited state, which in practice is achieved by iteratively finding higher and higher excited states.

For this reason, preparing highly excited states may be challenging with VQD given the states are prepared incrementally and thus errors in the lower excited states are expected to accumulate, i.e. at each step one is vulnerable to barren plateaus and to getting trapped in local minima.

Appendix B. Analysing dependence on hyper parameters

B.1. Effect of adiabatic time-step and circuit layers on the precision

In all simulations in the present work we used a hardware efficient ansatz illustrated in figure 6. Here, we present how the number of circuit layers influences the precision of the state preparation. Simulation results for both ground state energy and first excited states are presented for a 7-qubit spin problem from equation (18) in 7 for an increasing number of layers between 10 and 20 and for an increasing step size Δt .

While 10 variational layers in figure 7(left and right) are not sufficient for achieving a target precision of 10^{-3} , increasing the number of layers to 14 and 20 clearly improves the performance. Furthermore, it appears the performance of adiabatic CoVaR is fairly robust against the step size Δ as increasing Δ does not significantly affect the final energy difference.

Increasing the number of parameters does indeed improve the precision, however, comes with the trade-off of an increased circuit depth and also an increased runtime through an increased number of circuit runs. Thus, in practice it is worth minimising the number of layers as a hyper parameter.

Here, we present how the number of parameters in the variational circuit, specifically the number of variational layers, influences the precision of the final energy. Simulation results for both ground state energy and first excited states are presented in figure 7. All simulations in this section were conducted on a 7-qubit

system, with the number of layers ranging from 10 to 20. We illustrate results for 10, 14, and 20 layers, which were sufficient to illustrate the trend.




Figure 6. Example of a single layer of a 6-qubit hardware-efficient ansatz used in the present work. The vertical lines denote continuous-angle Pauli rotations of the form $\exp(-i \frac{\theta}{2} P_i P_j)$, where θ is a variational parameter. Here, P_i and P_j are single-qubit Pauli operators, e.g. $X_1 X_2$.




Figure 7. Final energy difference from the exact ground (left) and first excited state (right) for an increasing number of ansatz layers and an increasing step size Δt in a 7-qubit spin Hamiltonian from equation (18). The black horizontal dashed line represents a typical target precision, e.g. chemical accuracy.

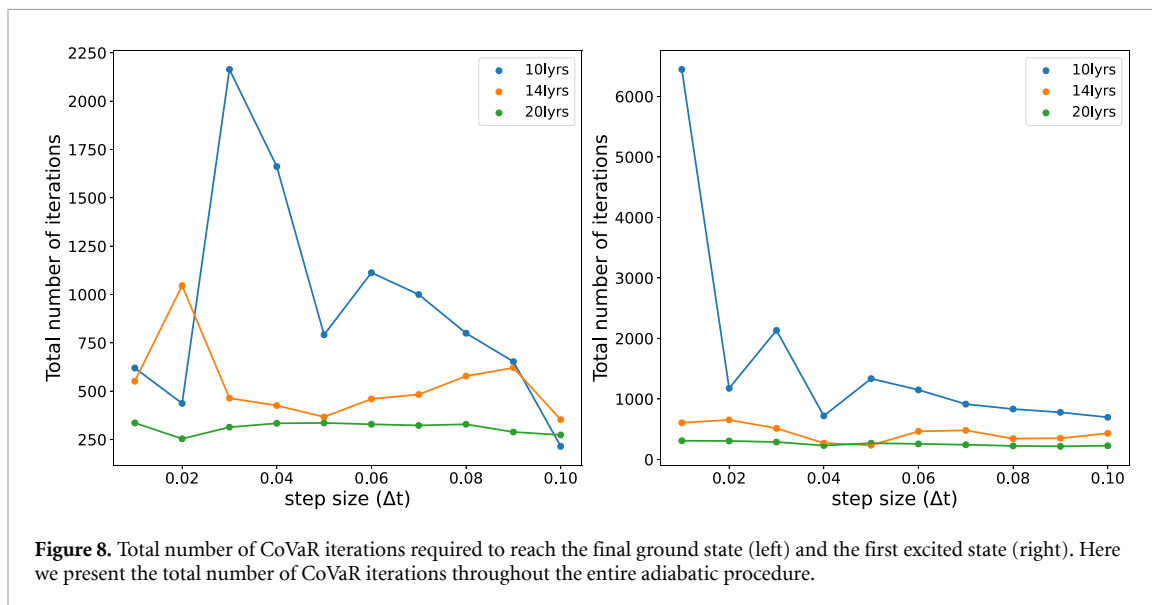
Figure 7 illustrates that 10 variational layers are insufficient for a 7-qubit system. For the first excited state computation, it does not achieve chemical accuracy for all adiabatic runtimes (different Δt values). For the ground state computation, chemical accuracy is only achieved for $\Delta t = 0.01$. Precision notably improves for 14 and 20 layers. In simulations with 14 layers, chemical accuracy is attained for the ground state when $\Delta t = 0.03$ and $\Delta t = 0.1$. For the first excited state, chemical accuracy is achieved for Δt ranging from $\Delta t = 0.03$ to $\Delta t = 0.08$. More accurate results are obtained with 20 layers. For the ground state computation, chemical accuracy is achieved for almost all Δt values except for $\Delta t = 0.01$ and $\Delta t = 0.07$. For the first excited state computation, chemical accuracy is achieved for almost all Δt values except for $\Delta t = 0.06$.

