

Randomized adaptive quantum state preparation

Alicia B. Magann¹, Sophia E. Economou,² and Christian Arenz¹³

¹*Quantum Algorithms and Applications Collaboratory, Sandia National Laboratories, Albuquerque, New Mexico 87185, USA*

²*Department of Physics, Virginia Tech, Blacksburg, Virginia 24061, USA*

³*School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, Arizona 85287, USA*



(Received 17 January 2023; accepted 1 September 2023; published 29 September 2023)

We develop an adaptive method for quantum state preparation that utilizes randomness as an essential component and that does not require classical optimization. Instead, a cost function is minimized to prepare a desired quantum state through an adaptively constructed quantum circuit, where each adaptive step is informed by feedback from gradient measurements in which the associated tangent space directions are randomized. We provide theoretical arguments and numerical evidence that convergence to the target state can be achieved for almost all initial states. We investigate different randomization procedures and develop lower bounds on the expected cost function change, which allows for drawing connections to barren plateaus and for assessing the applicability of the algorithm to large-scale problems.

DOI: 10.1103/PhysRevResearch.5.033227

I. INTRODUCTION

Methods for preparing quantum states are an integral component of any quantum technology. For example, the preparation of states that encode ground, excited, and thermal states of many-body systems is a key element in quantum simulation [1,2]. Ground state preparation can also be leveraged to solve combinatorial optimization problems, with a variety of applications including in routing and scheduling [3]. Although the task of ground state preparation is known to be hard, including for quantum computers, there is nonetheless significant interest in algorithms for quantum state preparation [4,5]. In particular, the growing availability of noisy, intermediate-scale quantum [6] devices has inspired immense interest in variational methods for preparing desired quantum states [7–10]. These variational quantum algorithms (VQAs) are heuristics that function by classically optimizing over a set of parameters that enter into a quantum circuit whose structure is typically fixed. The parameterized quantum circuit is executed on a quantum device and serves as an *Ansatz* to minimize a cost function, J , whose global minimum is achieved for the desired target state.

Even in the absence of noise, a variety of challenges are present in VQAs. On the quantum device, for example, one must select a quantum circuit *Ansatz* and associated initial state out of a formidably large design space. Meanwhile, the difficulty of the cost function minimization means that the challenges on the classical side can be even more significant [11], often involving the navigation of optimization

landscapes that contain barren plateaus and suboptimal local minima.

To overcome these challenges, approaches have been proposed that utilize feedback from qubit measurements to *adaptively* construct a quantum circuit to minimize J . The first such algorithm was the Adaptive Derivative-Assembled Problem-Tailored Variational Quantum Eigensolver (ADAPT-VQE) [12,13], which was also adapted to combinatorial optimization problems [14]. Instead of relying on a predefined *Ansatz*, ADAPT-VQE grows it in a layer-wise manner tailored to the problem, in tandem with classical optimization over the quantum circuit parameters.

Other methods have considered defining the structure of the *Ansatz* *a priori* and then performing layer-wise optimization to adaptively set the circuit parameter values [15–18]. Adaptive methods that do not require any classical optimization have also been developed, including the feedback-based algorithm for quantum optimization [19–21] and methods based on Riemannian gradient flows [22,23]. In methods

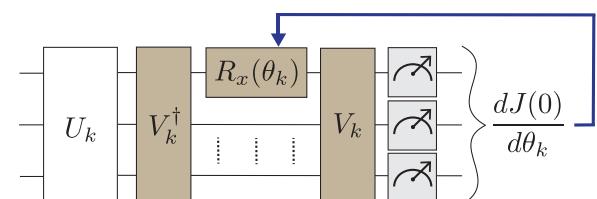


FIG. 1. Schematic representation of randomized adaptive quantum state preparation. Each adaptive step k involves first estimating the gradient $\frac{dJ(0)}{d\theta_k}$ at $\theta_k = 0$ of a cost function J with respect to a parameter θ_k , e.g., a rotation angle, whose corresponding direction is randomized through conjugation with a random unitary transformation V_k (brown). The parameter θ_k is then updated in the negative direction of the gradient (blue arrow), which moves the system closer to the target state, and the quantum circuit U_k is extended accordingly.

where each adaptive step is efficiently implementable, e.g., [12–20], and the approximate method in [22], the circuit growth can get stuck when the gradient vanishes. This means that these methods can face similar issues as conventional VQAs that are prone to converge to suboptimal solutions.

Here, we propose *randomized* adaptive quantum state preparation as a generic, adaptive quantum algorithm that minimizes up-front design choices, does not require classical optimization, and allows for preparing target quantum states from arbitrary (random) initial states. As a consequence of the latter point, the algorithm, or its key subroutine depicted in Fig. 1, can be readily combined with other state preparation methods to improve convergence and state preparation fidelities. We provide theoretical arguments and numerical evidence that substantiate our claims about convergence, explore different methods for achieving randomization in practice, and develop lower bounds on the expected change in J at each adaptive step. These bounds give guarantees for how much the cost function value changes when randomization is used, thereby allowing us to relate the efficiency of this randomized approach to the existence of barren plateaus [24]. We go on to discuss how this approach can be applied to cooling in open quantum systems and mixed state preparation in general.

II. RANDOMIZED ADAPTIVE QUANTUM ALGORITHMS

We consider minimizing cost functions of the form

$$J_k = \langle \psi_k | H_p | \psi_k \rangle, \quad (1)$$

by creating states $|\psi_k\rangle = U_k |\psi_0\rangle$ via an adaptively constructed quantum circuit U_k , $k = 0, 1, \dots$. The goal is to apply this circuit to a fixed initial state $|\psi_0\rangle$ to achieve $J_{k+1} \leq J_k$ in each adaptive step k . Here, H_p is a Hermitian operator whose ground state $|E_{\min}\rangle$, with corresponding eigenvalue E_{\min} , is taken to be the target state. We note that in general, knowledge of the initial and target states $|\psi_0\rangle$ and $|E_{\min}\rangle$ is not required. However, for the preparation of an arbitrary, known target state $|\psi_T\rangle$, H_p can be defined as $H_p = \mathbb{1} - |\psi_T\rangle \langle \psi_T|$.

