

Landscape approximation of low-energy solutions to binary optimization problems

Benjamin Y. L. Tan^{1,*}, Beng Yee Gan¹, Daniel Leykam^{1,†} and Dimitris G. Angelakis^{1,2,3,‡}

¹Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, Singapore 117543, Singapore

²School of Electrical and Computer Engineering, Technical University of Crete, Chania 73100, Greece

³AngelQ Quantum Computing, 531A Upper Cross Street, 04-95 Hong Lim Complex, Singapore 051531, Singapore



(Received 13 October 2023; accepted 18 December 2023; published 26 January 2024)

We show how the localization landscape, originally introduced to bound low-energy eigenstates of disordered wave media and many-body quantum systems, can form the basis for hardware-efficient quantum algorithms for solving binary optimization problems. Many binary optimization problems can be cast as finding low-energy eigenstates of Ising Hamiltonians. First, we apply specific perturbations to the Ising Hamiltonian such that the low-energy modes are bounded by the localization landscape. Next, we demonstrate how a variational method can be used to prepare and sample from the peaks of the localization landscape. Numerical simulations of problems of up to ten binary variables show that the localization landscape-based sampling can outperform quantum approximate optimization algorithm (QAOA) circuits of similar depth, as measured in terms of the probability of sampling the exact ground state.

DOI: [10.1103/PhysRevA.109.012433](https://doi.org/10.1103/PhysRevA.109.012433)

I. INTRODUCTION

Finding optimal solutions to quadratic unconstrained binary optimization (QUBO) problems is one proposed near-term application of quantum computers [1]. Solving large-scale QUBO problems has importance in scheduling and allocation tasks [2–4] and machine learning [5–9], amongst others [10–13]. The search for these optimal solutions is generally difficult, as QUBO problems are NP-hard [14–16]. However, in many cases obtaining approximate solutions close to the optimal can be sufficient. This is especially true within the context of industry applications where a higher quality solution, despite being suboptimal, may still result in significant cost savings [17–19].

Commonly employed techniques for solving QUBO problems using quantum computers are typically based on mapping the QUBO problem at hand to an Ising Hamiltonian, solving the problem by finding the ground state of the Ising Hamiltonian. Quantum algorithms to find the ground state include quantum annealing [20–22], variational problem-specific algorithms such as the quantum approximate optimization algorithm (QAOA) and its generalizations [23,24], variational quantum eigensolvers [25–27], and quantum assisted methods [1,28,29].

Methods such as quantum annealing and QAOA have shown provable convergence to the exact ground state in the limit of infinite annealing time and circuit depth. In many applications it is necessary to obtain a solution in a finite time, in which case these methods will ideally give a mixture of low-lying states. Hyperparameters such as the annealing

schedule, mixer Hamiltonian, number of QAOA steps, and sampling frequency thresholds for QAOA can affect the quality of the solutions obtained [30].

Here we consider a different approach. Instead of an exact method that, when run on a finite-sized circuit, gives approximate solutions whose quality is difficult to predict, we consider a scheme to sample from low-energy solutions with well-defined bounds, to solve the QUBO problem approximately using shallower-depth circuits. Our approach is inspired by the “localization landscape” used to study the Anderson localization of low-energy modes of disordered systems.

Anderson localization is the phenomenon where eigenfunction solutions to the Schrödinger equation with disordered potentials are confined due to wave interference [31]. Finding the locations where these quantum states localize typically requires solving the eigenvalue problem, as there is often seemingly little correlation between the potentials and the subregions where the peaks of these eigenfunctions can be found. Efforts into identifying these regions of localization resulted in the localization landscape function [32].

The localization landscape is a function that places a tight bound on the subregions where low-energy states tend to lie. The inverse of the landscape function serves as an effective potential that can be used to predict areas of confinement for low-energy eigenstates by identifying valleys within this effective potential. Since its introduction, efforts have gone into using the localization landscape to obtain the integrated density of states, thereby giving an estimate for the energies of the lowest eigenstates for the one-dimensional (1D) tight-binding model [33,34]. The original localization landscape function loses its accuracy when attempting to accurately identify localized regions of higher-energy eigenstates, motivating the development of related landscape functions such as the \mathcal{L}^2 landscape [35]. The \mathcal{L}^2 landscape is able to provide

*b.tan@u.nus.edu

†daniel.leykam@gmail.com

‡dimitris.angelakis@gmail.com

a tight bound on the localization of midspectrum eigenstates, and can be efficiently computed with a stochastic procedure using sparse matrix methods [36].

Remarkably, landscape functions can be generalized beyond low-dimensional disordered systems to more general families of real symmetric matrices (M matrices) [37]. The localization landscape theory has also been applied to many-body quantum systems [38], extending many of its well-known properties to Hamiltonians describing interacting spins, enabling the identification of regions of Hilbert space where the low-energy many-body eigenstates localize. Qualitative changes in the shape of the landscape, e.g., quantified using methods such as persistent homology, can be used as indicators of phase transitions in many-body quantum systems [39].

In this paper, we present a method of using the localization landscape to prepare a quantum state from which low-energy solutions to QUBO problems can be sampled with higher probabilities. We describe how this quantum state can be prepared on a near term quantum device, and demonstrate our methods for two problem instances—a nondegenerate randomly generated QUBO, and a degenerate MaxCut problem [2,40]. Existing efforts have explored the use of operators, such as the inverse of the Hamiltonian, to produce a quantum state with similar properties where eigenstates can be sampled with probabilities inversely related to their energy levels [41].

The outline of this paper is as follows: Sec. II reviews the localization landscape and its application to Anderson localization and many-body localization. Section III presents the mapping of QUBO problems to Ising Hamiltonians, showing how the Ising Hamiltonian can be perturbed such that its low-energy eigenstates are bounded by the localization landscape and proposing a heuristic using shallow variational circuits for sampling from this landscape suitable for noisy intermediate-scale quantum (NISQ) devices. Section IV presents numerical simulations showing how the method can be used to sample low-energy solutions with higher probability than shallow QAOA circuits. We analyze the effect of the two hyperparameters of the method (the energy offset and coupling strength) in Sec. V before concluding with Sec. VI.

II. LOCALIZATION LANDSCAPE

Given a disordered Hamiltonian \hat{H} , finding the regions where eigenstates localize typically requires solving the eigenvalue equation. However, Ref. [32] introduced a function called the *localization landscape*, u , that is able to predict these subregions where the eigenstates of \hat{H} peak, with the requirement that all the elements of its inverse are non-negative, i.e., $(\hat{H}^{-1})_{ij} \geq 0 \forall i, j$. The landscape function u is the solution to the following differential equation:

$$\hat{H}u = \vec{1} \quad (1)$$

where $\vec{1}$ is a vector of all 1's. For an eigenstate $|\phi^\beta\rangle$ of \hat{H} expressed in an orthonormal basis $\{|J\rangle\}$ with energy E^β , u can be expressed as [35]

$$u_J = \sum_{\beta} \frac{\langle J|\phi^\beta\rangle}{E^\beta} \sum_m \langle m|\phi^\beta\rangle \quad (2)$$

where u_J is the J th component of u and the summation is performed over all the basis states.

Originally developed to predict areas of localization for a single-particle system in a random potential with Dirichlet boundary conditions, u has the useful property of being able to bound the eigenstate amplitudes according to their energies.

An effective potential, W , can be defined from the inverse of the landscape function $W = \frac{1}{u}$, and the regions where low-energy eigenstates peak correspond to minima in W , providing greater insights into the regions of localizations compared to the original potentials, which are seemingly uncorrelated to these regions.

Reference [38] extended the concept of a localization landscape to many-body systems, showing that u bounds the eigenstate amplitudes of \hat{H} according to

$$|\langle J|\phi^\beta\rangle| = |E^\beta| \left| \sum_m (\hat{H}^{-1})_{Jm} \langle m|\phi^\beta\rangle \right| \quad (3)$$

$$\leq |E^\beta| \|\vec{\phi}^\beta\|_\infty \sum_m (\hat{H}^{-1})_{Jm} \quad (4)$$

$$= |E^\beta| \|\vec{\phi}^\beta\|_\infty u_J \quad (5)$$

where $\|\vec{\phi}^\beta\|_\infty = \max_m(|\langle m|\phi^\beta\rangle|)$ is the infinity norm of $\vec{\phi}^\beta$, and the definition of u in Eq. (1) was used to get from Eq. (4) to Eq. (5).

