

Universal Effectiveness of High-Depth Circuits in Variational Eigenproblems

Joonho Kim,^{1,*} Jaedeok Kim,^{2,*} and Dario Rosa^{3,4,*}

¹*School of Natural Sciences, Institute for Advanced Study, Princeton, NJ 08540, USA[†]*

²*AI Center, Samsung Research, Seoul 06765, Republic of Korea*

³*School of Physics, Korea Institute for Advanced Study,*

85 Hoegi-ro, Dongdaemun-gu, Seoul 02455, Republic of Korea

⁴*Department of Physics, Korea Advanced Institute of Science and Technology,*
291 Daehak-ro, Yuseong-gu, Daejeon 34141, Republic of Korea

We explore the effectiveness of variational quantum circuits in simulating the ground states of quantum many-body Hamiltonians. We show that generic high-depth circuits, performing a sequence of layer unitaries of the same form, can accurately approximate the desired states. We demonstrate their universal success by using two Hamiltonian systems with very different properties: the transverse field Ising model and the Sachdev-Ye-Kitaev model. The energy landscape of the high-depth circuits has a proper structure for the gradient-based optimization, *i.e.* the presence of local extrema – near any random initial points – reaching the ground level energy. We further test the circuit’s capability of replicating random quantum states by minimizing the Euclidean distance.

I. INTRODUCTION

Variational Quantum Eigensolver (VQE) [1, 2] is one of the most promising hybrid quantum-classical (HQC) algorithms, which may offer a precise approximation of the ground state of quantum systems. It is based on the iterative application of the following three steps: state preparation, measurement and optimization. Let us briefly describe each step. First, the preparation of a trial state $|\psi(\boldsymbol{\theta})\rangle$ is carried out by successive application of unitary quantum gates that depend on variational parameters $\boldsymbol{\theta}$. Second, the measurement step estimates the trial state mean energy,

$$E(\boldsymbol{\theta}) \equiv \langle \psi(\boldsymbol{\theta}) | \mathcal{H} | \psi(\boldsymbol{\theta}) \rangle, \quad (1)$$

by taking the expectation value of Hamiltonian \mathcal{H} of the target system over the trial state. Third, the optimization step adjusts the variational parameters $\boldsymbol{\theta}$ of the trial wavefunction to minimize the mean state energy $E(\boldsymbol{\theta})$ by applying a classical optimization algorithm. See [3] and references therein for more details. After a sufficient number of iterations, the variational state $|\psi(\boldsymbol{\theta}^*)\rangle$ at a convergence point $\boldsymbol{\theta} = \boldsymbol{\theta}^*$ is expected to reproduce well the ground state of the target Hamiltonian \mathcal{H} , under the assumption that the state Ansatz $|\psi(\boldsymbol{\theta})\rangle$ is *parametrically expressible* and *well-trainable* under the gradient based optimization. It is then a crucial question to find such an Ansatz.

Ideally, we are looking for a universally effective Ansatz capable of solving the VQE problem associated with an *arbitrary* target Hamiltonian. Note that it is generically a very challenging task: unless the target Hamiltonian is local, the ground states of non-local interacting systems are relatively close to typical quantum states that constitute most of the Hilbert space. Universality is thus

equivalent to demand a circuit Ansatz to approximate *any* random state $|\phi\rangle$ at certain values of the circuit parameters.

The existence of a parameter set $\boldsymbol{\varphi}$, such that $|\psi(\boldsymbol{\varphi})\rangle \simeq |\phi\rangle$, is therefore a necessary condition for the effectiveness of the variational circuit. On top of that, we must concern if the gradient descent optimization can actually reach the parameters $\boldsymbol{\varphi}$ from randomly initialized ones. Suppose that we use a layered quantum circuit composed of repetitive application of variational layers with the same architecture. Since increasing the number of layers extends the dimension of the parameter space, it can affect the circuit’s expressibility only positively. However, there has been a reported tension between increasing depth of layers and trainability of the circuit via minimizing the mean energy, (1), known as the barren plateau phenomenon [4].

When the layered circuit reaches a certain depth such that it evolves to an approximate 2-design, a numerical experiment [4] has shown the exponential decay of the variance of energy gradients $\nabla_{\boldsymbol{\theta}} E(\boldsymbol{\theta})$ with respect to the number n of qubits — for random circuit states obtained by uniformly sampling from the parameter space. Combined with another fact that the random energy gradient vanishes on average, Chebyshev’s inequality implies that the initial energy gradient can exponentially rarely deviate from zero. Such diminishing gradients hinder the beginning of efficient energy minimization, possibly causing the variational state to be stuck on non-optimal plateaus.

There have been several proposals to overcome the vanishing gradient problem and optimize the variational circuit. The most obvious approach is to incorporate some physics information on the target Hamiltonian [5–8], *e.g.* the ground state symmetry, for designing a less generic problem-tailored Ansatz. As an alternative direction towards the universal applicability of the variational algorithms, novel initialization [9, 10], architecture [11, 12], and optimization [13–15] of generic purpose circuits have been developed to enhance the circuit performance.

* All the authors contributed equally

[†] Correspondence to: joonhokim@ias.edu

The main goal of this paper is to demonstrate the inherent capability of variational circuits as universal and accurate eigensolvers for generic Hamiltonians – if they can even approximate close-to-random states. To this end, we will simply put aside the barren plateau problem by sufficiently increasing the classical computation capacity. Specifically, we will consider the high-depth regime of the layered quantum circuit, thus building an exponentially high-dimensional parameter space. In this regime, the variational circuit has sufficiently many layers and consequently many parameters. Hence, the norm of the gradient vector $\|\nabla_{\boldsymbol{\theta}} E(\boldsymbol{\theta})\|$ can still grow to a finite size, capable of moving the circuit parameter $\boldsymbol{\theta}$ from initial points, even though the magnitudes of the individual components $|\partial_i E(\boldsymbol{\theta})|$ are exponentially suppressed for the number n of qubits. We will illustrate the effectiveness of the high-depth circuits by solving two concrete VQE problems for the following quantum many-body Hamiltonians: the 1d Ising model in a constant transverse magnetic field, which is a prototypical model of locally interacting spin-chain systems, and the Sachdev-Ye-Kitaev (SYK) model, which is a strongly interacting quantum mechanical system of Majorana fermions [16–18]. Despite the striking contrast in these two Hamiltonians’ ground state properties, we will find that the high-depth circuit achieves a very high fidelity in approximating both ground states. Moreover, we will see that the high-depth circuits can narrow the Euclidean distance from random states up to arbitrary precision, indicating outstanding expressibility and trainability that stems from the high-dimensionality of $\boldsymbol{\theta}$.

The remarkable efficiency of the gradient descent optimization in the *over-parameterized* regime was also reported for approximating random unitary matrices [19] and ground states [8] with certain variational Ansatzes. More generally, over-parameterization is an active research topic that underlies the success of deep learning. It makes large neural networks capable to reach a global minimum during the optimization process, despite the non-convexity of the energy landscape [20, 21]. While searching the ground state via the energy minimization, we will observe some phenomena similar to what happens during the training of the neural networks with large parameter spaces, summarized as follows:

First, the energy landscape looks fairly simple in the local vicinity of randomly chosen points [15]. Generically, an initial point is already confined in a certain basin of attraction, such that an emanating optimization trajectory can quickly arrive at a nearby local extremum. Especially if the circuit has enough layers to become an approximate 2-design, almost all uniformly sampled initial states end up with rather homogeneous energy levels.

Second, for high-depth circuits, all the local extrema in the energy landscape, reachable by the VQE optimization from randomly initialized points, are substantially close to the exact ground energy, i.e., the value of the global minimum. These minima are not isolated individ-

ual points but develop multiple flat directions. It explains the robust success of the high-depth circuits in solving the VQE problems.