We consider a quantum circuit that is adaptively created according to

$$U_{k+1} = e^{-i\theta_k H_k} U_k, \quad U_0 = \mathbb{1}, \quad (2)$$

where in each step k we move into the negative direction of the gradient of J with respect to θ_k by setting $\theta_k \leftarrow -\gamma \frac{dJ(0)}{d\theta_k}$, where $\frac{dJ(0)}{d\theta_k} = i \langle \psi_k | [H_k, H_p] | \psi_k \rangle$ is the gradient of $J(\theta_k) = \langle \psi_k | e^{i\theta_k H_k} H_p e^{-i\theta_k H_k} | \psi_k \rangle$ evaluated at $\theta_k = 0$. Alternatively, in situations where J_k can be estimated via repeated measurements of H_p (e.g., for $H_p = \mathbb{1} - |\psi_T\rangle \langle \psi_T|$ with known $|\psi_T\rangle$), the derivative $\frac{dJ(0)}{d\theta_k}$ can be estimated via a finite difference approximation, by estimating J_k for different perturbations of θ_k . For sufficiently small learning rates γ , this ensures that $J_{k+1} \leq J_k$. References [12,14,19,20,22] consider adaptive procedures similar to Eq. (2), and in cases where the circuit growth gets stuck in suboptimal solutions, i.e., when the gradient vanishes, the utility of incorporating randomness into the circuit structure to escape these suboptima has been observed numerically in [22].

In this work, we utilize randomness to overcome challenges associated with convergence through the introduction

of an *intrinsically randomized* framework for quantum state preparation in which the H_k 's are selected at random. In the following, we provide theoretical arguments that this randomization enables convergence from almost all initial states to arbitrary target states.

We first note that the cost function gradient can be expressed as

$$\frac{dJ(0)}{d\theta_k} = \langle \text{grad}J[U_k], iH_k \rangle, \quad (3)$$

where $\langle \cdot, \cdot \rangle$ denotes the Hilbert-Schmidt inner product and $\text{grad}J[U_k] = [|\psi_k\rangle \langle \psi_k|, H_p]$ is the (Riemannian) gradient (up to multiplication with U_k from the right) of $J[U_k] = \langle \psi_0 | U_k^\dagger H_p U_k | \psi_0 \rangle$ with respect to the unitary transformation U_k [23,25,26]. Both iH_k and $\text{grad}J[U_k]$ belong to the special unitary algebra $\mathfrak{su}(2^n)$ consisting of all traceless and anti-Hermitian $2^n \times 2^n$ matrices where n is the number of qubits. From this geometric perspective we can now deduce two different cases (i) and (ii) for when $\frac{dJ(0)}{d\theta_k}$ vanishes [27]. In case (i), the gradient vanishes when iH_k is orthogonal to $\text{grad}J[U_k] \neq 0$, while in case (ii) $\text{grad}J[U_k] = 0$, which happens when H_p commutes with $|\psi_k\rangle \langle \psi_k|$.

We first discuss case (i). Typically, no assumptions can be made on whether $\text{grad}J[U_k]$ moves into a lower dimensional subspace of $\mathfrak{su}(2^n)$ when growing the circuit. As such, the situation that iH_k has no overlap with $\text{grad}J[U_k]$ can occur, e.g., when iH_k is an element of a subspace over which $\text{grad}J[U_k]$ has no support. To overcome this issue, we propose to create each H_k at random. This can be achieved by conjugating a traceless Hermitian operator H by a Haar random unitary transformation V_k in each adaptive step, i.e., such that $H_k = V_k^\dagger H V_k$. Since H_k is created uniformly randomly according to the Haar measure, the probability that iH_k is orthogonal to $\text{grad}J[U_k]$ is zero. That is, for almost all iH_k , but a set of measure zero, case (i) does not occur. While Haar random unitaries are not efficiently implementable, below we discuss the efficient implementation via approximate unitary 2-designs [28].

We now focus on case (ii). For cost functions of the form (1), the set of critical points where $\text{grad}J[U_k]$ vanishes consists of global optima and saddle points only [29,30]. Under mild assumptions on the nature of the saddle points (strict saddles), relevant works from the classical machine learning and optimization literature have found that saddle points are avoided for almost all initial conditions [31–33]. We thus expect that randomized adaptive quantum state preparation will almost surely converge to the ground state of H_p for almost all initial states $|\psi_0\rangle$. We remark that the convergence result cannot hold for *all* initial states, as we immediately see that for eigenstates of H_p , $\text{grad}J[U_0] = 0$. However, the situation that $\text{grad}J[U_0] = 0$ can be avoided with probability one when the initial state is randomized too.

Each step k of randomized adaptive quantum state preparation can now be summarized as follows. First, the unitary transformation $V_k^\dagger e^{-i\theta_k H} V_k$, whose generator H is randomized through conjugation with a random unitary V_k , is applied to the state $|\psi_k\rangle$. Second, the gradient, $\frac{dJ(0)}{d\theta_k}$, is estimated, e.g., using the parameter shift rule [34–38]. Third, the parameter θ_k is updated in the negative direction of the gradient, i.e., $\theta_k \leftarrow -\gamma \frac{dJ(0)}{d\theta_k}$.

III. EFFICIENCY AND RELATION TO BARREN PLATEAUS

Randomized adaptive quantum state preparation is not expected to be efficient in general, as ground state preparation is QMA complete [5,39]. Here, we show that for a given problem, the efficiency can be related to the scaling of the gradient (3) with the system size, and therefore, to the existence of barren plateaus [24], i.e., exponentially flat regions in the optimization landscape where the variance of the gradient vanishes exponentially in the number of qubits n .