This extension of the localization landscape to many-body systems also places additional considerations on \hat{H} for these bounds to hold, namely that sufficiently short-ranged hopping in \hat{H} is required. For a Fock space graph \mathcal{G}_F where nodes correspond to the N -spin states and edges connect state transitions according to the hopping terms in the potential, this can be realized by maintaining the maximum degree of the Fock space graph \mathcal{G}_F to be linear in N .

Further efforts in Ref. [37] explored the useful properties of the landscape function beyond disordered wave media, laying out additional constraints on the matrix form of \hat{H} for these bounds to hold. More generally, \hat{H} can be a positive semidefinite matrix with $\hat{H}_{ij} \leq 0$ for $i \neq j$, and $\hat{H}_{ij} \geq 0$ for $i = j$.

III. SAMPLING FROM THE LANDSCAPE FUNCTION

Our intention with this paper is to prepare a quantum state $|u\rangle$ that represents the localization landscape function u , from which exact solutions to Eq. (7) can be sampled with probability $\alpha |\langle \vec{x}^*|u\rangle|^2$, where α is the number of degenerate solutions to the problem. Other low-energy solutions can also be sampled with probabilities inversely proportional to their energy, as suggested by Eq. (2).

A. Quadratic unconstrained binary optimization

The QUBO problem can be represented as

$$\text{Find } \vec{x}^* = \underset{\vec{x}}{\operatorname{argmin}} \mathcal{C}_Q(\vec{x}), \quad (6)$$

$$\text{where } \mathcal{C}_Q(\vec{x}) = \vec{x}^\top \mathcal{A} \vec{x}. \quad (7)$$

The vector \vec{x} consists of N binary variables, $\vec{x} = (x_1, \dots, x_N) \in \{0, 1\}^N$, and \mathcal{A} is a real symmetric matrix that defines the problem. Finding optimal solutions to QUBO

problems, \vec{x}^* , is NP-hard in general [14], and serves as a strong impetus for designing classical and quantum heuristics to find approximate solutions.

Quantum algorithms used to solve QUBO problems typically begin by mapping the QUBO cost function to an Ising Hamiltonian of the form

$$\hat{H}_{\text{Ising}} = \frac{1}{4} \sum_{ij}^N \mathcal{A}_{ij} (\hat{\sigma}_i^z + \hat{I}) (\hat{\sigma}_j^z + \hat{I}) \quad (8)$$

where $\hat{\sigma}_i^z$ is the Pauli-Z operator acting on the i th qubit. By mapping each binary variable in \vec{x} to a qubit, the expectation value $\langle \hat{H}_{\text{Ising}} \rangle$ has a minimum value of $C_Q(\vec{x}^*)$ in Eq. (7), and the QUBO problem can be solved by finding the ground state, $|\vec{x}^*\rangle$, that minimizes $\langle \hat{H}_{\text{Ising}} \rangle$.

B. Fitting the constraints

In general, \hat{H}_{Ising} in Eq. (8) does not satisfy the aforementioned constraints for the landscape Eq. (5) to bound the support of the low-energy eigenstates. However, the constraints can be satisfied by introducing the following transformation accompanied by two hyperparameters Γ and λ :

$$\hat{H} = \hat{H}_{\text{Ising}} + \Gamma \hat{I} - \lambda \sum_i^N \hat{\sigma}_i^x \quad (9)$$

where \hat{I} is the identity matrix.

The role of Γ is to add a positive offset to the diagonal elements of \hat{H}_{Ising} that is at least as large as its largest negative eigenvalue. However, the largest negative eigenvalue is typically not known *a priori* as it requires finding the solution to Eq. (7), although in practice it is adequate to pick a sufficiently large value heuristically which can then be further fine tuned.

The ground state of \hat{H}_{Ising} in Eq. (8) is a basis state in the computational Z basis. For problems with symmetries, such as the \mathbb{Z}_2 symmetry in MaxCut problems [42], finding the exact ground state can lead to further ground states with the same energy. In general, being able to find the ground state or an approximate ground state provides little information on nearby states with similar energy values, although there are heuristics that attempt to find “nearby” solutions in terms of energy [43].

The role of λ is to introduce a mixing parameter into \hat{H}_{Ising} to increase the overlap between states that are similar in terms of energy levels. This is done so that the ground state of \hat{H} in Eq. (9) will contain components of surrounding low-energy eigenstates of \hat{H}_{Ising} . It is worth noting that by parametrizing $\lambda = \lambda(t)$, Eq. (9) is often used as the Hamiltonian in quantum annealing, where one starts in the ground state of an easy-to-solve Hamiltonian in the large λ limit and adiabatically decreases $\lambda(t)$ to zero to obtain the ground state of H_{Ising} . The conditions imposed on \hat{H} at the end of Sec. II and the negative sign in Eq. (9) limit $\lambda > 0$.

The Hamiltonian \hat{H} can be visualized using a Fock space graph, \mathcal{G}_F , where nodes representing states of \hat{H}_{Ising} are connected by an edge if they are one spin flip away, corresponding to the potential term $\lambda \sum_i^N \hat{\sigma}_i^x$ in Eq. (9). An example of \mathcal{G}_F for a $N = 4$ Hamiltonian with randomly generated \hat{H}_{Ising} with randomly chosen Γ and λ values satisfying these criteria is

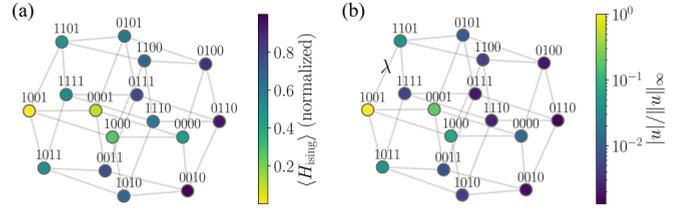


FIG. 1. Representation of the modified Hamiltonian in Eq. (9) as a graph in Fock space for an $N = 4$ random QUBO instance. The considered perturbation induces an N -dimensional hypergraph structure with each bitstring coupled to N nearest neighbors obtained by flipping one bit. Colors show (a) values of $\langle \hat{H}_{\text{Ising}} \rangle$ at each site normalized between 0 and 1, compared with (b) the amplitude of the landscape function $|u|/\|u\|_\infty$ (right).

shown in Fig. 1. The similarities between the peak amplitudes of the localization landscape and the low-energy states of \hat{H}_{Ising} at each site can be observed, along with their decay based on the Hamming distance to the optimal solution (although the rate of decay is different). The short-ranged hopping condition for the many-body localization landscape outlined at the end of Sec. II is satisfied by \mathcal{G}_F having a maximum degree of N .

Thus, we have shown that the QUBO problem can be mapped to a Hamiltonian whose low-energy eigenstates are bounded by the localization landscape, at the cost of introducing two hyperparameters Γ and λ , which control the tightness and extent in Hilbert space of the bounds provided by the landscape, respectively. While the process of finding the optimal values of λ and Γ for each problem instance is beyond the scope of this paper, we will show some results on how they can affect the probability of sampling the optimal solutions in Sec. V.

C. Preparing the landscape function

Once we have the transformed Hamiltonian \hat{H} the final step is to prepare $|u\rangle$, the state that corresponds to the landscape function of \hat{H} . Then, measuring $|u\rangle$ in the computational basis will sample bitstrings corresponding to peaks of $|u\rangle$ (or equivalently, valleys of the landscape) with higher probability. $|u\rangle$ can be obtained by solving the qubit analog of Eq. (1) as a linear system of equations using a quantum device:

$$\hat{H}|u\rangle = |+\rangle \quad (10)$$

where we use $|+\rangle$ to denote the N -qubit superposition state $|+\rangle^{\otimes N}$.