The rest of this paper is organized as follows. Section II introduces the architecture of the variational circuits used throughout this paper. Section III first reviews the occurrence of the barren plateau phenomenon, then examines the optimization of the variational circuits using the VQE example of the 1d Ising model coupled to a uniform transverse magnetic field. Section IV studies the same VQE problem for the SYK model, thus showing the universal effectiveness of the high-depth circuit. Section V addresses the ability of the variational circuits to reproduce typical states by minimizing the Euclidean distance. In particular, it shows that the high-depth Ansatz is highly expressible and trainable, being able to reach *any* random state. Finally, Section VI concludes with discussions.

II. CIRCUIT ANSATZES

Our focus in this work is to demonstrate the efficiency of high-depth layered circuits in a typical VQE problem, *i.e.* to approximate the ground state of a given Hamiltonian. To this purpose, here we specify the architecture of the variational circuit used in our numerical experiments. Our circuit state $|\psi(\boldsymbol{\theta})\rangle$ is composed of L unitary layers acting sequentially on the initial product state $|0\rangle$, *i.e.*,

$$|\psi(\boldsymbol{\theta})\rangle = U_L(\boldsymbol{\theta}_L) U_{L-1}(\boldsymbol{\theta}_{L-1}) \cdots U_2(\boldsymbol{\theta}_2) U_1(\boldsymbol{\theta}_1) |0\rangle, \quad (2)$$

where each layer $U_i(\boldsymbol{\theta}_i)$ has single-qubit y -rotation gates, parameterized by n periodic variables $\boldsymbol{\theta}_i$, and controlled- z gates operating on all pairs of n qubits. More precisely, once the single-qubit RY gates have acted upon all individual qubits, the entangling CZ gates acting on a -th controlling and b -th targeting qubits are arranged for every integer pair (a, b) satisfying $1 \leq a < b \leq n$. We have drawn the circuit state (2) with $n = 4$ qubits in Figure 1.

Note that we have deliberately chosen an unbiased circuit architecture, instead of embedding known properties of the ground states that are under search. Nonetheless, as we will see in Sections III, IV and V, the above circuit can solve the ground states of distinct Hamiltonians by minimizing $E(\boldsymbol{\theta})$ and accurately approximate random quantum states by minimizing the Euclidean distance, as long as the number L of layers is sufficiently large.

III. LOOKING INTO VQE TRAJECTORIES

This section is devoted to the detailed investigation of the VQE optimization procedure, with a particular focus on the effectiveness of the high-depth circuit Ansatz (2). Our exploration will be based on a concrete Hamiltonian system, commonly used in measuring the performance of




FIG. 1. The layered circuit Ansatz $|\psi(\boldsymbol{\theta})\rangle$ used in this paper.

variational circuits, *i.e.* the 1d Ising model in a transverse and uniform magnetic field [8, 22].

The 1d transverse field Ising model is defined over a spin lattice of length n , consisting of the spin-spin coupling between nearest neighbors in the z direction, as well as the spin interaction with a background uniform magnetic field along the transverse x direction. Assuming periodic boundary conditions, $\sigma_{n+1}^z \equiv \sigma_1^z$, the ferromagnetic Ising Hamiltonian reads

$$\mathcal{H} \equiv - \sum_{i=1}^n \sigma_i^z \sigma_{i+1}^z - g \sum_i \sigma_i^x, \quad (3)$$

where $\sigma_i^{x,y,z}$ is the Pauli operator acting on the i 'th spin, and g denotes the strength of the uniform magnetic field.

The physics of this model has been well-studied [23]. When the lattice size scales up to infinity, $n \rightarrow \infty$, the system undergoes a *quantum* phase transition at $|g| = 1$ between the ordered ($|g| < 1$) and disordered ($|g| > 1$) phases. The former phase has the spin-flip \mathbb{Z}_2 symmetry that connects the two opposite ferromagnetic ground states, while the latter phase has a unique ground state, with all the spins aligned along the x direction.

We will apply the VQE algorithm to the finite n Ising system, which exhibits some differences from the thermodynamic limit. Specifically, the \mathbb{Z}_2 degeneracy in the $0 < |g| < 1$ phase is broken at finite n . Our target state will be always non-degenerate and gapped at $|g| \neq 0$. All concrete calculations presented in this section have been done for $g = 2$.

1. Barren Plateaus and Classical Resolution

It was argued in [4] that optimization of the quantum variational circuits under the random parameter initialization comes with an inherent difficulty, given by the




FIG. 2. $\text{Tr}(\mathcal{H}^2)$ of the Ising model (3) and the SYK model (20) as a function of the system size n . The dashed lines are the regression lines with the functional form of $a \cdot n^b 2^n$, based on the numerical data denoted as small circles.

problem of vanishing gradients. For the circuit unitary ensemble that is quantum 2-design, random initial parameters are typically located on a plateau in the energy landscape. It means that the gradient-based optimization cannot even roll out. The odds of having non-vanishing gradients at random points decays exponentially with the system size n , impeding a large-scale application of the variational circuits. Its consequence would be even more detrimental on actual quantum devices, where the sampling noise must be taken into account [4, 14]. Given its prominence, we begin by reviewing the vanishing gradient problem in the VQE setting.

Consider the ensemble of the energy gradients over the parameter space of the variational circuit in Figure 1. For analyzing the k 'th component, $\partial_k E(\boldsymbol{\theta})$, that belongs to the ℓ 'th variational layer, it is convenient to group the L layer unitaries (2) into the following two blocks:

$$|\psi(\boldsymbol{\theta})\rangle = U_-(\boldsymbol{\theta}_-) U_+(\boldsymbol{\theta}_+) |0\rangle, \quad (4)$$

where $\boldsymbol{\theta}_q = \{\theta_p \mid \lfloor \frac{p}{n} \rfloor = q\}$, $\boldsymbol{\theta}_+ = \bigcup_{a \leq \ell} \boldsymbol{\theta}_a$, $\boldsymbol{\theta}_- = \bigcup_{a > \ell} \boldsymbol{\theta}_a$,

$$\begin{aligned} U_+(\boldsymbol{\theta}_+) &\equiv U_\ell(\boldsymbol{\theta}_\ell) U_{\ell-1}(\boldsymbol{\theta}_{\ell-1}) \cdots U_2(\boldsymbol{\theta}_2) U_1(\boldsymbol{\theta}_1), \\ U_-(\boldsymbol{\theta}_-) &\equiv U_L(\boldsymbol{\theta}_L) \cdots U_{\ell+2}(\boldsymbol{\theta}_{\ell+2}) U_{\ell+1}(\boldsymbol{\theta}_{\ell+1}). \end{aligned} \quad (5)$$

As the partial derivative ∂_k acts only on $U_\ell(\boldsymbol{\theta}_\ell)$, the variance of the k 'th gradient component, $\text{Var}_\theta[\partial_k E(\boldsymbol{\theta})]$, is

$$- \int \frac{d\boldsymbol{\theta}}{(2\pi)^{nL}} \langle \psi(\boldsymbol{\theta}) | [U_-(\boldsymbol{\theta}_-) V_k U_-(\boldsymbol{\theta}_-)^{\dagger}, \mathcal{H}] | \psi(\boldsymbol{\theta}) \rangle^2 \quad (6)$$

where V_k denotes the Pauli operator σ^y acting on the k 'th spin variable, such that $\text{Tr}(V_k) = 0$ and $\text{Tr}(V_k^2) = 2^n$.