In Appendix A, we show that if we select $\gamma = 1/(4\|H_p\|_2)$ and assume $\|H_k\|_2 = 1$, where $\|\cdot\|_2$ denotes the spectral norm, then the cost function change $\Delta J_k = J_k - J_{k+1}$ is lower bounded by

$$\Delta J_k \geq \frac{1}{8\|H_p\|_2} \left(\frac{dJ(0)}{d\theta_k} \right)^2. \quad (4)$$

If we assume that it takes M steps to create the ground state up to an error ϵ , i.e., $J_M = E_{\min} + \epsilon$, we find from (4) that M is upper bounded by

$$M \leq \frac{C_\epsilon}{\min_{k \in [0, M]} \langle \text{grad}J[U_k], iH_k \rangle^2}, \quad (5)$$

where the constant in the numerator is given by $C_\epsilon = 8\|H_p\|_2[J_0 - (E_{\min} + \epsilon)]$. Thus, if $\langle \text{grad}J_k[U_k], iH_k \rangle^2$ does not vanish faster than $1/\text{poly}(n)$, the ground state of H_p can be prepared up to precision ϵ (in the corresponding eigenvalue) in polynomially many steps. We remark that the minimum in the denominator of Eq. (5) explicitly depends on M , and thereby on the random path taken. It is interesting to note that similar expressions are obtained in adiabatic state preparation [40], where the scaling of the adiabatic state preparation time T is determined by the smallest value of the spectral gap $\Delta(t)$, taken over all times, i.e., $\min_{t \in [0, T]} \Delta(t)$ [41].

We proceed by investigating the efficiency of different randomization strategies. At each step k , the expected cost function change $\mathbb{E}_{H_k} \Delta J_k$ is lower bounded by the variance $\mathbb{E}_{H_k} \langle \text{grad}J[U_k], iH_k \rangle^2$ of the gradient, up to the prefactor in (4), assuming that the expectation vanishes. This suggests that sampling V_k from unitary 2-designs suffices to obtain convergence to the ground state. This observation is further substantiated by the numerical simulations in Fig. 2, which considers the task of preparing the ground state of an Ising Hamiltonian. We specifically consider a model in which we map each spin to a vertex on a 3-regular graph, and couplings are present between spins whose corresponding vertices are connected by an edge. This is equivalent to solving the combinatorial optimization problem MaxCut on an unweighted, 3-regular graph [42]. We consider the approximation ratio $\alpha = J_k/E_{\min}$ as our figure of merit. Figure 2 compares results of randomized adaptive quantum state preparation when V_k is sampled at random from the Haar measure, with results when V_k is sampled from a unitary 2-design. The results that are obtained are nearly identical, with the difference shown in the inset.

The appearance of the variance in the bound (4) means that we additionally expect the efficiency and practical utility of this method to be closely related to the existence of barren plateaus. While barren plateaus pose a major challenge to the

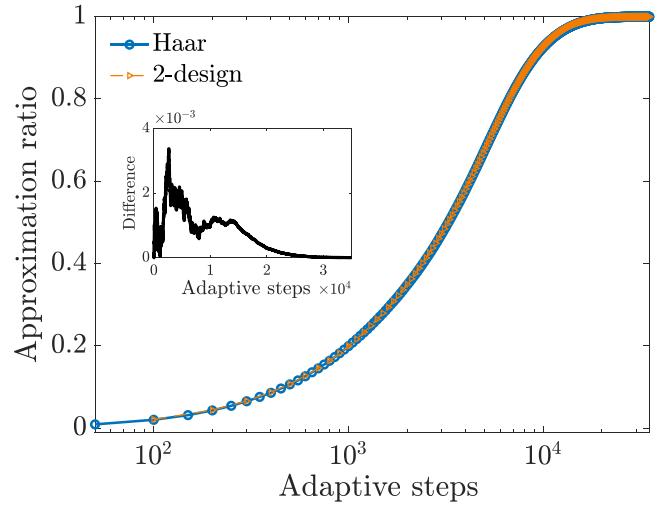


FIG. 2. Performance of randomized adaptive quantum state preparation when H_p is an Ising Hamiltonian with $n = 8$ spins. We map each spin to a vertex on a 3-regular graph with 8 vertices, and couple spins whose corresponding vertices are connected by an edge. Each data point corresponds to the average of the approximation ratio, α , taken over 100 different algorithm realizations and initial states. This is plotted as a function of the number of adaptive steps, M , on a logarithmic scale. In each step, the randomization of H_k is implemented by conjugating the Pauli operator $X_1 I_2 \cdots I_n$ with a Haar random unitary transformation (blue circles) and with a unitary transformation sampled from an approximate unitary 2-design (orange triangles), as depicted in Fig. 1, created using the sequence in [43] with $\ell = 1$. The two curves are nearly superimposed. Inset shows the difference between the two curves, computed as $|\alpha_{\text{Haar}} - \alpha_{\text{2-design}}|$, which never exceeds 4×10^{-3} .

scalability of VQAs [24,44,45], examples have been found where the variance of the gradient does not vanish faster than $1/\text{poly}(n)$ [46]. Leveraging these instances for efficient realizations of randomized adaptive quantum state preparation will be the subject of future studies.

We now consider lower bounds for $\mathbb{E}_{H_k} \Delta J_k$ to derive guarantees for how much J_k can be improved when randomization is used. When H_k is created through conjugation by a unitary V_k sampled from a unitary 2-design, we show in Appendix B that $\mathbb{E}_{V_k} \Delta J_k$ is lower bounded by

$$\mathbb{E}_{V_k} \Delta J_k \geq \frac{\text{Tr}\{H^2\}}{4\|H_p\|_2} \frac{\text{Var}_{\psi_k}(H_p)}{2^{2n} - 1}, \quad (6)$$

where $\text{Var}_{\psi_k}(H_p) = \langle \psi_k | H_p^2 | \psi_k \rangle - \langle \psi_k | H_p | \psi_k \rangle^2$ is the variance of H_p with respect to the state $|\psi_k\rangle$. Since $\text{Var}_{\psi_k}(H_p)$ is bounded from above by a constant that is independent of the system dimension, we see that the bound in (6) vanishes exponentially in n .

