

# Effect of barren plateaus on gradient-free optimization

Andrew Arrasmith<sup>1</sup>, M. Cerezo<sup>1,2</sup>, Piotr Czarnik<sup>1</sup>, Lukasz Cincio<sup>1</sup>, and Patrick J. Coles<sup>1</sup>

<sup>1</sup>Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA

<sup>2</sup>Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, NM, USA

Barren plateau landscapes correspond to gradients that vanish exponentially in the number of qubits. Such landscapes have been demonstrated for variational quantum algorithms and quantum neural networks with deep circuits, global cost functions, large entanglement, or hardware noise. For obvious reasons, it is expected that gradient-based optimizers will be significantly affected by barren plateaus. However, whether or not gradient-free optimizers are impacted is a topic of debate, with some arguing that gradient-free approaches are unaffected by barren plateaus. Here we show that, indeed, gradient-free optimizers do not solve the barren plateau problem. Our main result proves that cost function differences, which are the basis for making decisions in a gradient-free optimization, are exponentially suppressed in a barren plateau. Hence, without exponential precision, gradient-free optimizers will not make progress in the optimization. We numerically confirm this by training a parameterized quantum circuit in a barren plateau landscape with several gradient-free optimizers (Nelder-Mead, Powell, and COBYLA algorithms), and show that the number of shots required in the optimization grows exponentially with the number of qubits. These results provide new insight into the training landscapes of quantum neural networks and will inform the development of strategies to mitigate or avoid barren plateaus.

Andrew Arrasmith: The first two authors contributed equally to this work.

M. Cerezo: The first two authors contributed equally to this work.

## 1 Introduction

Parameterized quantum circuits offer a flexible paradigm for programming Noisy Intermediate Scale Quantum (NISQ) computers. These circuits are utilized in both Variational Quantum Algorithms (VQAs) [1–14] and Quantum Neural Networks (QNNs) [15–18]. Both VQA and QNN approaches involve efficiently evaluating a cost function  $C(\theta)$  or its gradient  $\nabla C(\theta)$  on a quantum computer. A classical optimizer is then employed to train the parameters  $\theta$  of a parameterized quantum circuit  $V(\theta)$  to minimize the cost.

Rigorous scaling results are urgently needed for this exciting approach to near-term quantum computing. Gradient scaling is one of the few directions of significant progress. The most famous gradient scaling result is the barren plateau phenomenon [19–33], whereby the gradient of the cost function shrinks exponentially with the number of qubits. Various issues lead to barren plateaus, such as deep ansatzes that lack structure [19, 21, 31], global cost functions [20, 21], high levels of noise [22, 32], scrambling target unitaries [24], and large entanglement [28, 29].

Without effort to avoid barren plateaus, this phenomenon can have a major impact on the scaling of one’s algorithm. Specifically, the exponential suppression of the gradient implies that one would need an exponential precision to make progress in the optimization, consequently, causing one’s algorithm to scale exponentially in the number of qubits. The standard goal of quantum algorithms is polynomial scaling, unlike the exponential scaling of classical algorithms. Hence, the exponential scaling due to barren plateaus could erase the possibility of a quantum speedup with a parametrized quantum circuit. It is therefore crucial to study barren plateaus in VQAs and QNNs in order to understand when quantum speedup is possible.

This has spawned an important research di-

rejection of finding strategies to avoid barren plateaus. Some examples include employing local cost functions [20], modifying the architecture [25, 26], pre-training [34], parameter correlation [35], layer-by-layer training [36], and initializing layers to the identity [37]. These strategies are promising. However, more analytical and numerical studies are needed to understand how effective they are in general, for example, as in Ref. [38].

One possible strategy to consider is the choice of optimizer. It is widely believed that gradient-based optimizers will be directly impacted by barren plateaus, for obvious reasons. Moreover, higher-order derivatives are also exponentially suppressed in a barren plateau [23], so optimizers based on such derivatives will also be impacted. Nevertheless, there still remains the question of whether gradient-free optimizers could somehow avoid the barren plateau problem. This is currently a topic of debate [20, 28]. The question is naturally made subtle by the fact that gradient-free optimizers can potentially use global information about the landscape, rather than being restricted to using local gradient information.