If we assume the quantum 2-design property of $U_+(\boldsymbol{\theta}_+)$ and/or $U_-(\boldsymbol{\theta}_-)$, the above integral (6) can be replaced with the unitary matrix integral. In that case, the matrix



(a) Sample variance of the energy derivative $\partial_k E(\theta)$ (b) Sample average of the Euclidean norm of $\nabla_\theta E(\theta)$

FIG. 3. The barren plateau experiment for the Ising Hamiltonian (3). The shades indicate (a) the variance across gradient components $\{\partial_k E(\theta)\}_{k=1}^{nL}$, (b) the first and third sample quantiles.

integral can be handled exactly and simplified to

$$\frac{2\text{Tr}(\mathcal{H}^2)}{2^{2n}}. \quad (7)$$

Such simplification (7) happens when both $U_\pm(\theta_\pm)$ are 2-designs. Instead, if the 2-design condition does not hold for either of $U_\pm(\theta_\pm)$, we have the following expression:

$$-\frac{1}{2^{2n}} \int \frac{d\theta_-}{(2\pi)^{n(L-\ell)}} \text{Tr}([V_k, U_-(\theta_-)^\dagger \mathcal{H} U_-(\theta_-)]^2) \quad (8)$$

if only $U_+(\theta_+)$ is a 2-design while $U_-(\theta_-)$ is not, or

$$-\frac{\text{Tr}(\mathcal{H}^2)}{2^{2n}} \int \frac{d\theta_+}{(2\pi)^{n\ell}} \text{Tr}([V_k, U_+(\theta_+)^\dagger \rho U_+(\theta_+)]^2) \quad (9)$$

where $\rho = |0\rangle\langle 0|$, if not $U_+(\theta_+)$ but only $U_-(\theta_-)$ is a 2-design. Asymptotically in the system size n , the above expressions (7)–(9) are all bounded as

$$0 \leq \text{Var}_\theta[\partial_k E(\theta)] \leq \frac{4\text{Tr}(\mathcal{H}^2)}{2^{2n}}, \quad (10)$$

where we find the upper bound by expanding the commutator inside the integral of (8) and (9), then applying the following trace inequality that holds for two Hermitian matrices A, B [24]:

$$|\text{Tr}(AB)^{2m}| \leq \text{Tr}(A^{2m} B^{2m}) \quad \text{for } m \in \mathbb{N}. \quad (11)$$

To determine the scaling behavior of the upper bound (10) with respect to the system size n , we examine how $\text{Tr}(\mathcal{H}^2)$ scales with n by extrapolating the values obtained from exact diagonalization of the Hamiltonian (3) up to $n \leq 17$. The result is presented in Figure 2. We see that the term $\text{Tr}(\mathcal{H}^2)$ scales as 2^n , such that the exponential factor in the denominator of the upper bound (10) cannot

fully be balanced by $\text{Tr}(\mathcal{H}^2)$. Then, inserting the upper bound (10) into Chebyshev's inequality, which states

$$\Pr(|\partial_k E(\theta)| > \epsilon) < \frac{\text{Var}_\theta[\partial_k E(\theta)]}{\epsilon^2} \quad (12)$$

with $\partial_k E(\theta)$ being a random variable of zero mean, the probability of having a non-zero and finite derivative is exponentially suppressed with growing n . Therefore, the large-scale VQE problems are expected to suffer from the vanishing gradients.

We have not yet justified the assumption that either of the unitaries $U_\pm(\theta_\pm)$ is quantum 2-design. To see if the problem of exponentially vanishing gradients happens in the circuit unitaries of Figure 1, we numerically estimate the variance of initial gradients by collecting 10^3 random energy gradients of the Ising model. Figure 3a clearly exhibits an exponential decay of the partial derivatives with the increasing number n of qubits, where the shaded region displays component-wise fluctuations of the variance for all $1 \leq k \leq nL$. It shows that the energy landscape for the circuit in Figure 1 indeed contain the barren plateaus, which hinders the circuit optimization towards the ground state.

On the other hand, Figure 3a shows that the variance, at a fixed value of n , converges to a constant by increasing the number L of layers. The variance is independent of L beyond a transition point L_0 , where the circuit unitaries with $L \geq L_0$ evolve to approximate 2-design. Due to the saturation, having exponentially many parameters can compensate for the exponential decay of individual components when calculating the gradient norm $\|\nabla_\theta E(\theta)\|$, which appears in the evolution of the circuit energy under the gradient descent with an infinitesimal rate α :

$$\frac{dE(\theta)}{d\tau} = -\alpha \|\nabla_\theta E(\theta)\|^2 \quad (13)$$

Therefore, the vanishing gradient problem inherent to the quantum Hilbert space can be resolved in a classical way, *i.e.* by using the exponentially high-dimensional θ -space.

In agreement with the reasoning above, Figure 3b exhibits a small initial drop dominated by the transient decrease of $\text{Var}_\theta[\partial_k E(\theta)]$ for $L \leq L_0$ layers, then a steady increase driven by the linear growth in the number of circuit parameters. One can approximate the asymptotic increase rate of the norm $\|\nabla_\theta E(\theta)\|$ as

$$\|\nabla_\theta E(\theta)\| \sim \sqrt{nL \times \text{Var}_\theta[\partial_k E(\theta)]} \quad (14)$$

which agrees well with the numerically estimated growth rates between $\log \|\nabla_\theta E(\theta)\|$ and $\log L$ in Figure 3b:

n qubits	4	6	8	10
rate	0.504	0.502	0.503	0.501

We remark that the barren plateau phenomenon concerns only the initial steps of the gradient descent update. In almost all physical systems of relevance, the Hamiltonian spectrum is symmetrically distributed around a certain value, which we canonically set to zero. A random state is therefore in a superposition of multiple Hamiltonian eigenstates, whose mean energy is almost certainly zero. On the other hand, the variational mean energy quickly becomes a negative value after a few steps of the VQE optimization, implying that the circuit state is no longer represented by sample statistics of random states.

Having found that the vanishing gradient problem can be trivially avoided at the cost of having exponentially many parameters, we turn to answer if the variational circuit with a sufficiently many layers can actually solve the VQE problems. We will examine the optimization error, the training curve, and the trajectory in the parameter space at different values of the circuit depth L .

2. Optimizing the Circuit

We now come back to original task of approximating the ground state of the Ising model, Eq. (3).

The VQE optimization results of the circuit states (2), with different numbers L of layers and under the random initialization of the circuit parameters θ , can be summarized to the following two features:

1. For small enough L , the minimized energies $E(\theta^*)$ that the circuit states can reach are highly variable.
2. For larger L , the $E(\theta^*)$ distribution turns concentrated around a mean value which gets smaller.

We illustrate them with the outcomes of the VQE algorithm for the case of a chain having $n = 10$ sites and for $L \in \{8, 10, 12, 14, 16, 18, 20\}$ layers. To find the optimal parameters θ^* that minimize the energy, we use the Adam optimization algorithm [25], which iteratively




FIG. 4. Optimized VQE energy ($E(\theta^*) - E_0$) density for the Ising model over 35 distinct runs with random initialization.

updates the variational parameters θ by the exponential moving averages of gradients and their squares. It has a clear advantage in convergence speed, being widely used in a variety of deep learning models. The parameter update rule is decided by a choice of three hyperparameters $(\alpha, \beta_1, \beta_2)$. We have collected 35 independent VQE runs for the Ising Hamiltonian (3), whose initial parameters are randomly sampled from the uniform distribution $\mathcal{U}(0, 2\pi)^{\otimes nL}$, after 500 parameter updates with the hyperparameters $(\alpha, \beta_1, \beta_2) = (0.05, 0.9, 0.999)$.

The sample distribution of the final VQE energy $E(\theta^*)$ is visualized in Figure 4 for different numbers L of layers. One can characterize it as follows. First, the energy distribution at the local extremum θ^* clearly exhibits the widespread spectrum for a shallow depth, *e.g.*,



$$\begin{aligned} 3.52 \leq E(\theta^*) - E_0 \leq 12.6 & \quad \text{for } L = 8, \\ 4.67 \leq E(\theta^*) - E_0 \leq 16.9 & \quad \text{for } L = 10. \end{aligned}$$

sometimes achieving a relatively good energy level while the gap always persists. Second, by stacking more layers, *i.e.*, $L \geq 12$, deeper than the 2-design transition point detected in Figure 3a, the energy distribution starts to concentrate around a value, which is far from the ground energy E_0 . Third, the average value of $E(\theta^*)$ continues to decrease for the growing number L of layers, suggesting that the high-depth variational circuits can possibly simulate the ground state using the VQE optimization.