Another way of creating random H_k 's is by sampling uniformly from an operator pool \mathcal{A} whose size we denote by $|\mathcal{A}|$. While in this case, situation (i) can occur when an operator is selected that is orthogonal to $\text{grad}J_k[U_k]$, on average, we have that

$$\mathbb{E}_{H_k \in \mathcal{A}} \Delta J_k \geq \frac{1}{8\|H_p\|_2 |\mathcal{A}|} \sum_{H_k \in \mathcal{A}} \langle \text{grad}J[U_k], iH_k \rangle^2, \quad (7)$$

which suggests that situation (i) can be avoided on average for sufficiently large operator pools, e.g., when $\text{span}\{\mathcal{A}\} = \mathfrak{su}(2^n)$. Note that in this case, convergence to the ground state can also be obtained by simply continuing to sample from \mathcal{A} when (i) occurs. However, in this situation we expect an exponential runtime, as M scales as $1/|\mathcal{A}|$. For $|\mathcal{A}| = \text{poly}(n)$, on the other hand, convergence to the ground state is no longer guaranteed, as the adaptive procedure can get stuck in suboptima where the gradient vanishes, due to (i). Furthermore, while the situation that $|\mathcal{A}| = \text{poly}(n)$ implies polynomial scaling of the denominator of Eq. (7), it does not necessarily imply that randomized adaptive quantum state preparation would be efficient in this setting, as it does not imply polynomial scaling of the numerator.

In Fig. 3, we numerically investigate this trade-off and study the convergence of randomized adaptive quantum state preparation for the problem of preparing the ground state of an Ising Hamiltonian with equal couplings present between all n spins, which is equivalent to solving the combinatorial optimization problem MaxCut on an unweighted, complete graph [42]. In Fig. 3(a), we plot the approximation ratio, α , as a function of the number of adaptive steps M for the different randomization strategies described above. In Fig. 3(b), we consider fixed $M = 35\,000$ and investigate α as a function of the pool size $|\mathcal{A}|$. The pool size is increased by adding successively higher weight Pauli operators to \mathcal{A} , until the full pool used in Fig. 3(a) containing all $2^{2n} - 1$ terms is formed. We observe that the curves in Fig. 3(a) are almost identical, suggesting that the three different randomization strategies converge in the same manner to the ground state for which $\alpha = 1$. Furthermore, the curves in Fig. 3(b) suggest that a full Pauli operator pool is not needed to obtain convergence to the ground state. However, the inset semilog plots in Figs. 3(a) and 3(b) do suggest an exponential scaling of the number of adaptive steps, M , and the number of operators in the pool, $|\mathcal{A}|$, with respect to n .

IV. MIXED STATES

Thus far, we have discussed the preparation of pure (ground) states from initial pure states. Here, we explore generalizations to preparing a target pure state from an initially mixed state, and vice versa. Since it is not possible to create pure quantum states from mixed quantum states in a closed quantum system through unitary transformations, we consider an extended, “dilated” space by coupling a set of system qubits, S , to a set of auxiliary qubits, A . We assume that the combined system is initially in a separable state $\rho_0^{SA} = \rho_0^S \otimes \rho_0^A$, where we denote the initial states of S and A by ρ_0^S and ρ_0^A , respectively. We then consider growing a quantum circuit over the full composite system according to (2), in order to adaptively create the state $\rho_k^{SA} = U_k \rho_0^{SA} U_k^\dagger$. The state ρ_k^S at the k th adaptive step is then obtained by tracing over the degrees of freedom of the auxiliary qubits in subsystem A , i.e., $\rho_k^S = \text{Tr}_A\{\rho_k^{SA}\}$.

In this setting, we now consider the task of preparing the system qubits in a target pure state $|\psi_T\rangle$. The cost function (1) becomes $J_k = 1 - \langle\psi_T|\rho_k^S|\psi_T\rangle$, and an adaptive change is described by $J(\theta_k) = 1 - \langle\psi_T|\text{Tr}_A\{e^{i\theta_k H_k} \rho_k^{SA} e^{-i\theta_k H_k}\}|\psi_T\rangle$.

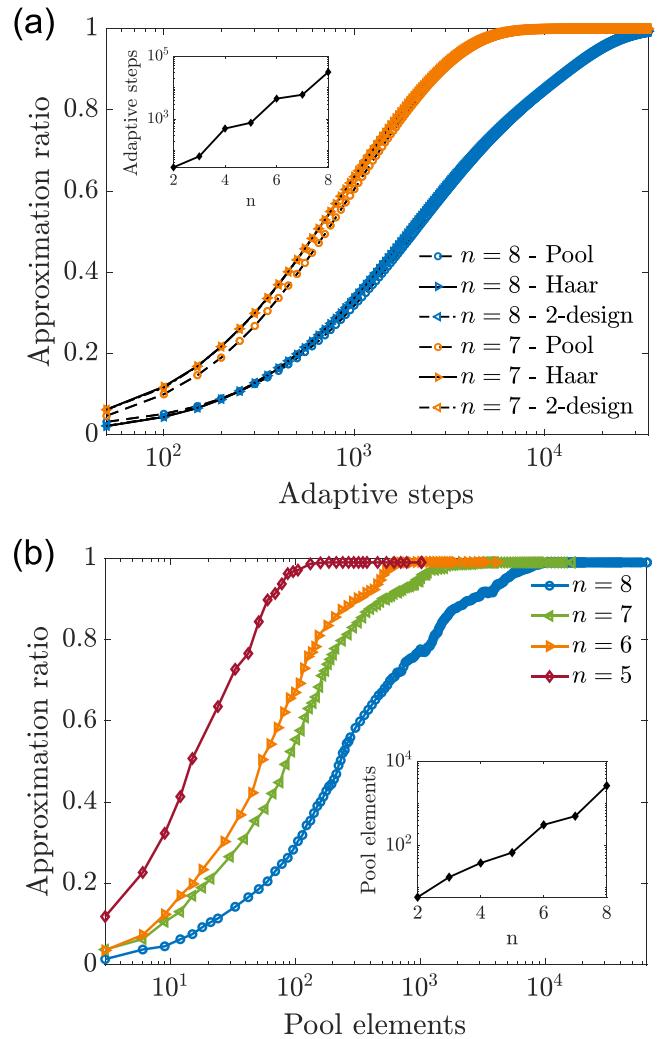


FIG. 3. Performance of randomized adaptive quantum state preparation when H_p is an Ising Hamiltonian with all-to-all couplings between n spins. Each data point corresponds to the average taken over 100 different algorithm realizations and initial states. In (a), the approximation ratio, α , is plotted as a function of the number of adaptive steps, M , shown on a logarithmic scale. In each step, the randomization of H_k is implemented by conjugating the Pauli operator $X_1 I_2 \cdots I_n$ with a Haar random unitary transformation (right triangles) and with a unitary transformation sampled from an approximate unitary 2-design (left triangles), as depicted in Fig. 1, created using the sequence in [43] with $\ell = 1$, and also by sampling H_k uniformly from an operator pool \mathcal{A} containing all $2^{2n} - 1$ Pauli operators (circles). In (b), α is plotted as a function of the number of pool elements, $|\mathcal{A}|$, shown on a logarithmic scale, for fixed $M = 35\,000$. Insets show semilog plots of the scaling with respect to n for both cases to achieve $\alpha > 0.99$ in (a) and $\alpha > 0.9$ in (b).