In this work, we present an analytical argument suggesting that gradient-free approaches will, indeed, be impacted by barren plateaus. Specifically, we show that cost function differences,  $C(\theta_B) - C(\theta_A)$ , will be exponentially suppressed in a barren plateau. This holds even when the points  $\theta_A$  and  $\theta_B$  are not necessarily close in parameter space. Gradient-free optimizers use such cost function differences to make decisions during the optimization. Hence, our results imply that such optimizers will either need to spend exponentially large resources to characterize cost function differences, or else these optimizers will not make progress in the optimization.

We confirm our analytical results with numerical simulations involving several gradient-free optimizers: Nelder-Mead, Powell, and COBYLA. For each of these optimizers, we attempt to train a deep parametrized quantum circuit, corresponding to the barren plateau scenario in Ref. [19]. In all cases, we find that the number of shots (i.e., the amount of statistics) required to begin to train the cost function grows exponentially in the number of qubits. This is the same behavior that one sees for gradient-based methods, and is a hallmark of the barren plateau

phenomenon.

## 2 Theoretical Background

Here we provide background needed to understand our results. We first consider the cost function used to train parameterized quantum circuits. Then we consider optimizers that can be used to optimize this cost function, with a specific focus on gradient-free optimizers. Finally, we give background on the barren plateau phenomenon.

### 2.1 Cost function

Consider a parameterized quantum circuit  $V(\theta)$ , whose parameters will be trained by minimizing a cost function  $C(\theta)$ . In this work, we consider a highly general cost function that can be expressed in the form

$$C(\theta) = \sum_{x=1}^S f_x(\theta, \rho_x). \quad (1)$$

Here,  $\theta$  is a vector of  $m$  continuous parameters,  $\{\rho_x\}_{x=1}^S$  are  $n$ -qubit input quantum states from a training set  $\mathcal{S}$  of size  $S$ , and  $f_x$  are functions that encode the problem and which can be different for each input state.

To ensure algorithmic efficiency, we assume that the number  $m$  of parameters in  $\theta$  is in  $\mathcal{O}(\text{poly}(n))$ . In addition we consider that any  $\theta_\mu \in \theta$  parametrizes a unitary of the form  $e^{-i\theta_\mu H_\mu}$ . We assume  $H_\mu$  is a Hermitian operator with two distinct non-zero eigenvalues (e.g.,  $H_\mu$  could be a Pauli operator).

We remark that the cost function in (1) contains as special cases many relevant applications. For instance, in a binary classification problem the cost function is given by the mean squared-error  $C(\theta) = \sum_x (y_x - \tilde{y}(\theta, \rho_x))^2$  [16, 17, 39, 40]. Here, the training set is given by  $\mathcal{S} = \{\rho_x, y_x\}$  where  $y_x$  are the true labels, and  $\tilde{y}(\theta, \rho_x)$  are the labels predicted by the Quantum Neural Network. In addition, the cost of several Variational Quantum Algorithms is covered by (1). In this case, the cost takes a simpler form, where the training set contains a single state ( $S = 1$ ) and the cost is  $C(\theta) = \text{Tr}[OV(\theta)\rho V^\dagger(\theta)]$ , with  $O$  a Hermitian operator [3–14].

The goal is then to solve the optimization problem

$$\theta_{\text{opt}} = \arg \min_{\theta} C(\theta). \quad (2)$$

This involves choosing an optimizer, which can either be a gradient-based or gradient-free optimizer. Various gradient-based approaches [41–44] have been proposed for training parameterized quantum circuits, and these will be directly impacted by barren plateaus. Optimizers employing higher-order derivatives are also impacted by barren plateaus [23]. In this work we consider the case when one employs a gradient-free optimization method. In the next section we review some widely-used gradient-free optimizers.

## 2.2 Gradient-Free Optimizers

We will refer to any optimization method that only accesses a zeroth-order oracle (i.e., does not directly access derivative information) as being gradient free. This is a very large class of methods, but they all depend on being able to distinguish cost function values at different points. Though our analytical results are general to any such optimizer, we now introduce three particular gradient-free optimizers that we will examine numerically: Nelder-Mead, Powell’s Method, and COBYLA.