The results imply that the precision of the final energy is heavily dependent on the number of parameters within the variational circuit. Increasing the number of parameters does tend to improve the precision. However, we have to note that there exists a trade-off between the number of parameters and the computation run-time. More parameters result in increased computational time for each CoVaR iteration, as we would have to compute more covariances. To optimize the computation process, it is important to find the most optimum number of layers that reach the chemical accuracy while minimizing the computation time.

B.2. Effect of adiabatic time-step and circuit layers on the number of total iterations

The total number of CoVaR iterations used for the entire adiabatic morphing is strongly influenced by the expressivity of the ansatz. In particular, in the present implementation of adiabatic CoVaR we demand that for each adiabatic timestep CoVaR must converge to the instantaneous eigenstate to precision $\epsilon = 0.002$ in the covariance vector norm. Passing this threshold, CoVaR terminates, and the computation proceeds to the next time step.

Figure 8 illustrates that circuits with 10 layers required the highest number of CoVaR iterations to reach the final eigenstate in contrast to circuits with more layers. The reason is likely that the shallow variational state was not able to sufficiently well express the instantaneous eigenstates throughout the evolution and thereby the threshold condition was not always passed as further elucidated in figure 10. As the number of layers is increased, the average number of CoVaR iterations decreases significantly.



Our analysis reveals a complex interplay between run-time, iteration count, and layer number in quantum computations. As run-time decreases due to increasing Δt , the total number of iterations generally decreases across all layers, an expected outcome given the reduced maximum possible iterations with larger Δt values. However, we observe a notable shift in this trend when Δt exceeds 0.05, where iterations begin to increase for both ground state and first excited state calculations, suggesting an optimal Δt for iteration minimization. In the case of first excited state computations, the 10-layer configuration required significantly more iterations compared to other layer counts, while 14 and 20 layers exhibited similar iteration counts. This similarity suggests that 14 layers may be sufficient to achieve energy within chemical accuracy, although it is important to note that similar iteration counts do not necessarily imply equivalent computation times due to the increased parameter count in the 20-layer configuration. The increase in iterations at larger Δt values can be attributed to the initialization of each CoVaR computation farther from the target eigenstate, necessitating more iterations for convergence. Intriguingly, we found that the optimal Δt for minimizing iterations coincides with the precision required to achieve chemical accuracy, a finding that merits further investigation in future studies.

Appendix C. Details of figure 2

A 30-layer hardware-efficient ansatz was used with $\Delta t = 0.05$. Performing a single iteration of CoVaR is nearly of the same quantum cost as performing an iteration of VQE. For a fair comparison, 50 iterations were performed for $t < 1$, and 100 iterations at $t = 1$ for the adiabatic CoVaR (blue) and adiabatic VQE (orange) while the same number 1050 of total iterations were used for conventional VQE.

Figure 9 presents the index of energy levels obtained initialising adiabatic CoVaR in the first excited state of the initial Hamiltonian in figure 2(right). We performed to different simulations using the time step $\Delta t = 0.05$, and in one case the energy levels were more rapidly increased than in case of all other simulations. This can be primarily attributed to the random nature of the initial variational parameter in CoVaR. The randomness inherent in CoVaR may lead to significant differences in the evolutionary trajectories. Demonstrably, even with identical hyper parameters, the resulting trajectories are radically different.