This yields the gradient

$$\frac{dJ(0)}{d\theta_k} = -\langle [\rho_k^{SA}, |\psi_T\rangle \langle\psi_T| \otimes \mathbb{1}_A], iH_k \rangle, \quad (8)$$

where we have made use of the fact that the cost function above can be obtained from (1) by setting $H_p = \mathbb{1} - |\psi_T\rangle \langle\psi_T| \otimes \mathbb{1}_A$ where $\mathbb{1}_A$ denotes the identity operator on subsystem A . Decreasing J_k in each step can now be achieved

by moving into the negative direction of the gradient given by Eq. (8).

If the auxiliary qubits are initially in a pure state, then due to the Stinespring dilation [47], there exists a unitary transformation over the composite system that allows for creating every state for the system qubits in subsystem S as long as subsystem A has dimension at most d_S^2 , where d_S is the dimension of subsystem S . For uniformly randomized H_k , the gradient can only vanish (with probability 1) when $\text{grad}J[U_k] = [\rho_k^{SA}, |\psi_T\rangle\langle\psi_T| \otimes \mathbb{1}_A] = 0$, i.e., at critical points that are given by saddle points and global optima, as in the closed system case [30]. Thus, when the initial system state ρ_0^S does not commute with the target system state $|\psi_T\rangle\langle\psi_T|$, we expect to obtain convergence almost surely to the target state. Although the fully mixed state $\rho_0^S = 1/d_A$ trivially commutes with every target state, if we additionally apply a random unitary transformation to ρ_0^{SA} , we can almost surely obtain convergence to a generic pure state. Thus, randomized adaptive quantum state preparation allows for “cooling” the system from a state of infinite temperature, i.e., the fully mixed state, to a pure state of zero temperature by adaptively “dumping” entropy into the auxiliary system.

V. CONCLUSION

We have introduced an algorithm for preparing quantum states that has favorable convergence properties and is applicable to almost all initial states. Knowledge of the initial and target quantum states is not required. The algorithm leverages randomization as the primary innovation, and operates by minimizing a cost function through an adaptively constructed quantum circuit. Each adaptive update step is informed by gradient measurements in which the associated tangent space directions are randomized. We have presented lower bounds on the average improvement that is obtained in each step, and have numerically studied the behavior of the algorithm for different randomization methods. We additionally discussed a generalization to mixed states that could be leveraged for thermal state preparation, an application area which is currently receiving significant interest [48–54].

A trade-off is that, on one hand, the consideration of a uniformly random tangent space direction in each step allows for convergence to the target state, while on the other hand, the gradients may become exponentially small in the system size, thereby causing the convergence time to diverge and making the algorithm impractical for large-scale problems. This should not be surprising, as selecting directions at random is far from being optimal. The largest gradient, and thereby the largest (guaranteed) cost function change, is obtained when $iH_k = \text{grad}J[U_k]$. This situation corresponds to the full gradient flow, which is, in general, not efficiently implementable. It is an interesting question how gradient flows can be efficiently approximated [22] while maintaining convergence. To that end, the algorithm we present could be modified to project into random subspaces, rather than into a single random direction.

Furthermore, we emphasize that the algorithm possesses significant design flexibility that can be harnessed to increase its practicality, for example, by developing application-specific adaptations that tailor the randomness to the problem instance, e.g., by taking into account the symmetries of the

target state [55–58], and by studying how much the rate of convergence can be improved by incorporating classical optimization in different ways [12]. More fundamentally, we view the randomized adaptive circuit update, which can be applied to arbitrary input states and is depicted in Fig. 1, as a *subroutine* that could be incorporated into or appended onto other algorithms [40,59] in a flexible way, e.g., those utilizing problem-specific *Ansätze* that are not random, to improve convergence and state preparation fidelities.

We conclude by noting that we have focused this work on randomized adaptive quantum state preparation in the circuit model of quantum computing. However, randomness and adaptive constructions can also be created outside the circuit model, e.g., through random fields [60,61]. Extensions to this latter setting could open up new approaches for state preparation in both closed and open quantum systems that can be driven by an applied field, including qubit systems, analog quantum simulators, molecules, and materials.

ACKNOWLEDGMENTS

We gratefully acknowledge feedback from Mohan Sarovar and Andrew Baczewski on this work. We also would like to thank Thomas Schulte-Herbrüggen, Gunther Dirr, and Emanuel Malvetti for fruitful discussions. A.B.M. acknowledges support from Sandia National Laboratories’ Laboratory Directed Research and Development Program under the Truman Fellowship. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy’s National Nuclear Security Administration under Contract No. DE-NA0003525. This article has been authored by an employee of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The employee owns all right, title and interest in and to the article and is solely responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan <https://www.energy.gov/downloads/doe-public-access-plan>. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. C.A. and S.E. acknowledge support from the National Science Foundation (Grant No. 2231328).

APPENDIX A: LOWER BOUNDING THE COST FUNCTION CHANGE

We consider

$$J(x) = \langle \psi_k | e^{ixH_k} H_p e^{-ixH_k} | \psi_k \rangle, \quad (\text{A1})$$

and note that $J(x=0) = \langle \psi_k | H_p | \psi_k \rangle = J_k$ and $J(x = -\gamma \frac{dJ(0)}{d\theta}) = J_{k+1}$.