In this paper, we use a variational method [44] to prepare $|u\rangle$ using a variational ansatz $|\psi(\vec{\theta})\rangle = \hat{U}(\vec{\theta})|0\rangle$. This is achieved by minimizing the following variational cost function:

$$f_v(\vec{\theta}) = [|\langle \psi(\vec{\theta}) | \hat{H} | \psi(\vec{\theta}) \rangle - \langle \psi(\vec{\theta}) | + \rangle|^2] \quad (11)$$

which is constructed from Eq. (10) by taking the inner product with $|u\rangle$ on both sides of the equation, squaring the difference between the two terms, and replacing $|u\rangle$ with a variational ansatz.

We note that using a variational approach comes with several potential issues, namely the risk of encountering barren

plateaus [45–47] or having limited expressibility where only a portion of the target state overlaps with states that can be produced. Notably, variational quantum algorithms can also be difficult to optimize [48]. Despite the plethora of issues, we pursue the variational approach here for its simplicity when implemented on NISQ devices and compatibility with shallow hardware-efficient circuits. Common techniques used to mitigate the effects of barren plateaus can also be applied [49–53], although this was not required in obtaining the presented results.

We contrast our variational approach here to traditional variational quantum approaches to QUBO problems, where the ground state of $\langle \hat{H}_{\text{Ising}} \rangle$ is typically not known, and the variational method is used to search for the optimal state, as opposed to preparing a known state. One major advantage of our method is that the target state in our case is known and $|u\rangle$ can be obtained using any of the existing approaches for solving linear equations of the form $A|x\rangle = |b\rangle$, such as the well known Harrow-Hassidim-Lloyd (HHL) algorithm [42]. Other, more NISQ-friendly methods for solving linear equations include variational methods such as the variational quantum linear solver (VQLS) [54,55], the classical combination of variational quantum states [56], and the hybrid classical-quantum linear solver [57]. Regardless, the choice of method used to prepare $|u\rangle$ does not affect the validity of the following results.

IV. RESULTS

To demonstrate the effectiveness of using the landscape function to solve QUBO problems and our variational approach to prepare $|u\rangle$, we apply the methods described above to two problem instances with $N = 10$ variables—a randomly generated fully connected QUBO matrix \mathcal{A} , with $\mathcal{A}_{ij} \in [-1, 1]$ to showcase the nondegenerate case, and a randomly generated three-regular MaxCut problem (formulated as a minimization problem) as a common example of a problem with multiple degenerate solutions. Using the landscape approximation for QUBO is generally problem agnostic, and in later sections, we use the nondegenerate case to further investigate the behavior of the landscape function, and the degenerate case as an example of how prior knowledge about a structured problem can be used to improve the quality of the solutions.

For small problem sizes, the exact solutions can be obtained exactly, and the minimum QUBO cost function for our MaxCut and randomly generated instances are $\mathcal{C}_Q^{\text{MC}}(\vec{x}^*) = -12$ and $\mathcal{C}_Q^{\text{rand}}(\vec{x}^*) \approx -7.895$ respectively. To fit our problems to the constraints, we used a value of $\Gamma = 8.5$ and $\lambda = 0.3$ for the randomly generated QUBO instance, and $\Gamma = 13$ and $\lambda = 0.3$ for the MaxCut instance.

Figure 2 shows the landscape function u of the respective perturbed Hamiltonians \hat{H} for both the randomly generated QUBO instance and the MaxCut problem, compared with the four lowest-energy eigenstates for the Ising Hamiltonians representing the two problem instances. We observe that the peaks of the landscape function u line up with basis states of \hat{H} which are the lowest-energy eigenstates of \hat{H}_{Ising} . Based off Fig. 2, we intend to prepare the target state $|u\rangle$ that will have amplitudes of a similar structure to u as shown in Fig. 2, from

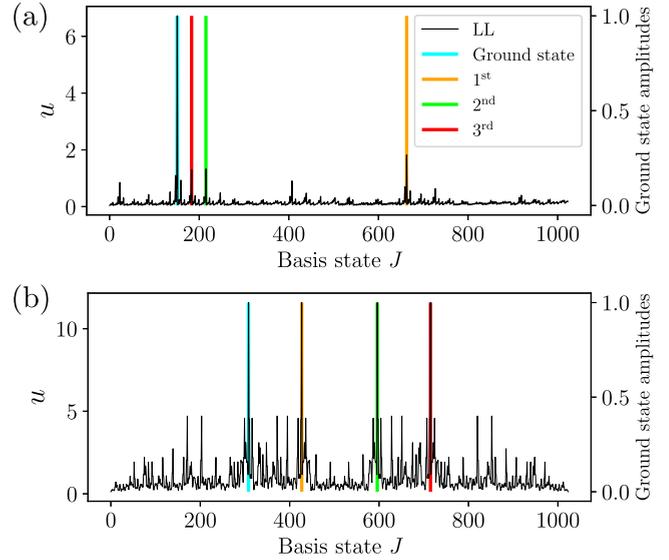


FIG. 2. Localization landscape, u , of \hat{H} for $N = 10$ qubits as constructed in Eq. (9), compared with the four lowest-energy states of \hat{H}_{Ising} for (a) a randomly generated QUBO instance (nondegenerate case) and (b) a randomly generated three-regular MaxCut problem (degenerate case). The peaks of the landscape function correspond to the low-energy eigenstates of \hat{H}_{Ising} .

which these low-energy states of \hat{H}_{Ising} can be sampled with high probability.

For both problem instances, we used the same variational ansatz consisting of an initial layer of Hadamard gates on all qubits, followed by four alternating layers of $R_y(\theta)$ rotations on all qubits and nearest-neighbor controlled-NOT (CNOT) entangling gates in a linear topology, keeping all the coefficients of the quantum state real. We used COBYLA [58–60], a gradient-free classical optimizer, to search for the optimal parameters that minimize Eq. (11) from ten initial starting sets of $\vec{\theta}$ angles. Gradient-based optimizers can also be used [61], and we show how the gradient of f_v can be obtained in Appendix A using a gate-based circuit that produces a quantum state with only real coefficients.

We compare the expectation value of $\langle \hat{H}_{\text{Ising}} \rangle$ obtained using our postoptimized state, $|\psi(\vec{\theta}^*)\rangle$, and from $|u\rangle$ obtained from inverting \hat{H} in Eq. (10). We also compare the final $\langle \hat{H}_{\text{Ising}} \rangle$ values obtained using both our landscape method and the QAOA with $p = 1$ layers. Finally, we compare the solutions obtained by sampling from our optimized ansatz, from $|u\rangle$, and from $|\psi(\gamma, \beta)\rangle_{\text{QAOA}}$ with $p = 1$ as defined in Appendix B. All simulations were conducted using the state vector simulator (i.e., number of shots $\rightarrow \infty$) in PENNYLANE [62].