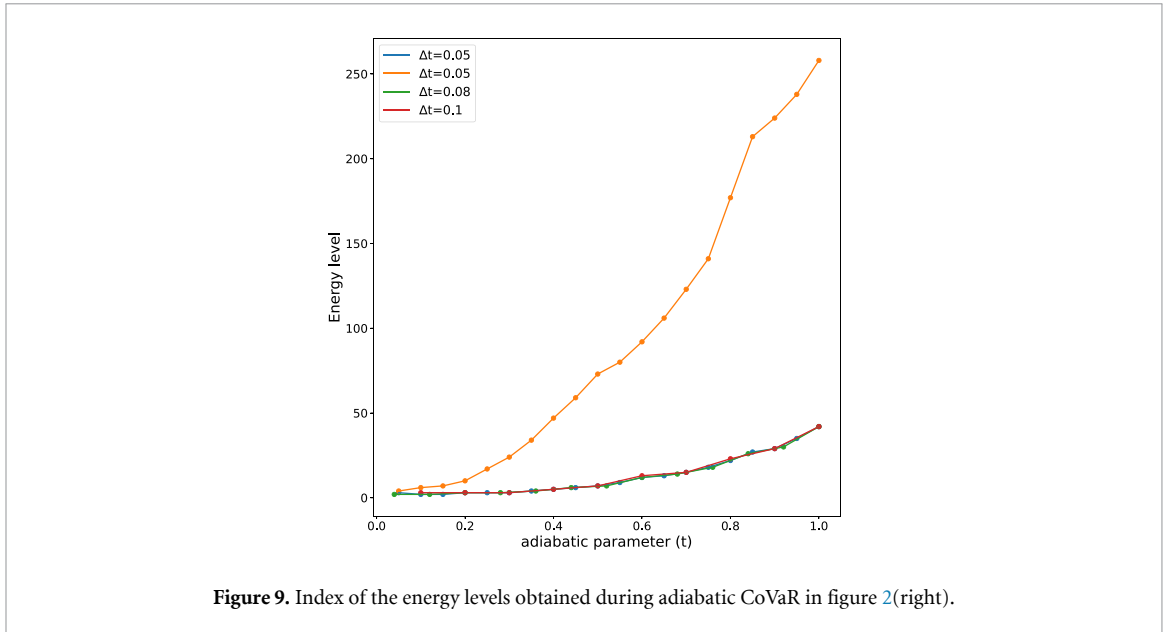


Figure 9. Index of the energy levels obtained during adiabatic CoVaR in figure 2(right).

Appendix D. Details of figure 10

Here, we present the evolution of the energy difference between the variational state and the target eigenstate with insufficient number of variational layers. A 10-layer hardware-efficient ansatz was used for computations of both ground state and the first excited state in the spin chain model. In order to illustrate the effect of insufficient variational layers on the precision of the final energy obtained by adiabatic CoVaR, we fixed the number of variational layers and ran multiple simulations varying the adiabatic runtime(Δt).