1. Recalling results from classical optimization

In order to lower bound the cost function change $\Delta J_k = J_k - J_{k+1}$, we first recall some standard tools from classical optimization, particularly the gradient descent algorithm. For a continuously differentiable cost function $J(x)$, $x \in \mathbb{R}$, with Lipschitz continuous gradient $\frac{\partial J(x)}{\partial x}$, a direct application of Taylor's theorem [62] yields

$$J(y) \leq J(x) + (y-x) \frac{\partial J(x)}{\partial x} + \frac{L_f^2}{2} (y-x)^2, \quad (\text{A2})$$

where L_f is the Lipschitz constant satisfying

$$\frac{\left| \frac{\partial J(x)}{\partial x} - \frac{\partial J(y)}{\partial y} \right|}{|x-y|} \leq L_f. \quad (\text{A3})$$

Setting $x=0$ and $y = -\gamma \frac{dJ(x)}{d\theta}|_{x=0}$ in (A2) and picking $\gamma = \frac{1}{L_f}$ gives

$$J_k - J_{k+1} \geq \frac{1}{2L_f} \left(\frac{\partial J(0)}{\partial x} \right)^2, \quad (\text{A4})$$

which is a well known result for gradient algorithms [62]. Thus, in order to lower bound the cost function change ΔJ_k , it remains to determine the Lipschitz constant L_f for our setting.

2. Determination of the Lipschitz constant

We first note that $\frac{\partial J(x)}{\partial x} = \langle \psi_k | e^{ixH_k} [iH_k, H_p] e^{-ixH_k} | \psi_k \rangle \equiv \langle \psi(x) | A | \psi(x) \rangle$, where we defined $|\psi(x)\rangle = e^{-ixH_k} |\psi_k\rangle$ and $A = [iH_k, H_p]$. We can then upper bound $\left| \frac{\partial J(x)}{\partial x} - \frac{\partial J(y)}{\partial y} \right|$ by

$$\left| \frac{\partial J(x)}{\partial x} - \frac{\partial J(y)}{\partial y} \right| = |\langle \psi(x) | A | \psi(x) \rangle - \langle \psi(y) | A | \psi(y) \rangle| \quad (\text{A5})$$

$$= \frac{1}{2} |[\langle \psi(x) | + \langle \psi(y) |] A [|\psi(x)\rangle - |\psi(y)\rangle] + [\langle \psi(x) | - \langle \psi(y) |] A [|\psi(x)\rangle + |\psi(y)\rangle]| \quad (\text{A6})$$

$$\leq 2 \|A[|\psi(x)\rangle - |\psi(y)\rangle]\|_2 \leq 2 \|A\|_2 \| |\psi(x)\rangle - |\psi(y)\rangle\|_2 \quad (\text{A7})$$

$$= 2 \|A\|_2 \| (U_x - U_y) |\psi_k\rangle \|_2 = 2 \|A\|_2 \| (U_y^\dagger U_x - \mathbb{1}) |\psi_k\rangle \|_2, \quad (\text{A8})$$

where we used unitary invariance of the vector norm, $U_x(s) = e^{-i s H_k}$ and $U_y(s) = e^{-i y s H_k}$, using the shorthand notation $U_x(1) = U_x$, $U_y(1) = U_y$, and $\|A\|_2$ and $\| |\psi\rangle\|_2$ denote the spectral norm of A and the Euclidean vector norm of $|\psi\rangle$, respectively. Since

$$U_y^\dagger U_x - \mathbb{1} = i \int_0^1 ds U_y^\dagger(s) (y-x) H_k U_x(s), \quad (\text{A9})$$

we have

$$\| (U_y^\dagger U_x - \mathbb{1}) |\psi_k\rangle \|_2 \leq \int_0^1 ds \| [U_y^\dagger(s) (y-x) H_k U_x(s)] |\psi_k\rangle \|_2 \quad (\text{A10})$$

$$= |x-y| \int_0^1 ds \|H_k\|_2 \| |\psi(x)\rangle\|_2 \\ = |x-y| \|H_k\|_2 = |x-y| \|H\|_2. \quad (\text{A11})$$

In the last step we used unitary invariance of the spectral norm, assuming that $H_k = V_k^\dagger H V_k$, where V_k is unitary and H is a Hermitian matrix. Since

$$\|A\|_2 = \| [H_k, H_p] \|_2 \leq 2 \|H_k\|_2 \|H_p\|_2 = 2 \|H_p\|_2 \|H\|_2, \quad (\text{A12})$$

we arrive at

$$\frac{\left| \frac{\partial J}{\partial x} - \frac{\partial J}{\partial y} \right|}{|x-y|} \leq 4 \|H_p\|_2 \|H\|_2^2, \quad (\text{A13})$$

from which we conclude that the Lipschitz constant is given by $L_f = 4 \|H_p\|_2 \|H\|_2^2$. Further assuming $\|H\|_2 = 1$, we arrive at the lower bound (4) given in the paper.

APPENDIX B: EXPECTED COST FUNCTION CHANGE WHEN SAMPLING FROM UNITARY 2-DESIGNS

We first note that the expected value of the gradient $\mathbb{E}_{V_k} \langle \text{grad}J[U_k], iH_k \rangle$ is given by

$$\mathbb{E}_{V_k} \langle \text{grad}J[U_k], iV_k^\dagger H V_k \rangle = \langle \text{grad}J[U_k], i\mathbb{1} \rangle \frac{\text{Tr}\{H\}}{d} = 0, \quad (\text{B1})$$

when V_k is sampled according to the Haar measure on $\text{SU}(d)$ or from a unitary 2-design. Consequently, in this case $\mathbb{E}_{V_k} \langle \text{grad}J[U_k], iH_k \rangle^2$ is equal to the variance.

In order to lower bound the expected cost function change $\mathbb{E}_{V_k} \Delta J_k$ when V_k is drawn according to the Haar measure on $\text{SU}(d)$ or sampled from a unitary 2-design, we first note that from the lower bound (4) we have $\mathbb{E}_{V_k} \Delta J_k \geq \frac{1}{8\|H_p\|_2} \mathbb{E}_{V_k} \langle \text{grad}J[U_k], iH_k \rangle^2 = \frac{1}{8\|H_p\|_2} \mathbb{E}_{V_k} \langle \psi_k | [iV_k^\dagger H V_k, H_p] |\psi_k \rangle^2$. To calculate the expectation on the right-hand side, we make use of the relation [46]

$$\begin{aligned} & \int_{V \in \text{SU}(d)} dV \text{Tr}\{V^\dagger B V C V^\dagger D V\} \\ &= \frac{1}{d^2 - 1} (\text{Tr}\{A\} \text{Tr}\{C\} \text{Tr}\{BD\} + \text{Tr}\{AC\} \text{Tr}\{B\} \text{Tr}\{D\}) \\ & \quad - \frac{1}{d(d^2 - 1)} (\text{Tr}\{AC\} \text{Tr}\{BD\} + \text{Tr}\{A\} \text{Tr}\{B\} \text{Tr}\{C\} \text{Tr}\{D\}). \end{aligned} \quad (\text{B2})$$