Encouraged by the observed shrinkage of the mean and variance of $E(\theta^*)$, we also have done the single VQE tryouts for a broader span of (n, L) , as summarized in Figure 5. It is clear that being deeper enhances the variational circuit's capability to replicate the ground state $|\psi_0\rangle$, reaching the ground state energy E_0 and achieving a high level of fidelity with $|\psi_0\rangle$. The high-depth circuits can achieve a zero error with the following precision:

$$|E(\theta) - E_0| \leq 10^{-5} \cdot \Delta E \quad (15)$$

where ΔE denotes the bandwidth of the target Hamiltonian \mathcal{H} , defined as the difference between the largest and smallest eigenvalues of \mathcal{H} . We note that such precision can be achieved only with an appropriately chosen learning rate α , in order to avoid too-large parameter updates that prevent the fine-level optimization. Here we refrain from the systematic hyperparameter search, which may

(a) The VQE optimization error $E(\theta^*) - E_0$ 

(b) Fidelity between the optimized and true ground states

FIG. 5. Single-run VQE outcomes for the Ising model (3) using the layered circuit Ansatz (2) at different depths L .

be more relevant for the case where the average gap between nearby energy levels shrinks, but simply stick to $\alpha = 0.05$ ($L = 4, 6, 8$) and $\alpha = 0.01$ ($L = 10$). We have found that (15) can be achieved when the circuit depth L passes through the following threshold value:

n qubits	4	6	8	10
L_v layers	10	24	68	250

We also remark that deeper circuits do not necessarily lead to better performance with VQE, as displayed by the gentle ramp after passing the threshold point L_v . This can be understood as follows: As the space of variational parameters θ has more dimensions, the basin of attractor to local extrema becomes narrower [11], giving a larger value of the estimated inverse volume [26]

$$V_k^{-1} = \sum_{i=1}^k \log \lambda_i(H) \quad (16)$$

where the summation is taken over the top- k eigenvalues $\{\lambda_i(H)\}_{i=1}^k$ of the Hessian matrix $H_{ij} \equiv \partial_i \partial_j E(\theta)$. For

instance, the positive correlation between L and $V_k^{-1}(L)$ can clearly be identified in Figure 7a, drawn for $k = 100$. As a result, the VQE trajectory is unlikely to land at an exact extremum but wander around nearby points, whose deviation gets larger as the attractor basin becomes narrower and steeper.

Even with the optimal number of layers, such that the circuit state can reach an exact extremum and accurately represent the ground state $|\psi_0\rangle$ of the Ising Hamiltonian (3), the narrowness of the attractor basin still makes the VQE trajectory somewhat unstoppable, passing through the best point θ^* and then hopping around in the local neighborhood. Figure 6a shows that the VQE error $E(\theta_\tau) - E_0$ slightly increases on average and mildly fluctuates after achieving the minimum error $E(\theta^*) - E_0$. This residual error can be reduced by making use of popular optimization tricks, such as early stopping or learning rate scheduling, which causes the VQE optimization to stop at the optimal point. For instance, by introducing the exponential decay of the learning rate α , *i.e.*

$$\alpha_\tau = \alpha_0 c^{\tau/500}, \quad \tau \geq 0 \quad (17)$$

at the optimization step τ with a constant value $c = 0.3$, we could reduce the late time fluctuations as in Figure 6b.

So far, we have discussed some important aspects of the VQE optimization. The main observation is that, when supported by the high-dimensional parameter space, randomly initialized variational circuits can approximate the Ising ground state with remarkably high accuracy. The energy gradient neither vanishes nor randomly fluctuates along the optimization trajectory, making it quickly converge to a local minimum that exhibits a small energy gap from the ground energy E_0 and high fidelity with the exact ground state $|\psi_0\rangle$. We will explore this efficiency of the high-depth circuit in some details by visualizing the VQE trajectory on the energy landscape.

3. Visualizing the Trajectory

A key characteristic that contributes to the success of the high-depth circuit is the fact that randomly initialized points are likely to be already confined in the basin of attraction of a good attractor, *i.e.* a local extremum that is close enough to the ground state energy. We illustrate now this point by looking at the actual optimization trajectories under the energy minimization.

For a given initial point θ_0 and the trajectory $\mathcal{T}(\theta_0) = \{\theta_\tau\}$ thereafter, we identify the optimal parameter θ^* as the point in the trajectory $\mathcal{T}(\theta_0)$ having the minimum energy expectation value,

$$E(\theta^*) \leq E(\theta_\tau) \quad \text{for any } \tau \geq 0, \quad (18)$$

which is the best possible representative point of the local extrema that the trajectory $\mathcal{T}(\theta_0)$ converges around. The associated basin of attractor can be estimated by

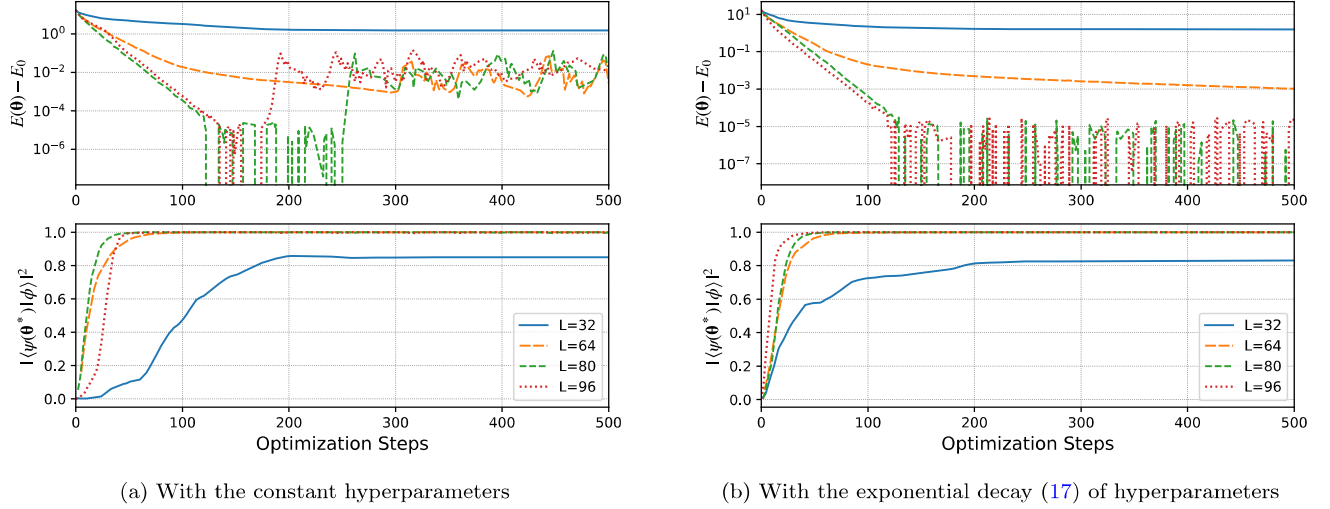
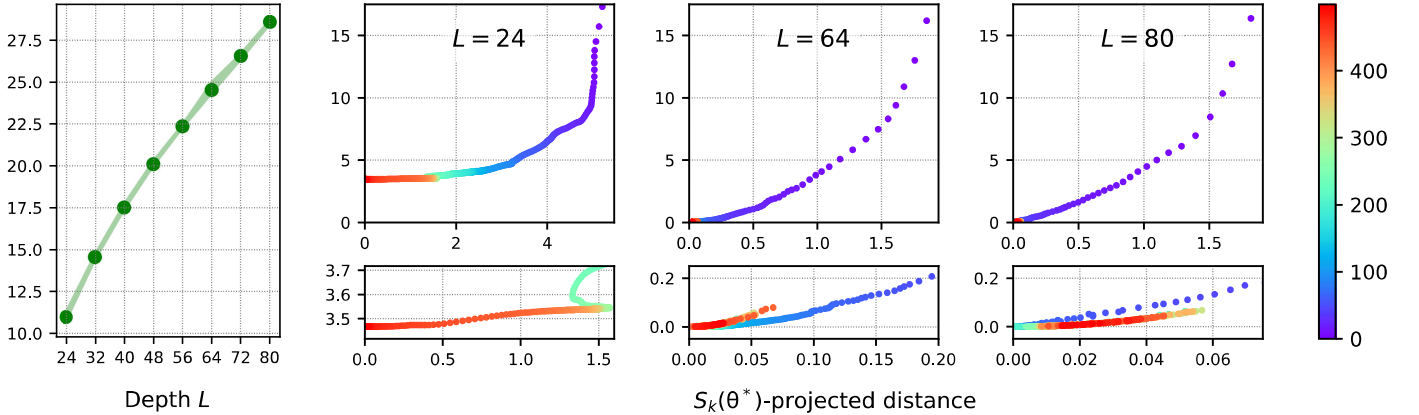