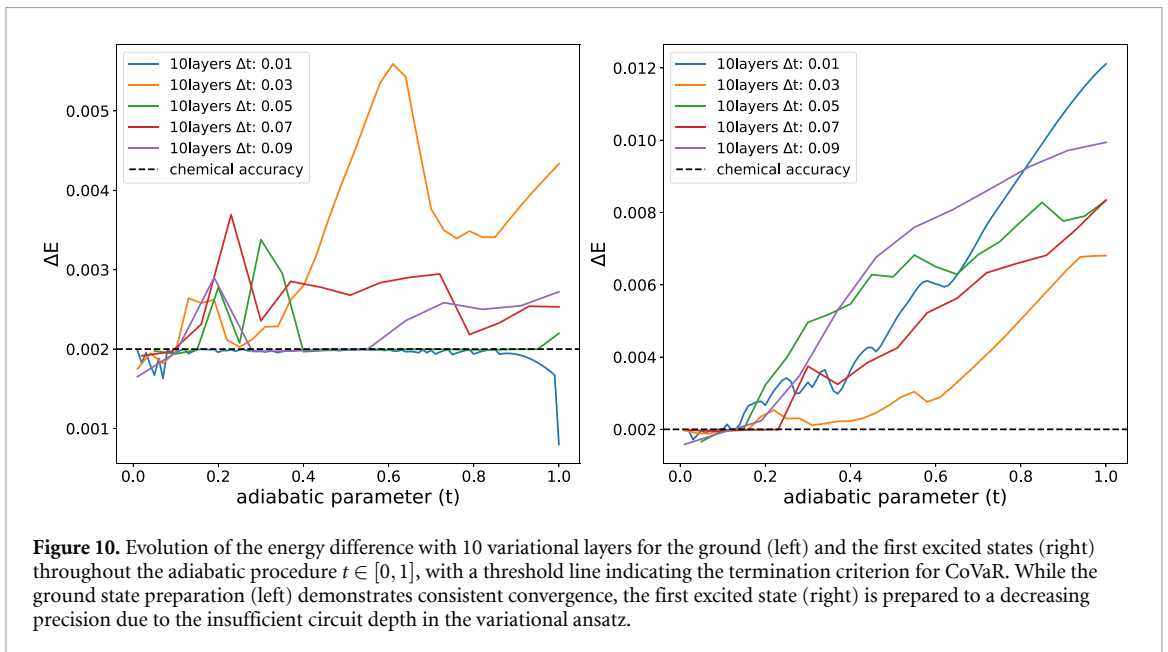


Figure 10. Evolution of the energy difference with 10 variational layers for the ground (left) and the first excited states (right) throughout the adiabatic procedure $t \in [0, 1]$, with a threshold line indicating the termination criterion for CoVaR. While the ground state preparation (left) demonstrates consistent convergence, the first excited state (right) is prepared to a decreasing precision due to the insufficient circuit depth in the variational ansatz.

Figure 10 depicts the energy difference evolution for the spin chain system. For the ground state computation, most runtimes failed to achieve chemical accuracy, with $\Delta t = 0.01$ being the sole exception. The first excited state computations fared worse, with all runtimes failing to reach chemical accuracy and the energy difference diverging throughout the adiabatic evolution. These results stem from our use of a coupling constant $J = 1.0$, which led to a highly entangled and mixed final state, markedly different from the initial unperturbed eigenstate. This increased entanglement presented a formidable challenge for eigenstate computation, necessitating greater circuit depths and resulting in a gradual increase in energy deviation. Furthermore, the adiabatic evolution of the Hamiltonian caused a shift in the system's eigenstates, expanding

the Hilbert space that the variational ansatz needed to span. This expansion further underscored the need for larger variational depths in adiabatic CoVaR. These findings highlight the crucial importance of carefully selecting both the number of variational layers and runtime parameters in adiabatic CoVaR computations, especially when dealing with highly entangled systems. To address these challenges, future research could explore adaptive layer techniques or alternative ansatz structures, potentially offering more robust solutions for these complex scenarios.