In Fig. 3 we show the optimization runs used to prepare $|u\rangle$ using our variational ansatz, and the quality of the solutions obtained after every 200 iterations of COBYLA used to calculate the classical QUBO cost function, \mathcal{C}_Q in Eq. (7) for both the MaxCut and random QUBO instances. Also shown in Fig. 3 are comparisons between $\langle \hat{H}_{\text{Ising}} \rangle$ from preparing the exact $|u\rangle$, from randomly sampling bitstrings over a uniform distribution, from optimizing the QAOA with $p = 1$, and

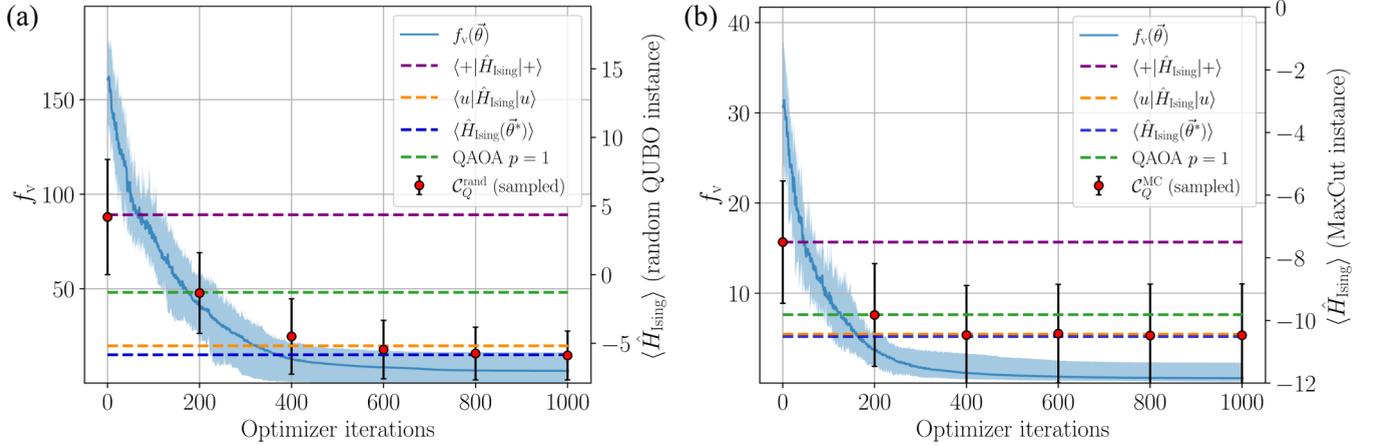


FIG. 3. Variational search for $|u\rangle$ for (a) a randomly generated QUBO problem with $-1 < \mathcal{A}_{ij} < 1$ (nondegenerate case) and (b) a randomly generated three-regular MaxCut problem (degenerate case). The variational search was performed using COBYLA to minimize $f_v(\vec{\theta})$ in Eq. (11) over ten unique initial sets of $\vec{\theta}$ using the state vector simulator in PENNYLANE. Solid lines show the average cost function at each optimizer iteration over ten unique initial sets of $\vec{\theta}$. Shaded areas show the minimum and maximum of $f_v(\vec{\theta})$ over the ten runs at each iteration. Dashed lines show different $\langle \hat{H}_{\text{Ising}} \rangle$ values, obtained from randomly sampling bitstrings over a uniform distribution (purple), preparing $|u\rangle$ exactly (orange), from the output state of our variational circuit after attempting to minimize $f_v(\vec{\theta})$ (blue), and from the optimized state of the QAOA with $p = 1$ (green). Red markers and error bars show the average and standard deviation of classical QUBO cost function C_Q obtained from ten bitstrings sampled every 200 iterations per optimization run.

from our variational ansatz after optimization $\langle \hat{H}_{\text{Ising}}(\vec{\theta}^*) \rangle = \langle \psi(\vec{\theta}^*) | \hat{H}_{\text{Ising}} | \psi(\vec{\theta}^*) \rangle$. Further information regarding our implementation of the QAOA and the number of CNOT gates used can be found in Appendix B.

As observed in Fig. 3, the mean QUBO cost function from sampled bitstrings obtained every 200 iterations tends towards $\langle u | \hat{H}_{\text{Ising}} | u \rangle$ as the ansatz converges to a state representing $|u\rangle$. In both problem instances, being able to prepare $|u\rangle$, whether exactly or using our simple circuit ansatz, allows one to obtain a lower cost function value compared with the QAOA.

For the nondegenerate case, being able to prepare and sample from the exact landscape function state $|u\rangle$ brings us closer to the optimal C_Q value compared to $p = 1$ of the QAOA. This is likely due to our specific problem and choice of hyperparameters, where the mixing introduced by λ in Eq. (9) is small compared to the difference between the lowest two energy states of $\langle \hat{H}_{\text{Ising}} \rangle$ for the nondegenerate case. This causes the ground state of $\langle \hat{H}_{\text{Ising}} \rangle$ to be the dominant basis state in the ground state of \hat{H} and preparing $|u\rangle$ will produce a strong peak at $|\vec{x}^*\rangle$.

V. EFFECT OF HYPERPARAMETERS

In this section we explore how the hyperparameters Γ and λ affect the quality of the solutions obtained for the nondegenerate case, although similar properties hold for degenerate problem instances as well.

To properly characterize the capabilities of $|u\rangle$, the results presented from this section on are limited to states that can be prepared exactly. As mentioned in Sec. III C, the variational method in Sec. IV was mainly an example of how $|u\rangle$ can be prepared quickly using NISQ-friendly methods, and other methods can be used for $|u\rangle$ with potentially higher accuracy.

We begin by noting that the optimal solution to the QUBO problem in Eq. (6) can be represented by a computational basis state of \hat{H}_{Ising} in Eq. (8) used to construct \hat{H} . Using Eq. (5), we can find the probability amplitude associated with sampling $|\vec{x}^*\rangle$ if we have prepared the ground state of \hat{H} :

$$u_{x^*} |E^\beta| \|\phi_\beta\|_\infty \geq |\langle \vec{x}^* | \phi^\beta \rangle| \quad (12)$$

where in this case we let $|\phi^\beta\rangle$ and E^β be the ground state and ground-state energy of \hat{H} , respectively. For small values of λ , we can expand the denominator using perturbation theory and express E^β in terms of Γ , λ , and the ground-state energy of \hat{H}_{Ising} [i.e., $E^* = C_Q(\vec{x}^*)$]. To first order, this gives

$$u_{x^*} \geq \frac{|\langle \vec{x}^* | \phi^\beta \rangle|}{|E^\beta| \|\phi_\beta\|_\infty} \quad (13)$$

$$\approx \frac{|\langle \vec{x}^* | \phi^\beta \rangle|}{|E^* + \Gamma + \lambda \langle \vec{x}^* | \sum_i \hat{\sigma}_i^x | \vec{x}^* \rangle| \|\phi_\beta\|_\infty}. \quad (14)$$

The left-hand side of Eq. (13), u_{x^*} , is not $|\langle \vec{x}^* | u \rangle|$, since the landscape function u from Eq. (5) is not a normalized state. Nevertheless, it is related to the probability amplitude of sampling $|\vec{x}^*\rangle$ from $|u\rangle$ and it is still in our interest to maximize it. According to Eq. (14), this can be done by choosing Γ to be as close to $-E^*$ as possible.

Figure 4(a) shows the probability of sampling the optimal solution \vec{x}^* from $|u\rangle$ as a function of both Γ and λ for the randomly generated QUBO instance. Values of $\lambda > 1$ are beyond the perturbative regime used in the approximation in Eq. (14).

Shown in Fig. 4(b) is the Hamming distance between the optimal solution \vec{x}^* and the vector \vec{x} that has the highest probability of being sampled from $|u\rangle$ for the randomly generated QUBO instance. This can be expressed more succinctly as

$$d(\vec{x}^*, \underset{\vec{x}}{\operatorname{argmax}} |\langle \vec{x} | u \rangle|^2), \quad (15)$$

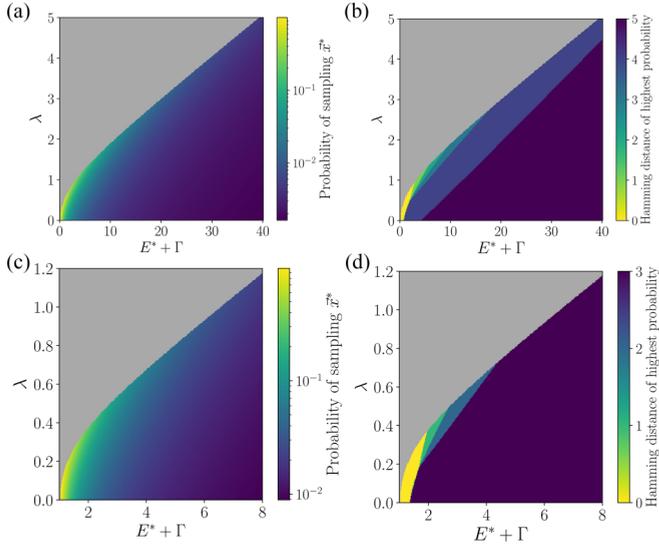


FIG. 4. (a), (c) Probability of sampling x^* from $|u\rangle$ as a function of hyperparameters λ and Γ for (a) the random QUBO instance and (c) the three-regular MaxCut instance. (b), (d) Bitstrings with Hamming distance to solution most likely to be sampled from $|u\rangle$ as a function of hyperparameters λ and Γ for (b) the random QUBO instance and (d) the three-regular MaxCut instance. Gray areas show values of λ and Γ for which there exist elements of $\hat{H}^{-1} < 0$, where the conditions for the landscape to bound the low-energy eigenstates are not satisfied.

where $d(\bar{x}_1, \bar{x}_2)$ is the Hamming distance between \bar{x}_1 and \bar{x}_2 . Figures 4(b) and 4(d) show the same plots but for the MaxCut instance. The results in Fig. 4 also suggest that sampling solutions close to the optimum favor having λ to be as large as possible while still respecting the constraints for a given $E^* + \Gamma$, which should be as close to zero as possible.