FIG. 6. Optimization curves of the circuit state on the Ising model at 8 qubits. The upper and lower plots denote the VQE error $E(\theta_\tau) - E_0$ and the fidelity $|\langle \psi(\theta^*) | \phi \rangle|^2$ between the circuit and true ground states, respectively, as a function of the optimization steps τ . Notice the late-time fluctuation has been alleviated with the learning rate scheduling.



(a) The mean of the top 100 Hessian eigenvalues (b) The horizontal and vertical axes denote the $\mathcal{S}_{100}(\theta^*)$ -projected Euclidean distance and the VQE error $E(\theta_\tau) - E_0$, respectively. The bottom boxes magnify the trajectory in the vicinity of the optimal point θ^* .

FIG. 7. The visualization of the optimization trajectory in the 100 steepest directions for the Ising model with $n = 8$ qubits.

calculating the Hessian $H_{ij} \equiv \partial_i \partial_j E(\theta^*)$ at the optimum θ^* . We are interested in the eigenvalue spectrum, $\{\lambda_i(H)\}$, from which we distinguish steep and flat directions and calculate the degree of steepness.

Similar to the case of deep neural networks [26–28], the local extrema in the VQE energy landscape of high-depth circuits are often interconnected in multiple flat directions, whose corresponding Hessian eigenvalues are zero, looking like valleys rather than isolated singular points. Success of the VQE algorithm is not affected by a specific position in flat directions, but only concerned with if a ball, initially at a considerable height, can roll off in steep directions and reach a sufficiently deep gorge. It motivates us to consider the k -dimensional

hypersurface $\mathcal{S}_k(\theta^*)$ along the k steepest directions, *i.e.*, spanned by the top- k Hessian eigenvectors. More precisely, we want to examine if the $\mathcal{S}_k(\theta^*)$ -projected Euclidean distance between θ^* and θ_τ decreases along the trajectory $\mathcal{T}(\theta_0)$ as the optimization step τ progresses, until $E(\theta_\tau) = E(\theta^*)$. This will tell us if the entire trajectory $\mathcal{T}(\theta_0)$ is confined in the k -dimensional basin of attraction, while ignoring the movement in the directions of less or zero attraction.

Figure 7b displays the VQE error $E(\theta_\tau) - E_0$ and the projection distance $\|\theta_\tau - \theta^*\|_{\mathcal{S}_{100}(\theta^*)}$ along the actual optimization trajectories, whose step number τ is indicated by color. We have made the following observations: First, when an appropriate value of k is selected, both the dis-

tance $\|\boldsymbol{\theta}_\tau - \boldsymbol{\theta}^*\|_{\mathcal{S}_k(\boldsymbol{\theta}^*)}$ and the loss value $E(\boldsymbol{\theta}_\tau) - E_0$ continuously decrease on a macroscopic scale. It exhibits that the trajectory $\mathcal{T}(\boldsymbol{\theta}_0)$ converges without escaping from a specific basin of the attractor that encloses a randomly initialized point $\boldsymbol{\theta}_0$. Second, for a shallow circuit state, the optimization trajectory often makes a slight detour in some orthogonal directions. In contrast, the steady convergence occurs typically for large L . It implies that the vicinity of any randomly initialized parameters is effectively convex. Finally, we find from Figure 7a that the attractor basin in the direction of $\mathcal{S}_k(\boldsymbol{\theta}^*)$ evolves rapidly steeper and narrower as the depth L increases, thereby causing the rapid convergence and substantial late-time fluctuation around $\boldsymbol{\theta}^*$. It is noticeable that the initial convergence follows a milder path than later fluctuations in the $L = 64$ case, while the convergence in the $L = 80$ case happens along a much steeper route. Taken together, these observations show the quick convergence of high-depth circuits [8, 19] under the VQE algorithm.

IV. SOLVING THE SYK MODEL

Our discussions so far have been based on just a particular Hamiltonian, (3), the Ising model in a transverse uniform magnetic field. Since the variational circuit that we use has no features particularly tailored for the Ising model, we expect it to closely replicate the ground states of other Hamiltonians in the high-depth regime. To check the generality of the prior discussions on the efficiency of the high-depth circuits, we will now solve the VQE problem, defined for another Hamiltonian of very distinct nature: the Sachdev-Ye-Kitaev (SYK) model.

1. The SYK Model

The SYK model [16–18] is built out of $2n$ Majorana fermions in $1d$, *i.e.* the operators γ_i , with $i = 1, \dots, 2n$, satisfying the following anti-commutation relations

$$\{\gamma_i, \gamma_j\} = \delta_{ij}, \quad (19)$$

where δ_{ij} denotes the Kronecker delta. The SYK Hamiltonian is an all-to-all Hamiltonian, which couples all the Majorana fermions together in a fully non-local fashion, consisting of the following q -body interaction terms with $q \geq 2$ being an even integer:

$$\mathcal{H} \equiv (i)^{q/2} \sum_{i_1 < \dots < i_q} J_{i_1 \dots i_q} \gamma_{i_1} \cdots \gamma_{i_q}, \quad (20)$$

where the coupling constants $J_{i_1 \dots i_q}$ are randomly sampled from the Gaussian distribution of mean 0 and variance

$$\langle J_{i_1 \dots i_q}^2 \rangle \equiv \frac{J^2 (q-1)!}{(2n)^{q-1}}, \quad (21)$$

where J^2 is a constant which we set to be equal to one. The model has recently attracted a widespread attention from different communities, due to some peculiar features it enjoys. It has been shown that when $q \geq 4$ the model is highly chaotic [17, 18, 29, 30], although solvable in the large n limit [17, 18], thus creating a perfect situation to study relevant questions on quantum chaos which are usually out of reach for other chaotic models. Moreover, the SYK model has intriguing connections with the physics of black holes and quantum gravity, promoting itself as an ideal candidate to address new questions on holography and the AdS/CFT correspondence.

We will focus our attention on the SYK model with $q = 4$. It has two notable features that can be a source of trouble for any eigensolver algorithms, regardless of being classical or quantum mechanical, whose goal is to reach the ground state. First, the energy spectrum of the SYK model is very dense, especially having a small energy gap between the ground state E_0 and all other excited states. Second, the SYK ground state is much less distinguishable from generic quantum states in the Hilbert space, supporting the volume law scaling of the entanglement entropy. Therefore, the VQE computation with the SYK model must be seen as a highly challenging benchmark [31].