ORCID iDs

Wooseop Hwang  <https://orcid.org/0009-0006-4717-4026>

Bálint Koczor  <https://orcid.org/0000-0002-4319-6870>

References

- [1] Preskill J 2018 Quantum computing in the NISQ era and beyond *Quantum* **2** 79
- [2] Peruzzo A, McClean J, Shadbolt P, Yung M-H, Zhou X-Q, Love P J, Aspuru-Guzik A and O'Brien J L 2014 A variational eigenvalue solver on a photonic quantum processor *Nat. Commun.* **5** 4213
- [3] Endo S, Cai Z, Benjamin S C and Yuan X 2021 Hybrid quantum-classical algorithms and quantum error mitigation *J. Phys. Soc. Japan* **90** 032001
- [4] Cerezo M *et al* 2021 Variational quantum algorithms *Nat. Rev. Phys.* **3** 625–44
- [5] Li Y and Benjamin S C 2017 Efficient variational quantum simulator incorporating active error minimization *Phys. Rev. X* **7** 021050
- [6] Higgott O, Wang D and Brierley S 2019 Variational quantum computation of excited states *Quantum* **3** 156
- [7] Gao Q, Jones G O, Motta M, Sugawara M, Watanabe H C, Kobayashi T, Watanabe E, Ohnishi Y-Y, Nakamura H and Yamamoto N 2021 Applications of quantum computing for investigations of electronic transitions in phenylsulfonyl-carbazole tADF emitters *npj Comput. Mater.* **7** 70
- [8] McArdle S, Jones T, Endo S, Li Y, Benjamin S C and Yuan X 2019 Variational ansatz-based quantum simulation of imaginary time evolution *npj Quantum Inf.* **5** 75
- [9] Koczor B and Benjamin S C 2022 Quantum natural gradient generalized to noisy and nonunitary circuits *Phys. Rev. A* **106** 062416
- [10] Barraza N, Alvarado Barrios G, Peng J, Lamata L, Solano E and Albarrán-Arriagada F 2022 Analog quantum approximate optimization algorithm *Quantum Sci. Technol.* **7** 045035
- [11] Tilly J *et al* 2022 The variational quantum eigensolver: a review of methods and best practices *Phys. Rep.* **986** 1
- [12] Bittel L and Kliesch M 2021 Training variational quantum algorithms is NP-hard *Phys. Rev. Lett.* **127** 120502
- [13] Anschuetz E R and Kiani B T 2022 Quantum variational algorithms are swamped with traps *Nat. Commun.* **13** 7760
- [14] McClean J R, Boixo S, Smelyanskiy V N, Babbush R and Neven H 2018 Barren plateaus in quantum neural network training landscapes *Nat. Commun.* **9** 4812
- [15] Boyd G and Koczor B 2022 Training variational quantum circuits with CoVaR: Covariance root finding with classical shadows *Phys. Rev. X* **12** 041022
- [16] Huang H-Y, Kueng R and Preskill J 2020 Predicting many properties of a quantum system from very few measurements *Nat. Phys.* **16** 1050–7
- [17] Larocca M, Thanasilp S, Wang S, Sharma K, Biamonte J, Coles P J, Cincio L, McClean J R, Holmes Z and Cerezo M 2024 A review of barren plateaus in variational quantum computing (arXiv:2405.00781[quant-ph])
- [18] Koczor B and Benjamin S C 2022 Quantum analytic descent *Phys. Rev. Res.* **4** 02301
- [19] Dennis J E and Schnabel R B 1996 *Numerical Methods for Unconstrained Optimization and Nonlinear Equations* (Society for Industrial and Applied Mathematics)
- [20] Chan H H S, Meister R, Goh M L and Koczor B 2022 Algorithmic shadow spectroscopy *PRX Quantum* accepted
- [21] Born M and Fock V 1928 Beweis des adiabatsatzes *Z. Phys.* **51** 165
- [22] Farhi E, Goldstone J, Gutmann S and Sipser M 2000 Quantum computation by adiabatic evolution (arXiv:quant-ph/0001106[quant-ph])
- [23] Shevchenko S, Ashhab S and Nori F 2010 Landau-zener-stückelberg interferometry *Phys. Rep.* **492** 1–30
- [24] Yuan X, Endo S, Zhao Q, Li Y and Benjamin S C 2019 Theory of variational quantum simulation *Quantum* **3** 191
- [25] Berry D W, Childs A M and Kothari R 2015 Hamiltonian simulation with nearly optimal dependence on all parameters *2015 IEEE 56th Annual Symp. on Foundations of Computer Science (IEEE)*
- [26] Low G H and Chuang I L 2019 Hamiltonian simulation by qubitization *Quantum* **3** 163
- [27] Wecker D, Hastings M B and Troyer M 2015 Progress towards practical quantum variational algorithms *Phys. Rev. A* **92** 042303
- [28] Mc Keever C and Lubasch M 2024 Towards adiabatic quantum computing using compressed quantum circuits *PRX Quantum* **5** 020362
- [29] Gocho S, Nakamura H, Kanno S, Gao Q, Kobayashi T, Inagaki T and Hatanaka M 2023 Excited state calculations using variational quantum eigensolver with spin-restricted ansätze and automatically-adjusted constraints *npj Comput. Mater.* **9** 13
- [30] Zener C 1932 Non-adiabatic crossing of energy levels *Proc. R. Soc. A* **137** 833
- [31] Bergou E, Diouane Y, Kungurtsev V and Royer C W 2021 A stochastic levenberg-marquardt method using random models with complexity results (arXiv:1807.02176[math.OA])
- [32] Liew S S, Khalil-Hani M and Bakhteri R 2016 An optimized second order stochastic learning algorithm for neural network training *Neurocomputing* **186** 74
- [33] Pagano G *et al* 2020 Quantum approximate optimization of the long-range ising model with a trapped-ion quantum simulator *Proc. Natl Acad. Sci.* **117** 25396–401
- [34] Bharti K *et al* 2022 Noisy intermediate-scale quantum algorithms *Rev. Mod. Phys.* **94** 015004
- [35] Harrigan M P *et al* 2021 Quantum approximate optimization of non-planar graph problems on a planar superconducting processor *Nat. Phys.* **17** 332–6
- [36] Lucas A 2014 Ising formulations of many NP problems *Front. Phys.* **2**
- [37] Luitz D J, Laflorencie N and Alet F 2015 Many-body localization edge in the random-field Heisenberg chain *Phys. Rev. B* **91** 081103