In Fig. 5, we compare the total probability of sampling solutions with Hamming distance d away from the optimal solution for the randomly generated QUBO problem from three states— $|u\rangle$, the ground state of the perturbed Hamiltonian \hat{H} , and the same optimal state of the QAOA with $p = 1$ in Fig. 3. These probabilities change as a function of Γ and λ .

In Fig. 5(a) Γ and λ are too small, and there is insufficient mixing between the low-energy states of the corresponding \hat{H}_{Ising} , and preparing the landscape function provides similar probabilities to sample the ground state of \hat{H}_{Ising} from the perturbed Hamiltonian. This can be desirable in most cases where we are only interested in the optimal solution to the QUBO problem.

In contrast, for cases where Γ and/or λ are too large, such as in Figs. 5(b) and 5(c), the majority of the solutions sampled will be approximately $\frac{N}{2}$ Hamming distances away. For large values of λ , the perturbation term in \hat{H} becomes dominant compared to the ZZ interactions in \hat{H}_{Ising} , and $|u\rangle$ tends toward the uniform superposition state in the computational basis.

However, there is also an interesting regime in Fig. 5(d) where, for well-chosen values of Γ and λ , sampling from $|u\rangle$ is able to provide the optimal solution with a high probability along with nearby solutions in terms of Hamming distance. In practice, this can be used to find the optimal bitstring from just a handful of samples on $|u\rangle$.

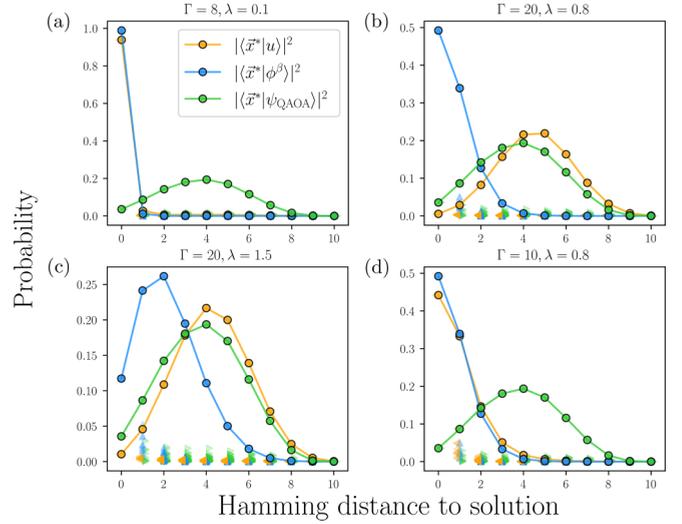


FIG. 5. Probability of sampling a solution with Hamming distance d away from the optimal solution \bar{x}^* for the randomly generated QUBO problem for different values of Γ and λ . Colored triangles show the probability of sampling these individual solutions from $|u\rangle$ (orange), from the ground state of \hat{H} (blue), and from the optimal state of the QAOA with $p = 1$ (green). Connected lines show the total probability of all solutions of Hamming distance d away from the optimal solution (i.e., the sum of all the triangles at a given d).

On the other hand, sampling from the optimal state produced by the QAOA with $p = 1$ will result in the majority of the samples being $\frac{N}{2}$ Hamming distance away from the optimal solution.

We note that in all of these cases, the probability of obtaining the optimal solution \bar{x}^* from $|u\rangle$ is still higher than any other bitstring, although it may not form the majority of the samples obtained.

The main results presented so far mainly concerned two different problem instances for $N = 10$. Figure 6 represents an initial foray into how using the localization landscape scales with problem sizes, as well as how prior knowledge of the problem can be used to increase the probability of sampling optimal solutions.

For each problem type (random QUBO and MaxCut instances), the probability of sampling eigenstates of \hat{H}_{Ising} (\bar{x}^* , \bar{x}_1 , and \bar{x}_2) corresponding to the three distinct lowest-energy values (E^* , E_1 , and E_2) is plotted against the problem size N , and compared against the probability obtained from sampling the optimal solution from a uniform distribution. For the MaxCut instances in Fig. 6(b), the curves show the total probability, i.e., $\alpha |\langle \bar{x}_i | u \rangle|^2$, where α is the number of degenerate states corresponding to energy E_i .

At first glance, it is worth noting that while the exponentially decreasing probability of sampling the optimal solution may pose a glaring issue, especially for the random QUBO instance, the probability of sampling \bar{x}^* remains consistently above that of random sampling.

As mentioned earlier, a “good” choice of Γ would be one that is as close to $-E^*$ as possible. For any QUBO problem, $-\sum_{i,j} |A_{ij}|$ is a lower bound of \mathcal{C}_Q and an initial value $\Gamma = 1.1 \times \sum_{i,j} |A_{ij}|$ can be used for unstructured problems.

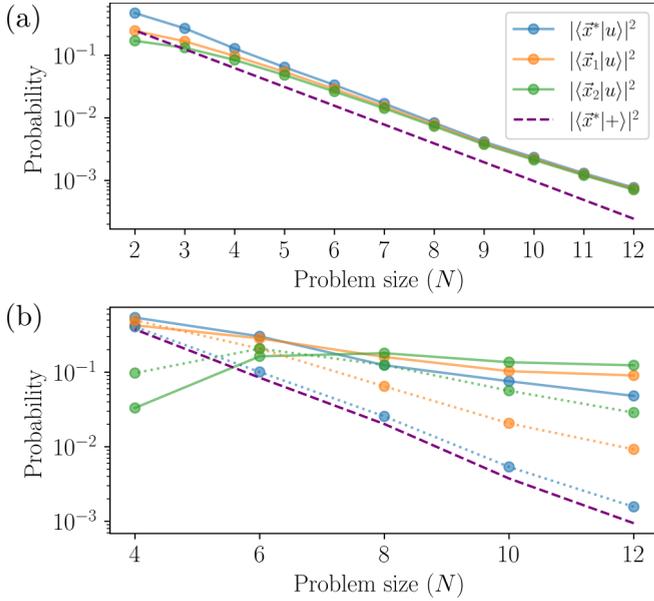


FIG. 6. Probability of sampling eigenstates of \hat{H}_{Ising} corresponding to the first three lowest-energy states from $|u\rangle$ for (a) randomly generated QUBO problems (nondegenerate case) and (b) randomly generated three-regular MaxCut problems (degenerate case). Solid lines in (b) show these probabilities using Γ and λ values chosen with some prior knowledge of the problem instances. Dotted lines in (b) show these probabilities using the same method of choosing Γ and λ values as in (a). Purple dashed lines show the probability of sampling \vec{x}^* from a uniform distribution. Each plot point was obtained by averaging over 100 randomly generated problem instances, and $|u\rangle$ was found by solving Eq. (10) exactly.

For a MaxCut problem, one can use the maximum possible number of edge bisections in a graph as the lower bound for C_Q . This is equal to the total number of edges, $n_e = \frac{Nd}{2}$, for a d -regular graph. For three-regular graphs, one can choose $\Gamma = \frac{3N}{2} + 1$ which is typically less than $\sum_{i,j} |A_{ij}|$ to obtain a much higher probability in sampling \vec{x}^* , as observed when comparing the solid and dotted lines in Fig. 6(b).