As another characteristic, the SYK model is known to have two-fold degenerate eigenstates, if and only if $n = 4k + 2$ for any positive integer k [29, 30]. Denoting the two-fold degenerate ground states by $|\phi_1\rangle$ and $|\phi_2\rangle$, that we assume to be normalized and orthogonal to each other, the VQE target states for the SYK model are then given by all the possible linear combinations of the form:

$$|\phi_0\rangle = \frac{\alpha}{\sqrt{|\alpha|^2 + |\beta|^2}} |\phi_1\rangle + \frac{\beta}{\sqrt{|\alpha|^2 + |\beta|^2}} |\phi_2\rangle \quad (22)$$

with α and β being complex numbers. Therefore, the distance between the circuit state $|\psi(\boldsymbol{\theta})\rangle$ and the closest state of the form $|\phi_0\rangle$ can be measured by computing

$$|\langle \psi(\boldsymbol{\theta}) | \phi_1 \rangle|^2 + |\langle \psi(\boldsymbol{\theta}) | \phi_2 \rangle|^2 \leq 1 \quad (23)$$

with the inequality which is saturated whenever $|\psi(\boldsymbol{\theta})\rangle$ takes exactly the form (22).

We will numerically show that the high-depth quantum circuit can effectively learn the ground state of the SYK model. It will imply that even complicated systems, involving non-local interactions and a high level of entanglement, can be universally simulated through the variational circuits.

2. Optimizing the Circuit

We start by discussing if the SYK Hamiltonian causes the vanishing gradient problem for the variational circuit in Figure 1. When at least one of $U_\pm(\boldsymbol{\theta}_\pm)$ is a 2-design, the scaling behavior of $\text{Tr}(\mathcal{H}^2)$ will determine if the SYK

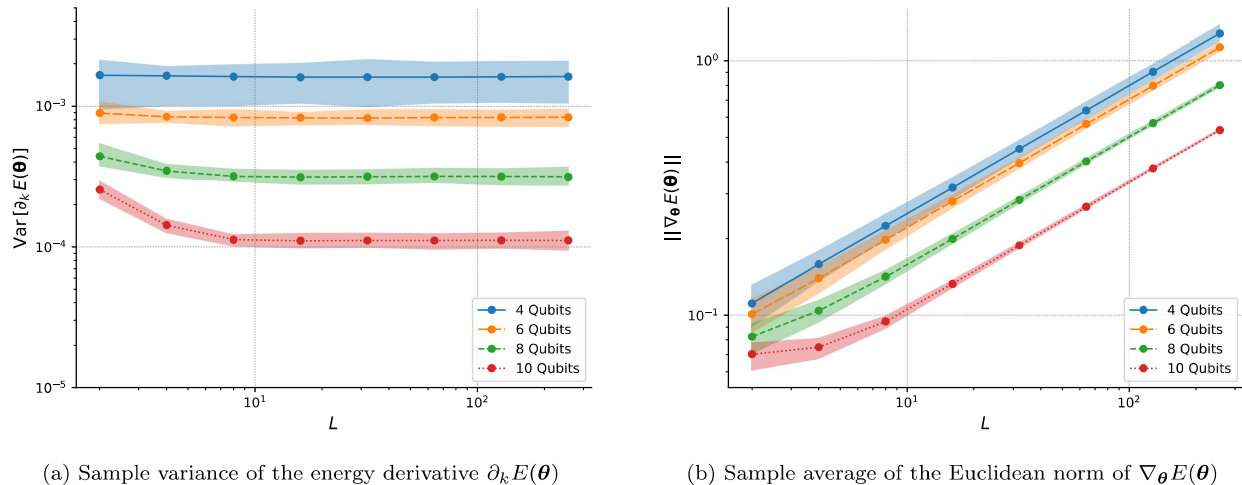


FIG. 8. The barren plateau experiment for the SYK Hamiltonian (20). The shades indicate (a) the variance across gradient components $\{\partial_k E(\theta)\}_{k=1}^{nL}$, (b) the first and third sample quantiles.

energy gradient will be exponentially suppressed or not. The Hamiltonian (20) has been exactly diagonalized up to $n = 15$. Moreover, an approximate analytical formula of the spectral density is at disposal in [32], which we use to extrapolate $\text{Tr}(\mathcal{H}^2)$ for the large system size n . The results are presented in Figure 2. We clearly see, as in the Ising model, that $\text{Tr}(\mathcal{H}^2)$ scales like 2^n , which indicates the abundance of the barren plateaus in the VQE energy landscape of the SYK model.

The vanishing gradient problem has also been observed numerically by computing the sample variance of $\partial_k E(\theta)$ over a collection of 1000 random parameters, as displayed in Figure 8a. Clearly, the energy gradient $\partial_k E(\theta)$ is exponentially suppressed by increasing n . We also see that for the SYK model, contrary to the Ising Hamiltonian, there is almost no transient regime where $\partial_k E(\theta)$ is still large, yet decreasing for the growing number L of layers. The lack of the transient regime, which is a consequence of the non-local nature of the SYK interactions [11], is a clear obstacle in using the variational circuit to approximate the SYK ground state in the low-depth regime.

On the other hand, as we can see from Figure 8b, the norm of the gradient vector is again increasing for the growing number L of layers, due to the saturation of the $\text{Var}_{\theta}[\partial_k E(\theta)]$ with respect to L . The empirically measured growth rates between $\log \|\nabla_{\theta} E(\theta)\|$ and $\log L$ are

n qubits	4	6	8	10
rate	0.503	0.503	0.501	0.502

matching the simple estimation formula (14). It tells that the gradient-based optimization can at least launch for high-depth quantum circuits, which bypass the vanishing gradient problem using an exponentially high-dimensional parameter space.

As the next step, we examine the performance of the variational circuit Ansatz (2) in reaching the ground state






FIG. 9. Optimized VQE energy ($E(\theta^*) - E_0$) density for the SYK model over 35 distinct runs with random initialization.

of the SYK Hamiltonian, by repeating the same numerical experiments conducted in Section III 2 under the same choice of hyperparameters $(\alpha, \beta_1, \beta_2) = (0.05, 0.9, 0.999)$.

First, the sample distribution of the minimized energy $E(\theta^*)$ for randomly initialized circuits with $4 \leq L \leq 16$ layers is illustrated in Figure 9, based on 35 sample VQE runs at each L . A notable observation is the small variance of the final energy, compared to the Ising VQE energy distribution in Figure 4, even at very low depth. We interpret it as another manifestation of the fact that the ground state that we aim to approximate is less distinguishable from generic quantum states. Hence, developing an optimization trajectory demands more classical computing power through the high-dimensional θ -space [31]. Another – and perhaps the most important – point to stress is that, the mean value of the minimized circuit energy $E(\theta^*)$ decreases by stacking more layers, just like what has happened to the Ising VQE problem. It suggests that the high-depth circuits can reach a very good approximation of the SYK ground states.

Second, we have measured the performance of the layered circuit Ansatz (2) in approximating a ground state of the SYK model (20), as a function of the circuit depth. Specifically, the VQE single-run error $E(\theta^*) - E_0$ is drawn

(a) The VQE optimization error $E(\theta^*) - E_0$ (b) Fidelity between the optimized and true ground states. As explained near (22)–(23), for the two-fold degenerate case $n = 4k + 2$, the sum $|\langle \psi(\theta^*) | \phi_1 \rangle|^2 + |\langle \psi(\theta^*) | \phi_2 \rangle|^2$ is drawn.FIG. 10. Single-run VQE outcomes for the SYK model (20) using the layered circuit Ansatz (2) at different depths L .

in Figure 10. The high-depth circuit performs very well, reaching a zero error with the following accuracy,

$$|E(\theta) - E_0| \leq 10^{-5} \cdot \Delta E,$$

when the depth L arrives at the following values.

n qubits	4	6	8	10
L_v layers	12	30	96	220

We also note from Figure 10b that the fidelity between the optimized and ground states tends to decrease at an intermediate scale of depth, while the VQE error continues to reduce without temporary increase. Such contrasting behavior is due to the dense energy spectrum of the SYK model near the ground energy level E_0 . With an intermediate-depth circuit, the VQE algorithm is going to approximate not the exact ground states, but some low-lying excited states. In this way, the energy continues to decrease but the fidelity does not improve. However, for sufficiently deep circuits, the VQE algorithm can

overcome the difficulty and reach an excellent agreement with the ground state.