- [38] Childs A M, Maslov D, Nam Y, Ross N J and Su Y 2018 Toward the first quantum simulation with quantum speedup *Proc. Natl Acad. Sci.* **115** 9456–61
- [39] Bosse J L and Montanaro A 2021 Probing ground state properties of the kagome antiferromagnetic Heisenberg model using the variational quantum eigensolver (arXiv:2108.08086[quant-ph])
- [40] Wang H, Ashhab S and Nori F 2012 Quantum algorithm for obtaining the energy spectrum of a physical system *Phys. Rev. A* **85** 062304
- [41] Li Z, Liu X, Wang H, Ashhab S, Cui J, Chen H, Peng X and Du J 2019 Quantum simulation of resonant transitions for solving the eigenproblem of an effective water hamiltonian *Phys. Rev. Lett.* **122** 090504
- [42] Chakraborty B, Honda M, Izubuchi T, Kikuchi Y and Tomiya A 2022 Classically emulated digital quantum simulation of the schwinger model with topological term via adiabatic state preparation (arXiv:2001.00485[hep-lat])
- [43] Farhi E, Goldstone J and Gutmann S 2014 A quantum approximate optimization algorithm (arXiv:1411.4028[quant-ph])
- [44] Djordjevic I B 2023 Chapter 12 - quantum machine learning *Quantum Communication, Quantum Networks and Quantum Sensing*, ed I B Djordjevic (Academic) pp 491–561
- [45] Boyd G, Koczor B and Cai Z 2024 High-dimensional subspace expansion using classical shadows (arXiv:2406.11533)
- [46] Cai Z, Babbush R, Benjamin S C, Endo S, Huggins W J, Li Y, McClean J R and O'Brien T E 2023 Quantum error mitigation *Rev. Mod. Phys.* **95** 045005
- [47] Koczor B 2021 Exponential error suppression for near-term quantum devices *Phys. Rev. X* **11** 031057
- [48] Koczor B 2021 The dominant eigenvector of a noisy quantum state *New J. Phys.* **23** 123047
- [49] Inane H, Steinberg J, Cai Z, Nguyen H C and Koczor B 2024 Quantum error mitigated classical shadows *PRX Quantum* **5** 010324
- [50] Bertoni C, Haferkamp J, Hinsche M, Ioannou M, Eisert J and Pashayan H 2022 Shallow shadows: expectation estimation using low-depth random clifford circuits (arXiv:2209.12924[cond-mat, physics:quant-ph])
- [51] Wan K, Huggins W J, Lee J and Babbush R 2023 Matchgate shadows for fermionic quantum simulation *Commun. Math. Phys.* **404** 629
- [52] Wu B and Koh D E 2024 Error-mitigated fermionic classical shadows on noisy quantum devices *npj Quantum Inf.* **10** 1
- [53] McArdle S, Endo S, Aspuru-Guzik A, Benjamin S C and Yuan X 2020 Quantum computational chemistry *Rev. Mod. Phys.* **92** 015003
- [54] Yoshioka N, Sato T, Nakagawa Y O, Ohnishi Y-Y and Mizukami W 2022 Variational quantum simulation for periodic materials *Phys. Rev. Res.* **4** 013052
- [55] Wang L, Jiang S, Mao Y, Li Z, Zhang Y and Li M 2024 Lithium-ion battery state of health estimation method based on variational quantum algorithm optimized stacking strategy *Energy Rep.* **11** 2877
- [56] Koczor B, Morton J J L and Benjamin S C 2024 Probabilistic interpolation of quantum rotation angles *Phys. Rev. Lett.* **132** 130602
- [57] Koczor B 2024 Sparse probabilistic synthesis of quantum operations *PRX Quantum* **5** 040352
- [58] Zimborás Z *et al* 2025 Myths around quantum computation before full fault tolerance: what no-go theorems rule out and what they don't (arXiv:2501.05694)
- [59] Jones T, Brown A, Bush I and Benjamin S C 2019 Quest and high performance simulation of quantum computers *Sci. Rep.* **9** 10736
- [60] Jones T and Benjamin S 2020 Questlink-mathematica embiggened by a hardware-optimised quantum emulator *Quantum Sci. Technol.* **5** 034012
- [61] Jones T, Endo S, McArdle S, Yuan X and Benjamin S C 2019 Variational quantum algorithms for discovering hamiltonian spectra *Phys. Rev. A* **99** 062304