We used a value of $\lambda = 0.07 \Gamma$ for the random QUBO instances, and $\lambda = 0.03 \Gamma$ for the MaxCut problems to fulfill the constraints in Sec. II for the instances considered in Fig. 6. However, these values may not be valid or ideal for all QUBO problems in general. In all problem instances, the exponential decrease in probability can be ameliorated with further tuning of Γ and λ for the specific instance.

Another interesting observation of Fig. 6(b) is how higher-energy states can have a greater overall probability of being sampled compared to the optimal solution. This can be explained by the increase in number of degenerate states closer to the middle of the energy spectrum. As shown in Fig. 2(c), the probability of sampling individual ground states is still dominant compared to the other states.

As mentioned in Sec. III C, quantum approaches such as the HHL or VQLS algorithm can be used to prepare $|u\rangle$. The runtime complexity of these algorithms grows with the condition number of the perturbed Hamiltonian, $\kappa(\hat{H})$, defined as the ratio between the largest and smallest singular values [42,63]. Here, we investigate how the condition number of

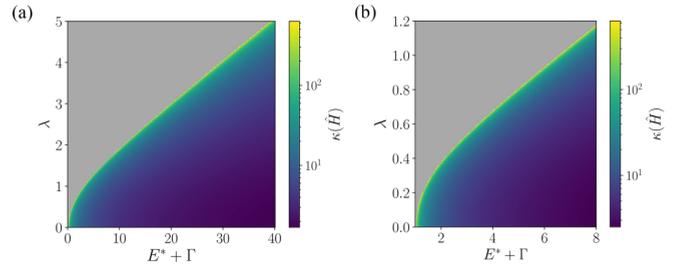


FIG. 7. Condition number of the perturbed Hamiltonian, $\kappa(\hat{H})$, for (a) the randomly generated QUBO instance and (b) the three-regular MaxCut instance, as a function of the hyperparameters Γ and λ . Gray areas show values of λ and Γ for which there exist elements of $\hat{H}^{-1} < 0$ and the conditions outlined in Sec. II are not satisfied. We observe that $\kappa(\hat{H})$ diverges when both λ and $E^* + \Gamma$ are small, suggesting that having $|u\rangle$ to tightly bound the exact ground state of \hat{H}_{Ising} can make it more difficult to prepare.

\hat{H} , and therefore the difficulty of preparing $|u\rangle$, depend on the perturbation terms Γ and λ . Figure 7 shows $\kappa(\hat{H})$ for different values of Γ and λ used to perturb \hat{H}_{Ising} according to Eq. (9). We observe, for both the random QUBO instance in Fig. 7(a) and the MaxCut instance in Fig. 7(b), that the condition number decreases for increasing values of Γ and λ .

In Eq. (14), we saw that the localization landscape more tightly bounds the ground state of \hat{H}_{Ising} as $E^* + \Gamma$ (i.e., the smallest eigenvalue of H) and λ are close to zero. The divergence of the condition number in this limit suggests there is no free lunch from using a landscape that very tightly bounds the exact ground state, since it will also be difficult to prepare. Conversely, the decrease in condition number in Fig. 7 further suggests that adding offset and mixing terms to perturb the Hamiltonian allows for approximate solutions such as $|u\rangle$ to be more easily prepared and sampled from.

VI. DISCUSSION AND CONCLUSION

A key part in obtaining the results in this paper was perturbing (\hat{H}_{Ising}), which is diagonal in the computational basis, with a uniform transverse magnetic field $\sum_i^N \hat{\sigma}_i^z$, equivalent to a uniform nearest-neighbor hopping on an N -dimensional hypercube. This was done to controllably smear out the eigenstates of \hat{H}_{Ising} in the Fock space, allowing for the QUBO problem to be solved approximately by sampling from the solutions of the easier-to-solve landscape problem.

As we have observed, the quality of the resulting solutions will depend on the strength and form of the perturbing potential, and the properties of alternative perturbation terms provide interesting avenues for further exploration. For example, one may replace the perturbative term with a number-conserving perturbation (arising in the case of models of many-body localization), such as $\sum_{i,j} (\hat{\sigma}_i^+ \hat{\sigma}_j^- + \text{H.c.})$ where $\hat{\sigma}^\pm = \frac{1}{2}(\hat{\sigma}^x \pm i\hat{\sigma}^y)$, leading to landscape functions that explore decoupled subspaces of the full Hilbert space as in Ref. [38]. Such a perturbative term may be more useful when sampling solutions to QUBO problems involving hard constraints, such as those requiring the number of spin excitations to be preserved.

Another interesting avenue for exploration would be to consider QUBO problems where the eigenvalue spectrum of the Hamiltonian encoding the problem is skewed towards having a few low-energy states separated from many high-energy states by a large gap. These types of QUBO problems are typically present in industry-relevant contexts, where the use of a penalty term when constructing the unconstrained problem causes all solutions that do not satisfy any constraints to have very high costs. By preparing the landscape function, it should be possible to prepare a state such that solutions satisfying the constraints can be sampled more easily, and the optimal solution can be easily found from this smaller, finite group of samples.

In conclusion, we showed how to apply the localization landscape theory used to find localized regions of low-energy eigenstates in many-body systems to prepare quantum states that can be used to sample low-energy solutions to the QUBO problem with high probability. We demonstrated our methods on two problem instances, a randomly generated MaxCut problem [2,40] exemplifying the degenerate case and a randomly generated QUBO problem for the nondegenerate case, and showed that by preparing a state, $|u\rangle$, representing the landscape function, low-energy solutions to the Ising Hamiltonian corresponding to the QUBO problem can be sampled with higher probability. An advantage of the approach is that the good solutions can be sampled using relatively shallow circuits, minimizing the effect of gate noise and decoherence present in current noisy intermediate-scale quantum processors.

ACKNOWLEDGMENTS

We acknowledge support from the National Research Foundation, Prime Minister's Office, Singapore and the Agency for Science, Technology and Research, Singapore under the CQT Bridging Grant and the Quantum Engineering Programme (NRF2021QEP2-02-P02), and by the European Union's Horizon Programme (HORIZON-CL4-2021-DIGITALEMERGING-02-10), Grant Agreement No. 101080085, QCFD.

APPENDIX A: GRADIENT CALCULATION

In this Appendix, we show how the derivative of the cost function, $\frac{\partial f_v}{\partial \theta_i}$, can be obtained *in situ* using a quantum device and the parameter shift rule:

$$f_v = [\langle \psi(\vec{\theta}) | \hat{H} | \psi(\vec{\theta}) \rangle - \langle \psi(\vec{\theta}) | + \rangle]^2 \quad (\text{A1})$$

$$\frac{\partial f_v}{\partial \theta_i} = 2[\langle \psi(\vec{\theta}) | \hat{H} | \psi(\vec{\theta}) \rangle - \langle \psi(\vec{\theta}) | + \rangle] \times \frac{\partial}{\partial \theta_i} [\langle \psi(\vec{\theta}) | \hat{H} | \psi(\vec{\theta}) \rangle - \langle \psi(\vec{\theta}) | + \rangle]. \quad (\text{A2})$$

From here, we will proceed term by term. Using the parameter shift rule,

$$\frac{\partial}{\partial \theta_i} \langle \psi(\vec{\theta}) | \hat{H} | \psi(\vec{\theta}) \rangle = \frac{1}{2} \left[\langle \hat{H} \rangle \left(\theta_i + \frac{\pi}{2} \right) - \langle \hat{H} \rangle \left(\theta_i - \frac{\pi}{2} \right) \right]. \quad (\text{A3})$$