Interestingly, we have seen that the necessary number L_v of layers to reach the high precision (15) is roughly in the same order, both for the SYK and Ising models. We will also see that the circuit (2) with $L \geq L_v$ achieves high precision in replicating random states by minimizing the Euclidean distance, as shown in Figure 11. This compatibility highlights the universal effectiveness of the high-depth circuit in approximating any quantum state in the Hilbert space, both *generic* and *non-generic* ones.

V. APPROXIMATING RANDOM STATES USING EUCLIDEAN LOSS FUNCTION

Taking one step further, we will look into the high-depth circuit's capability to approximate generic quantum states $|\phi\rangle$ by minimizing the Euclidean distance.

Recall that the variational circuit $|\psi(\theta)\rangle$ with L layers depends on nL parameters, which are periodic over the finite interval $[0, 2\pi)$ up to an overall sign. It is a mapping from the nL -dimensional torus to the Hilbert space of n qubits. As a practical measure for the expressibility and trainability of the circuits in simulating quantum states, we consider whether, for a given quantum state $|\phi\rangle$, there exist a point θ^* in the parameter space, reachable by the gradient-based optimization such that $|\psi(\theta^*)\rangle \simeq |\phi\rangle$. Such measure of expressibility is motivated by [33] but tailored for the hybrid algorithms. Its definition follows:

Let us apply the gradient descent to find the minimum distance at the closest point

$$\theta^* = \arg \min_{\theta} \| |\psi(\theta)\rangle - |\phi\rangle \| \quad (24)$$

between the circuit and target states, where $\|\cdot\|$ is the Euclidean norm of a complex vector. The (in)expressibility of the variational circuit is written as the average of the minimum distances over all quantum states:

$$\varepsilon = \frac{1}{\text{Vol}(U(2^n))} \int d\mu_{\phi} \min_{\theta} \| |\psi(\theta)\rangle - |\phi\rangle \|. \quad (25)$$

Notice that the closeness between two quantum states is usually defined by using the trace distance, which is highly sensitive to small parameter changes. Instead, we have adopted the Euclidean distance to improve the convexity of the optimization landscape, such that the closest point θ^* can be found as easily as possible. The Euclidean norm defines trivially a convex landscape. Hence, any non-convexity of the optimization landscape is inherited from the variational circuit itself that fixes how to embed the nL -dimensional torus into the Hilbert space.

For an actual estimation of ε , we substitute the Haar unitary integral (25) with the sample mean over m states:

$$\varepsilon_m = \frac{1}{m} \sum_{i=1}^m \min_{\theta} \| |\psi(\theta)\rangle - |\phi_i\rangle \| \quad (26)$$




FIG. 11. The parametric (in)expressibility of the layered circuit Ansatz over $m = 10$ random target samples, $\varepsilon_{m=10}/2^n$, normalized by the number 2^n of state components. The narrow shade denotes the fluctuation across $m = 10$ target states.

Given a target state $|\phi_i\rangle$, we start with an Ansatz state $|\psi(\theta)\rangle$ with nL initial parameters θ , randomly sampled from the uniform distribution $\mathcal{U}(0, 2\pi)^{\otimes nL}$. To reach the optimal parameters θ^* that minimize the distance, we use the Adam optimization algorithm with the hyperparameters $(\alpha, \beta_1, \beta_2) = (0.05, 0.9, 0.999)$. Note that the overwhelming majority of Hilbert space is filled with generic states with a high degree of entanglement, so the sample states $|\phi_i\rangle$ will almost never be non-generic, low-entangling states similar to the Ising ground state.

Figure 11 shows the sample mean ε_m over $m = 10$ random target states, divided by the number 2^n of state components. The narrow shade displays the fluctuation range of the component-averaged distance $\| |\psi(\theta^*)\rangle - |\phi_i\rangle \| \cdot 2^{-n}$ across different target states indexed by $i = 1, \dots, 10$. We have selected the minimum distance among the first $\tau \leq 500$ optimization steps, which is sufficient since the optimal parameter θ^* is empirically always found in an early stage of the optimization.

We find that the (in)expressibility approaches to 0 as the circuit depth L grows. To achieve the following high precision,

$$\varepsilon_{m=10} \leq 10^{-5} \cdot 2^n, \quad (27)$$

the depth L has to be greater than the threshold values:

n qubits	4	6	8	10
L_ε layers	10	24	56	150

Note that the slight ramp-up after the steep fall of $\varepsilon_{m=10}$ is caused by the fluctuation of θ around the optimal point θ^* , since the fluctuation size gets amplified for bigger L . It is a common phenomenon in gradient-based optimization, where a suitable reduction of the learning rate can turn it to the stable convergence $\theta \rightarrow \theta^*$.

In Figure 11, we have found the following overall trend: Deeper circuits are not only a superset of shallow circuits but also behave more effectively in the gradient-based optimization. Though the functional form of the variational circuit in Figure 1 is composed of only two types of gates, it can accurately express and reach most quantum states in Hilbert space in the high-depth regime.

VI. DISCUSSION

From the viewpoint of the vanishing gradient problem, adding more parametric gates to the circuit architecture brings both positive and negative effects on trainability. On the negative side, it makes the circuit unitary ensemble closer to quantum 2-design, whose one- and two-point correlators are equivalent to those of Haar random unitaries, so that initial random gradients may decay exponentially with the system size n [4]. On the positive side, however, it enlarges the dimensionality of the parameter space and maintains the norm of random gradients to be finite.

Until now, we have explored one limiting regime, where the impact of exponentially many parameters dominates over the barren plateau phenomenon. It turns out that the high-depth circuit has been very effective in solving the ground states of the Ising and SYK Hamiltonians, as well as in simulating generic random states. One may note a certain degree of qualitative similarity between the VQE optimization trajectory of high-depth circuits, demonstrated in Section III 2, and the lazy learning [34] in over-parameterized neural networks. We speculate it as naturally emerging in any systems involving the high-dimensional parameter space. It would be interesting to know why local extrema on the energy landscape of the high-depth circuits can reliably reach a zero VQE error, as studied in [35] for the neural networks.

Some interesting phenomena that we found during the optimization of low- and intermediate-depth circuits need to be further understood. For low-depth circuits, it is important to characterize what features of the initial points lead to the observed big difference in their minimized energy levels. We are also curious to know what makes the local minima on the energy landscape of intermediate-depth circuits so homogeneous. More generally, we want to understand how the circuit states evolve along a gradient optimization trajectory on average, in terms of various quantum information measures.

We need to concern two types of errors for the actual use of the variational circuits on the near-term quantum devices [36]. Firstly, for accurate estimation of the energy gradient, it is necessary to sample the variational wavefunction repeatedly, ideally exponentially many times in the system size n . Under the assumption that we can collect only a limited number of samples, the energy gradient estimation will be inevitably noisy. Therefore, the expressibility and trainability of the variational circuit

needs to be addressed with noisy gradients, or even without direct gradient computation as in [37, 38]. We expect the simplicity of the energy landscape in the high-depth regime may bring robustness against the sampling noise. Secondly, and more severely, quantum hardware noise restricts our ability to maintain quantum states when a sequence of layer unitary acts on them. A noise-induced mechanism that causes the vanishing gradient problem has also been studied in [39]. We would like to investigate if the variational circuit approach can be successful under decoherence of quantum states in the near future.

ACKNOWLEDGMENTS

We thank Boris Hanin and Yaron Oz for helpful discussions. Joonho Kim acknowledges the sup-

port from the NSF grant PHY-1911298 and the Si-vian fund. The work of DR is supported by the National Research Foundation of Korea (NRF) grant NRF-2020R1C1C1007591. Our Python scripts, which are available <https://github.com/phyjoon/high-depth-vqe>

for the numerical experiments written in JAX [40], QuTip [41]. The experimental data are managed in [Weights & Biases](#) [42].