We find

$$\begin{aligned} \mathbb{E}_{V_k} \langle \psi_k | [iV_k^\dagger H V_k, H_p] | \psi_k \rangle^2 \\ = \mathbb{E}_{V_k} \text{Tr} \{ [iV_k^\dagger H V_k, H_p] | \psi_k \rangle \langle \psi_k | [iV_k^\dagger H V_k, H_p] | \psi_k \rangle \langle \psi_k | \} \\ = -2 \text{Tr} \{ H^2 \} \left(\frac{1}{d^2 - 1} \langle \psi_k | H_p | \psi_k \rangle^2 \right. \end{aligned}$$

$$\begin{aligned} & - \frac{1}{d(d^2 - 1)} \langle \psi_k | H_p | \psi_k \rangle^2 \\ & \times - \frac{1}{d^2 - 1} \langle \psi_k | H_p^2 | \psi_k \rangle + \frac{1}{d(d^2 - 1)} \langle \psi_k | H_p | \psi_k \rangle^2 \Big) \\ & = \frac{2 \text{Tr} \{ H^2 \}}{d^2 - 1} \left(\langle \psi_k | H_p^2 | \psi_k \rangle - \langle \psi_k | H_p | \psi_k \rangle^2 \right), \end{aligned} \quad (B3)$$

which yields the desired result given in (6) in the paper.

[1] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, *Rev. Mod. Phys.* **92**, 015003 (2020).

[2] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. Sawaya *et al.*, Quantum chemistry in the age of quantum computing, *Chem. Rev.* **119**, 10856 (2019).

[3] C. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity* (Dover Publications, Mineola, NY, 1998).

[4] A. Kitaev, A. Shen, and M. N. Vyalyi, *Classical and Quantum Computation* (American Mathematical Society, Providence, Rhode Island, 2002).

[5] J. Kempe, A. Kitaev, and O. Regev, The complexity of the local Hamiltonian problem, *SIAM J. Comput.* **35**, 1070 (2006).

[6] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).

[7] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).

[8] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio *et al.*, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).

[9] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).

[10] A. B. Magann, C. Arenz, M. D. Grace, T.-S. Ho, R. L. Kosut, J. R. McClean, H. A. Rabitz, and M. Sarovar, From pulses to circuits and back again: A quantum optimal control perspective on variational quantum algorithms, *PRX Quantum* **2**, 010101 (2021).

[11] L. Bittel and M. Kliesch, Training Variational Quantum Algorithms is NP-hard, *Phys. Rev. Lett.* **127**, 120502 (2021).

[12] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, An adaptive variational algorithm for exact molecular simulations on a quantum computer, *Nat. Commun.* **10**, 3007 (2019).

[13] H. L. Tang, V. O. Shkolnikov, G. S. Barron, H. R. Grimsley, N. J. Mayhall, E. Barnes, and S. E. Economou, Qubit-ADAPT-VQE: An adaptive algorithm for constructing hardware-efficient ansätze on a quantum processor, *PRX Quantum* **2**, 020310 (2021).

[14] L. Zhu, H. L. Tang, G. S. Barron, F. A. Calderon-Vargas, N. J. Mayhall, E. Barnes, and S. E. Economou, Adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer, *Phys. Rev. Res.* **4**, 033029 (2022).

[15] J. Carolan, M. Mohseni, J. P. Olson, M. Prabhu, C. Chen, D. Bunandar, M. Y. Niu, N. C. Harris, F. N. Wong, M. Hochberg *et al.*, Variational quantum unsampling on a quantum photonic processor, *Nat. Phys.* **16**, 322 (2020).

[16] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, and M. Leib, Layerwise learning for quantum neural networks, *Quantum Mach. Intell.* **3**, 5 (2021).

[17] E. Campos, A. Nasrallah, and J. Biamonte, Abrupt transitions in variational quantum circuit training, *Phys. Rev. A* **103**, 032607 (2021).

[18] E. Campos, D. Rabinovich, V. Akshay, and J. Biamonte, Training saturation in layerwise quantum approximate optimization, *Phys. Rev. A* **104**, L030401 (2021).

[19] A. B. Magann, K. M. Rudinger, M. D. Grace, and M. Sarovar, Feedback-Based Quantum Optimization, *Phys. Rev. Lett.* **129**, 250502 (2022).

[20] A. B. Magann, K. M. Rudinger, M. D. Grace, and M. Sarovar, Lyapunov-control-inspired strategies for quantum combinatorial optimization, *Phys. Rev. A* **106**, 062414 (2022).

[21] J. B. Larsen, M. D. Grace, A. D. Baczewski, and A. B. Magann, Feedback-based quantum algorithm for ground state preparation of the Fermi-Hubbard model, [arXiv:2303.02917](https://arxiv.org/abs/2303.02917).

[22] R. Wiersema and N. Killoran, Optimizing quantum circuits with Riemannian gradient flow, *Phys. Rev. A* **107**, 062421 (2023).

[23] T. Schulte-Herbrüggen, S. J. Glaser, G. Dirr, and U. Helmke, Gradient flows for optimization in quantum information and quantum dynamics: Foundations and applications, *Rev. Math. Phys.* **22**, 597 (2010).

[24] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).

[25] U. Helmke and J. B. Moore, *Optimization and Dynamical Systems* (Springer Science & Business Media, New York, 2012).

[26] P.-A. Absil, R. Mahony, and R. Sepulchre, Optimization algorithms on matrix manifolds, in *Optimization Algorithms on Matrix Manifolds* (Princeton University Press, Princeton, NJ, 2009).

[27] We neglect the trivial third case that $iH_k = 0$.

[28] C. Dankert, R. Cleve, J. Emerson, and E. Livine, Exact and approximate unitary 2-designs and their application to fidelity estimation, *Phys. Rev. A* **80**, 012304 (2009).