To evaluate $\frac{\partial}{\partial \theta_i} [\langle \psi(\vec{\theta}) | + \rangle]$, we observe that for a real quantum state $|\psi(\vec{\theta})\rangle$

$$\frac{\partial}{\partial \theta_i} |\langle \psi(\vec{\theta}) | + \rangle|^2 = 2 \langle \psi(\vec{\theta}) | + \rangle \frac{\partial}{\partial \theta_i} [\langle \psi(\vec{\theta}) | + \rangle] \quad (\text{A4})$$

and

$$\frac{\partial}{\partial \theta_i} |\langle \psi(\vec{\theta}) | + \rangle|^2 = \frac{\partial}{\partial \theta_i} [\langle \psi(\vec{\theta}) | + \rangle \langle + | \psi(\vec{\theta}) \rangle]. \quad (\text{A5})$$

Putting Eqs. (A4) and (A5) together, and letting $\hat{M} = |+\rangle\langle +| = (\frac{I + \hat{\sigma}_x}{2})^{\otimes N}$, we obtain

$$\frac{\partial}{\partial \theta_i} [\langle \psi(\vec{\theta}) | + \rangle] = \frac{1}{2} \frac{1}{\langle \psi(\vec{\theta}) | + \rangle} \frac{\partial}{\partial \theta_i} \langle \psi(\vec{\theta}) | \hat{M} | \psi(\vec{\theta}) \rangle \quad (\text{A6})$$

$$= \frac{1}{2} \frac{1}{\langle \psi(\vec{\theta}) | + \rangle} \frac{\partial}{\partial \theta_i} \langle \hat{M} \rangle(\vec{\theta}) \quad (\text{A7})$$

$$= \frac{1}{4} \frac{\langle \hat{M} \rangle(\theta_i + \frac{\pi}{2}) - \langle \hat{M} \rangle(\theta_i - \frac{\pi}{2})}{\langle \psi(\vec{\theta}) | + \rangle} \quad (\text{A8})$$

where we have used the parameter shift rule in Eq. (A8) to evaluate $\frac{\partial}{\partial \theta_i} \langle \hat{M} \rangle$. Evaluating the gradient $\frac{\partial f_v}{\partial \theta_i}$ therefore requires three state preparations per variational parameter, at θ_i , $\theta_i + \frac{\pi}{2}$, and $\theta_i - \frac{\pi}{2}$.

APPENDIX B: QUANTUM APPROXIMATE OPTIMIZATION ALGORITHM

The QAOA is a variational quantum algorithm for finding approximate solutions to combinatorial optimization problems. The QAOA state is parametrized by two sets of angles, $\vec{\gamma} = \{\gamma_1, \dots, \gamma_p\}$ and $\vec{\beta} = \{\beta_1, \dots, \beta_p\}$:

$$|\psi(\vec{\gamma}, \vec{\beta})\rangle_{\text{QAOA}} = \prod_i^p U_x(\beta_i) U_H(\gamma_i) |+\rangle \quad (\text{B1})$$

where

$$U_H(\gamma) = e^{-i\gamma \hat{H}_{\text{Ising}}}, \quad (\text{B2})$$

$$U_x(\beta) = e^{-i\beta \sum_i \hat{\sigma}_i^x}. \quad (\text{B3})$$

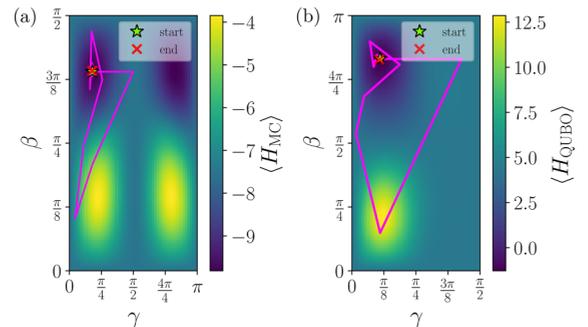


FIG. 8. Grid search for the QAOA with $p = 1$ for (a) the MaxCut problem and (b) a randomly generated QUBO problem. The expectation values $\langle \hat{H}_{\text{Ising}} \rangle$ for the respective problems were calculated for a 100×100 grid using the state vector simulator in PENNYLANE. The optimal parameters found using the grid search (green star) were used as starting parameters for further fine tuning of $\langle \hat{H}_{\text{Ising}} \rangle$ using COBYLA (magenta line). The red cross denotes the final parameters obtained using COBYLA.

In Sec. IV, we compared the results obtained from preparing the landscape state $|u\rangle$ with results obtained from $p = 1$ of QAOA. For $p = 1$, the state in Eq. (B1) only contains two variational parameters, and the optimal parameters to obtain the QAOA results in Fig. 3 were found by using a grid search with a resolution of 100×100 and then using COBYLA to perform a local search, further fine tuning the parameters. Figure 8 shows the grid and fine tuning needed to obtain the optimal parameters.

The variational ansatz described in Sec. IV to produce the results in Fig. 3 uses $4 \times (N - 1) = 36$ CNOT gates. By comparison, decomposing the unitaries in the QAOA to similar gatesets requires $2 \times n_e$ number of CNOT gates per depth p , where n_e is the number of edges in the problem graph [64]. For $p = 1$, this amounts to 30 and 90 CNOT gates for the MaxCut and random QUBO graph before accounting for measures to handle long-ranged interactions between qubits.

-
- [1] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke *et al.*, Noisy intermediate-scale quantum algorithms, *Rev. Mod. Phys.* **94**, 015004 (2022).
- [2] F. Glover, G. Kochenberger, and Y. Du, A tutorial on formulating and using QUBO models, [arXiv:1811.11538](https://arxiv.org/abs/1811.11538).
- [3] A. I. Pakhomchik, S. Yudin, M. R. Perelshtein, A. Alekseyenko, and S. Yarkoni, Solving workflow scheduling problems with QUBO modeling, [arXiv:2205.04844](https://arxiv.org/abs/2205.04844).
- [4] C. Papalitsas, T. Andronikos, K. Giannakis, G. Theocharopoulou, and S. Fanarioti, A QUBO model for the traveling salesman problem with time windows, *Algorithms* **12**, 224 (2019).
- [5] P. Date, D. Arthur, and L. Pusey-Nazzaro, Qubo formulations for training machine learning models, *Sci. Rep.* **11**, 10029 (2021).
- [6] C. Bauckhage, N. Piatkowski, R. Sifa, D. Hecker, and S. Wrobel, A QUBO formulation of the k-medoids problem, in *Lernen. Wissen. Daten. Analysen.* (Fraunhofer, Berlin, 2009), pp. 54–63.
- [7] H. Neven, V. S. Denchev, G. Rose, and W. G. Macready, Training a binary classifier with the quantum adiabatic algorithm, [arXiv:0811.0416](https://arxiv.org/abs/0811.0416).
- [8] F. Bapst, W. Bhimji, P. Calafiura, H. Gray, W. Lavrijsen, L. Linder, and A. Smith, A pattern recognition algorithm for quantum annealers, *Comput. Software Big Sci.* **4**, 1 (2020).
- [9] N. Matsumoto, Y. Hamakawa, K. Tatsumura, and K. Kudo, Distance-based clustering using QUBO formulations, *Sci. Rep.* **12**, 2669 (2022).
- [10] F. Glover, G. Kochenberger, and Y. Du, Applications and computational advances for solving the QUBO model, in *The Quadratic Unconstrained Binary Optimization Problem: Theory, Algorithms, and Applications* (Springer, New York, 2022), pp. 39–56.
- [11] W. Guan, G. Perdue, A. Pesah, M. Schuld, K. Terashi, S. Vallecorsa, and J.-R. Vlimant, Quantum machine learning in high energy physics, *Mach. Learn.: Sci. Technol.* **2**, 011003 (2021).
- [12] Z. Huang, Q. Li, J. Zhao, and M. Song, Variational quantum algorithm applied to collision avoidance of unmanned aerial vehicles, *Entropy* **24**, 1685 (2022).
- [13] H. Wang, D. Huo, and B. Alidaee, Position unmanned aerial vehicles in the mobile ad hoc network, *J. Intell. Rob. Syst.* **74**, 455 (2014).
- [14] Y. Fu and P. W. Anderson, Application of statistical mechanics to NP-complete problems in combinatorial optimisation, *J. Phys. A* **19**, 1605 (1986).
- [15] C. H. Papadimitriou, Computational complexity, in *Encyclopedia of Computer Science* (Wiley, New York, 2003), pp. 260–265.
- [16] F. Barahona, On the computational complexity of Ising spin glass models, *J. Phys. A* **15**, 3241 (1982).
- [17] D. S. Hochba, Approximation algorithms for np-hard problems, *ACM Sigact News* **28**, 40 (1997).
- [18] J. K. Lenstra, D. B. Shmoys, and É. Tardos, Approximation algorithms for scheduling unrelated parallel machines, *Math. Program.* **46**, 259 (1990).
- [19] D. P. Williamson and D. B. Shmoys, *The Design of Approximation Algorithms* (Cambridge University, New York, 2011).
- [20] A. Rajak, S. Suzuki, A. Dutta, and B. K. Chakrabarti, Quantum annealing: An overview, *Phil. Trans. R. Soc. A* **381**, 20210417 (2023).
- [21] G. E. Santoro and E. Tosatti, Optimization using quantum mechanics: quantum annealing through adiabatic evolution, *J. Phys. A* **39**, R393 (2006).
- [22] A. Das and B. K. Chakrabarti, Colloquium: Quantum annealing and analog quantum computation, *Rev. Mod. Phys.* **80**, 1061 (2008).
- [23] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
- [24] S. Hadfield, Z. Wang, B. O’gorman, E. G. Rieffel, D. Venturelli, and R. Biswas, From the quantum approximate optimization algorithm to a quantum alternating operator ansatz, *Algorithms* **12**, 34 (2019).
- [25] 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).
- [26] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [27] R. Shaydulin, H. Ushijima-Mwesigwa, C. F. A. Negre, I. Safro, S. M. Mniszewski, and Y. Alexeev, A hybrid approach for solving optimization problems on small quantum computers, *Computer* **52**, 18 (2019).
- [28] O. Kyriienko, Quantum inverse iteration algorithm for programmable quantum simulators, *npj Quantum Inf.* **6**, 7 (2020).
- [29] K. Bharti and T. Haug, Iterative quantum-assisted eigensolver, *Phys. Rev. A* **104**, L050401 (2021).
- [30] D. Lykov, J. Wurtz, C. Poole, M. Saffman, T. Noel, and Y. Alexeev, Sampling frequency thresholds for quantum advantage of quantum approximate optimization algorithm, [arXiv:2206.03579](https://arxiv.org/abs/2206.03579).
- [31] P. W. Anderson, Absence of diffusion in certain random lattices, *Phys. Rev.* **109**, 1492 (1958).