-
- [1] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, *Nature Communications* **5** (2014), [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213), 1304.3061.
- [2] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, *New Journal of Physics* **18** (2016), [10.1088/1367-2630/18/2/023023](https://doi.org/10.1088/1367-2630/18/2/023023).
- [3] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, and et al., *Chemical Reviews* **119**, 10856–10915 (2019).
- [4] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, *Nature Communications* **9** (2018), [10.1038/s41467-018-07090-4](https://doi.org/10.1038/s41467-018-07090-4).
- [5] J.-G. Liu, Y.-H. Zhang, Y. Wan, and L. Wang, *Physical Review Research* **1** (2019), [10.1103/physrevresearch.1.023025](https://doi.org/10.1103/physrevresearch.1.023025).
- [6] B. T. Gard, L. Zhu, G. S. Barron, N. J. Mayhall, S. E. Economou, and E. Barnes, *npj Quantum Information* **6** (2020), [10.1038/s41534-019-0240-1](https://doi.org/10.1038/s41534-019-0240-1).
- [7] K. Seki, T. Shirakawa, and S. Yunoki, *Physical Review A* **101** (2020), [10.1103/physreva.101.052340](https://doi.org/10.1103/physreva.101.052340).
- [8] R. Wiersema, C. Zhou, Y. de Sereville, J. F. Carrasquilla, Y. B. Kim, and H. Yuen, “Exploring entanglement and optimization within the Hamiltonian Variational Ansatz,” (2020), [arXiv:2008.02941 \[quant-ph\]](https://arxiv.org/abs/2008.02941).
- [9] E. Grant, L. Wossnig, M. Ostaszewski, and M. Benedetti, *Quantum* **3**, 214 (2019).
- [10] G. Verdon, M. Broughton, J. R. McClean, K. J. Sung, R. Babbush, Z. Jiang, H. Neven, and M. Mohseni, “Learning to learn with quantum neural networks via classical neural networks,” (2019), [arXiv:1907.05415 \[quant-ph\]](https://arxiv.org/abs/1907.05415).
- [11] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, “Cost-Function-Dependent Barren Plateaus in Shallow Quantum Neural Networks,” (2020), [arXiv:2001.00550 \[quant-ph\]](https://arxiv.org/abs/2001.00550).
- [12] T. Volkoff and P. J. Coles, *Quantum Science and Technology* **6**, 025008 (2021).
- [13] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, *Quantum* **4**, 269 (2020).
- [14] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, and M. Leib, “Layerwise learning for quantum neural networks,” (2020), [arXiv:2006.14904 \[quant-ph\]](https://arxiv.org/abs/2006.14904).
- [15] B. Koczor and S. C. Benjamin, “Quantum Analytic Descent,” (2020), [arXiv:2008.13774 \[quant-ph\]](https://arxiv.org/abs/2008.13774).
- [16] S. Sachdev and J. Ye, *Physical Review Letters* **70**, 3339–3342 (1993).
- [17] A. Kitaev, “A simple model of quantum holography,” <http://online.kitp.ucsb.edu/online/entangled15/kitaev/> & <http://online.kitp.ucsb.edu/online/entangled15/kitaev2/> (2015).
- [18] J. Maldacena and D. Stanford, *Phys. Rev. D* **94**, 106002 (2016), [arXiv:1604.07818 \[hep-th\]](https://arxiv.org/abs/1604.07818).
- [19] B. T. Kiani, S. Lloyd, and R. Maity, “Learning Unitaries by Gradient Descent,” (2020), [arXiv:2001.11897 \[quant-ph\]](https://arxiv.org/abs/2001.11897).
- [20] M. Belkin, D. Hsu, S. Ma, and S. Mandal, (2019), [arXiv:1812.11118 \[stat.ML\]](https://arxiv.org/abs/1812.11118).
- [21] C. Liu, L. Zhu, and M. Belkin, (2020), [arXiv:2003.00307 \[cs.LG\]](https://arxiv.org/abs/2003.00307).
- [22] K. Nakaji and N. Yamamoto, “Expressibility of the alternating layered ansatz for quantum computation,” (2020), [arXiv:2005.12537 \[quant-ph\]](https://arxiv.org/abs/2005.12537).
- [23] S. Sachdev, *Quantum Phase Transitions*, 2nd ed. (Cambridge University Press, 2011).
- [24] Z. Yang and X. Feng, *JIPAM. Journal of Inequalities in Pure & Applied Mathematics [electronic only]* **3**, Paper No. 78, 12 p., electronic only (2002).
- [25] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” (2017), [arXiv:1412.6980 \[cs.LG\]](https://arxiv.org/abs/1412.6980).
- [26] L. Wu, Z. Zhu, and W. E, “Towards Understanding Generalization of Deep Learning: Perspective of Loss Landscapes,” (2017), [arXiv:1706.10239 \[cs.LG\]](https://arxiv.org/abs/1706.10239).
- [27] L. Sagun, L. Bottou, and Y. LeCun, “Eigenvalues of the Hessian in Deep Learning: Singularity and Beyond,” (2017), [arXiv:1611.07476 \[cs.LG\]](https://arxiv.org/abs/1611.07476).
- [28] P. Chaudhari, A. Choromanska, S. Soatto, Y. LeCun, C. Baldassi, C. Borgs, J. Chayes, L. Sagun, and R. Zecchina, “Entropy-SGD: Biasing Gradient Descent

- Into Wide Valleys,” (2017), [arXiv:1611.01838 \[cs.LG\]](#).
- [29] A. M. García-García and J. J. M. Verbaarschot, *Phys. Rev. D* **94**, 126010 (2016), [arXiv:1610.03816 \[hep-th\]](#).
- [30] J. S. Cotler, G. Gur-Ari, M. Hanada, J. Polchinski, P. Saad, S. H. Shenker, D. Stanford, A. Streicher, and M. Tezuka, *JHEP* **05**, 118 (2017), [Erratum: *JHEP* 09, 002 (2018)], [arXiv:1611.04650 \[hep-th\]](#).
- [31] J. Kim and Y. Oz, (2021), [arXiv:2102.12534 \[quant-ph\]](#).
- [32] A. M. García-García and J. J. M. Verbaarschot, *Phys. Rev. D* **96**, 066012 (2017), [arXiv:1701.06593 \[hep-th\]](#).
- [33] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, *Advanced Quantum Technologies* **2**, 1900070 (2019).
- [34] S. S. Du, X. Zhai, B. Poczos, and A. Singh, “Gradient Descent Provably Optimizes Over-parameterized Neural Networks,” (2019), [arXiv:1810.02054 \[cs.LG\]](#).
- [35] D. Soudry and Y. Carmon, “No bad local minima: Data independent training error guarantees for multilayer neural networks,” (2016), [arXiv:1605.08361 \[stat.ML\]](#).
- [36] J. Preskill, *Quantum* **2**, 79 (2018).
- [37] N. Maheswaranathan, L. Metz, G. Tucker, D. Choi, and J. Sohl-Dickstein (PMLR, Long Beach, California, USA, 2019) pp. 4264–4273.
- [38] S. Welleck and K. Cho, “MLE-guided parameter search for task loss minimization in neural sequence modeling,” (2020), [arXiv:2006.03158 \[cs.LG\]](#).
- [39] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, “Noise-Induced Barren Plateaus in Variational Quantum Algorithms,” (2020), [arXiv:2007.14384 \[quant-ph\]](#).
- [40] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, and S. Wanderman-Milne, “JAX: composable transformations of Python+NumPy programs,” (2018).
- [41] J. R. Johansson, P. Nation, and F. Nori, *Comput. Phys. Commun.* **184**, 1234 (2013).
- [42] L. Biewald, “Experiment Tracking with Weights and Biases,” (2020), software available from wandb.com.