[29] R. W. Brockett, Least squares matching problems, *Lin. Alg. Appl.* **122-124**, 761 (1989).

[30] R. Wu, H. Rabitz, and M. Hsieh, Characterization of the critical submanifolds in quantum ensemble control landscapes, *J. Phys. A* **41**, 015006 (2007).

[31] R. Ge, F. Huang, C. Jin, and Y. Yuan, Escaping from saddle points—online stochastic gradient for tensor decomposition, in *Proceedings of the 28th Conference on Learning Theory*, edited

by P. Grünwald, E. Hazan, and S. Kale, *Proceedings of Machine Learning Research* Vol. 40 (PMLR, 2015), pp. 797–842.

[32] C. Jin, R. Ge, P. Netrapalli, S. M. Kakade, and M. I. Jordan, How to escape saddle points efficiently, in *Proceedings of the 34th International Conference on Machine Learning*, edited by D. Precup and Y. W. Teh, *Proceedings of Machine Learning Research* Vol. 70 (PMLR, 2017), pp. 1724–1732.

[33] J. D. Lee, I. Panageas, G. Piliouras, M. Simchowitz, M. I. Jordan, and B. Recht, First-order methods almost always avoid strict saddle points, *Math. Program.* **176**, 311 (2019).

[34] J. Li, X. Yang, X. Peng, and C.-P. Sun, Hybrid Quantum-Classical Approach to Quantum Optimal Control, *Phys. Rev. Lett.* **118**, 150503 (2017).

[35] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Quantum circuit learning, *Phys. Rev. A* **98**, 032309 (2018).

[36] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, Evaluating analytic gradients on quantum hardware, *Phys. Rev. A* **99**, 032331 (2019).

[37] D. Wierichs, J. Izaac, C. Wang, and C. Y.-Y. Lin, General parameter-shift rules for quantum gradients, *Quantum* **6**, 677 (2022).

[38] O. Kyriienko and V. E. Elfving, Generalized quantum circuit differentiation rules, *Phys. Rev. A* **104**, 052417 (2021).

[39] A. Y. Kitaev, A. Shen, and M. N. Vyalyi, *Classical and Quantum Computation* (American Mathematical Society, Providence, Rhode Island, 2002).

[40] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, [arXiv:quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106).

[41] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Rev. Mod. Phys.* **90**, 015002 (2018).

[42] F. Barahona, M. Grötschel, M. Jünger, and G. Reinelt, An application of combinatorial optimization to statistical physics and circuit layout design, *Oper. Res.* **36**, 493 (1988).

[43] Y. Nakata, C. Hirche, C. Morgan, and A. Winter, Unitary 2-designs from random X- and Z-diagonal unitaries, *J. Math. Phys.* **58**, 052203 (2017).

[44] C. Ortiz Marrero, M. Kieferová, and N. Wiebe, Entanglement-induced barren plateaus, *PRX Quantum* **2**, 040316 (2021).

[45] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, Noise-induced barren plateaus in variational quantum algorithms, *Nat. Commun.* **12**, 6961 (2021).

[46] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, *Nat. Commun.* **12**, 1791 (2021).

[47] M. A. Nielsen and I. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, UK, 2002).

[48] J. Wu and T. H. Hsieh, Variational Thermal Quantum Simulation via Thermofield Double States, *Phys. Rev. Lett.* **123**, 220502 (2019).

[49] J. E. Moussa, Low-depth quantum metropolis algorithm, [arXiv:1903.01451](https://arxiv.org/abs/1903.01451).

[50] M. Motta, C. Sun, A. T. K. Tan, M. J. O’Rourke, E. Ye, A. J. Minnich, F. G. S. L. Brandão, and G. K.-L. Chan, Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution, *Nat. Phys.* **16**, 205 (2020).

[51] Y. Wang, G. Li, and X. Wang, Variational Quantum Gibbs State Preparation with a Truncated Taylor Series, *Phys. Rev. Appl.* **16**, 054035 (2021).

[52] A. Warren, L. Zhu, N. J. Mayhall, E. Barnes, and S. E. Economou, Adaptive variational algorithms for quantum Gibbs state preparation, [arXiv:2203.12757](https://arxiv.org/abs/2203.12757).

[53] A. N. Chowdhury, G. H. Low, and N. Wiebe, A variational quantum algorithm for preparing quantum Gibbs states, [arXiv:2002.00055](https://arxiv.org/abs/2002.00055).

[54] T. J. Sewell, C. D. White, and B. Swingle, Thermal multi-scale entanglement renormalization ansatz for variational Gibbs state preparation, [arXiv:2210.16419](https://arxiv.org/abs/2210.16419).

[55] B. T. Gard, L. Zhu, G. S. Barron, N. J. Mayhall, S. E. Economou, and E. Barnes, Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm, *npj Quantum Inf.* **6**, 10 (2020).

[56] V. O. Shkolnikov, N. J. Mayhall, S. E. Economou, and E. Barnes, Avoiding symmetry roadblocks and minimizing the measurement overhead of adaptive variational quantum eigensolvers *Quantum* **7**, 1040 (2023).

[57] T.-C. Yen, R. A. Lang, and A. F. Izmaylov, Exact and approximate symmetry projectors for the electronic structure problem on a quantum computer, *J. Chem. Phys.* **151**, 164111 (2019).

[58] G.-L. R. Anselmetti, D. Wierichs, C. Gogolin, and R. M. Parrish, Local, expressive, quantum-number-preserving VQE ansätze for fermionic systems, *New J. Phys.* **23**, 113010 (2021).

[59] D. Poulin and P. Wocjan, Preparing Ground States of Quantum Many-Body Systems on a Quantum Computer, *Phys. Rev. Lett.* **102**, 130503 (2009).

[60] L. Banchi, D. Burgarth, and M. J. Kastoryano, Driven Quantum Dynamics: Will it Blend? *Phys. Rev. X* **7**, 041015 (2017).

[61] P. Yang, M. Yu, R. Betzholz, C. Arenz, and J. Cai, Complete Quantum-State Tomography with a Local Random Field, *Phys. Rev. Lett.* **124**, 010405 (2020).

[62] A. Beck, *First-Order Methods in Optimization* (SIAM, Philadelphia, 2017).