- [32] M. Filoche and S. Mayboroda, Universal mechanism for Anderson and weak localization, *Proc. Natl. Acad. Sci. USA* **109**, 14761 (2012).
- [33] D. N. Arnold, G. David, M. Filoche, D. Jerison, and S. Mayboroda, Computing spectra without solving eigenvalue problems, *SIAM J. Sci. Comput.* **41**, B69 (2019).
- [34] G. David, M. Filoche, and S. Mayboroda, The landscape law for the integrated density of states, *Adv. Math. (NY)* **390**, 107946 (2021).
- [35] L. Herviou and J. H. Bardarson, \mathcal{L}^2 localization landscape for highly excited states, *Phys. Rev. B* **101**, 220201(R) (2020).
- [36] M. Kakoi and K. Slevin, A stochastic method to compute the 12 localisation landscape, *J. Phys. Soc. Jpn.* **92**, 054707 (2023).
- [37] M. Filoche, S. Mayboroda, and T. Tao, The effective potential of an m -matrix, *J. Math. Phys.* **62**, 041902 (2021).
- [38] S. Balasubramanian, Y. Liao, and V. Galitski, Many-body localization landscape, *Phys. Rev. B* **101**, 014201 (2020).
- [39] G. A. Hamilton and B. K. Clark, Analysis of many-body localization landscapes and Fock space morphology via persistent homology, [arXiv:2302.09361](https://arxiv.org/abs/2302.09361).
- [40] M. R. Garey, D. S. Johnson, and L. Stockmeyer, Some simplified NP-complete problems, in *Proceedings of the Sixth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, 1974), pp. 47–63.
- [41] D. Amaro, C. Modica, M. Rosenkranz, M. Fiorentini, M. Benedetti, and M. Lubasch, Filtering variational quantum algorithms for combinatorial optimization, *Quantum Sci. Technol.* **7**, 015021 (2022).
- [42] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum algorithm for linear systems of equations, *Phys. Rev. Lett.* **103**, 150502 (2009).
- [43] S. T. Goh, J. Bo, S. Gopalakrishnan, and H. C. Lau, Techniques to enhance a QUBO solver for permutation-based combinatorial optimization, in *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (Association for Computing Machinery, New York, 2022), pp. 2223–2231.
- [44] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [45] 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).
- [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] 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).
- [48] L. Bittel and M. Kliesch, Training variational quantum algorithms is NP-hard, *Phys. Rev. Lett.* **127**, 120502 (2021).
- [49] X. Liu, G. Liu, J. Huang, H.-K. Zhang, and X. Wang, Mitigating barren plateaus of variational quantum eigensolvers, [arXiv:2205.13539](https://arxiv.org/abs/2205.13539).
- [50] H.-Y. Liu, T.-P. Sun, Y.-C. Wu, Y.-J. Han, and G.-P. Guo, Mitigating barren plateaus with transfer-learning-inspired parameter initializations, *New J. Phys.* **25**, 013039 (2023).
- [51] E. Grant, L. Wossnig, M. Ostaszewski, and M. Benedetti, An initialization strategy for addressing barren plateaus in parametrized quantum circuits, *Quantum* **3**, 214 (2019).
- [52] A. Mari, T. R. Bromley, J. Izaac, M. Schuld, and N. Killoran, Transfer learning in hybrid classical-quantum neural networks, *Quantum* **4**, 340 (2020).
- [53] 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).
- [54] C. Bravo-Prieto, R. LaRose, M. Cerezo, Y. Subasi, L. Cincio, and P. J. Coles, Variational quantum linear solver *Quantum* **7**, 1188, (2020).
- [55] H. Patil, Y. Wang, and P. S. Krstić, Variational quantum linear solver with a dynamic ansatz, *Phys. Rev. A* **105**, 012423 (2022).
- [56] H.-Y. Huang, K. Bharti, and P. Rebentrost, Near-term quantum algorithms for linear systems of equations with regression loss functions, *New J. Phys.* **23**, 113021 (2021).
- [57] C.-C. Chen, S.-Y. Shiao, M.-F. Wu, and Y.-R. Wu, Hybrid classical-quantum linear solver using noisy intermediate-scale quantum machines, *Sci. Rep.* **9**, 16251 (2019).
- [58] M. J. D. Powell, A direct search optimization method that models the objective and constraint functions by linear interpolation, in *Advances in Optimization and Numerical Analysis*, edited by S. Gomez and J.-P. Hennart (Springer, New York, 1994), pp. 51–67.
- [59] M. J. D. Powell, Direct search algorithms for optimization calculations, *Acta Numerica* **7**, 287 (1998).
- [60] M. J. Powell, A view of algorithms for optimization without derivatives, *Mathematics Today* **43**, 170 (2007).
- [61] S. Ruder, An overview of gradient descent optimization algorithms, [arXiv:1609.04747](https://arxiv.org/abs/1609.04747).
- [62] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, C. Blank, K. McKiernan, and N. Killoran, PennyLane: Automatic differentiation of hybrid quantum-classical computations, [arXiv:1811.04968](https://arxiv.org/abs/1811.04968).
- [63] D. Dervovic, M. Herbster, P. Mountney, S. Severini, N. Usher, and L. Wossnig, Quantum linear systems algorithms: a primer, [arXiv:1802.08227](https://arxiv.org/abs/1802.08227).
- [64] R. Majumdar, D. Madan, D. Bhoumik, D. Vinayagamurthy, S. Raghunathan, and S. Sur-Kolay, Optimizing ansatz design in QAOA for Max-Cut, [arXiv:2106.02812](https://arxiv.org/abs/2106.02812).