### 2.2.1 Nelder-Mead

One popular gradient-free optimization strategy is the Nelder-Mead algorithm [45]. In this approach, one constructs a simplex in the space to be optimized over. Then one modifies it with a sequence of reflect, expand, contract, and shrink operations to move the simplex and then shrink it around the minimum. These operations are chosen based on conditional comparisons of the cost function values at each vertex as well as proposed new vertices. See Figure 1a for an illustration these operations.

When used in an environment where the errors in those cost function values are large enough to cause mistakes in these comparisons, however, this algorithm is vulnerable to performing shrink operations prematurely, which slows the optimization down and may lead to a false appearance of convergence [46]. Due to this difficulty, one would expect the number of iterations required to converge with Nelder-Mead to be especially bad in limited precision environments, though we note that there are a number of modifications that attempt to improve the method’s robustness to noise [46, 47].

### 2.2.2 Powell’s Method

The Powell algorithm [48] is another popular gradient-free optimizer that performs sequential line searches. This method starts with some input set of search vectors  $V = \{\mathbf{v}_i\}$ , usually just the coordinate directions in the parameter space. Searching along each of these directions in sequence, this method looks for the displacement  $\{a_i\}$  along each direction that would minimize the cost when only varying parameters along the current direction. Finding the displacements  $\{a_i\}$  is typically done with Brent’s parabolic interpolation method [49], though in principle one could use any univariate gradient free optimizer.

After sweeping through all of the search vectors, the iteration is completed by replacing the search vector  $\mathbf{v}_j$  that corresponds to the greatest displacement,  $a_j = \max(a_i)$ , and replacing it with

$$\mathbf{v}_j \rightarrow \sum_i a_i \mathbf{v}_i. \quad (3)$$

By making this replacement, convergence is accelerated and the method avoids getting stuck in a cyclic pattern of updates. See Figure 1b for a sketch of two iterations of this method.

### 2.2.3 COBYLA

Constrained Optimization BY Linear Approximation (COBYLA) is another popular gradient-free optimizer by Powell [50]. This algorithm constructs a simplex and uses the  $m + 1$  points in the parameter space, with  $m$  being the number of parameters, to define a hyperplane to capture the local slope of the cost function. The algorithm then replaces the highest cost function value point on the simplex by stepping from the lowest cost point along the direction of the slope. The method steps as far as possible along this estimated slope while staying within a lower bound on the radius of the trust region.

The lower bound on the size of the trust region is decreased when the algorithm detects that it has stopped making progress, allowing the method to converge. Note, however, that the size of the trust region never increases in COBYLA. An iteration of this method (showing a shrinking trust region) is sketched in Figure 1c.




Figure 1: Graphical depiction of the gradient-free optimizers considered. Panel **a** shows the different operations that the Nelder-Mead algorithm performs on the initial (grey) simplex: reflection (red), reflection with expansion (blue), reflection with contraction (green), and shrinking (turquoise). Panel **b** shows two iterations of the Powell method (with black for the first iteration and red for the second). Note that the direction of the final step is changed, reflecting a modified search direction. Finally, panel **c** shows an illustration of a COBYLA step from an initial (grey) simplex. After fitting a plane to the initial simplex, the method steps along the fitted slope to form a new simplex (red). The trust region is shown as a solid blue circle. A smaller trust region which might be used later in the optimization is illustrated with the dashed blue circle (though for this particular step the trust region would likely not be contracted).

### 2.3 Barren Plateaus

When the cost function exhibits a barren plateau, the cost function gradient vanishes exponentially with the system size. Without loss of generality we consider here the following generic definition of a barren plateau.

**Definition 1** (Barren Plateau). *Consider the cost function defined in Eq. (1). This cost exhibits a barren plateau if, for all  $\theta_\mu \in \boldsymbol{\theta}$ , the expectation value of the cost function partial derivative  $\partial C(\boldsymbol{\theta})/\partial\theta_\mu = \partial_\mu C(\boldsymbol{\theta})$  is  $E_{\boldsymbol{\theta}}[\partial_\mu C(\boldsymbol{\theta})] = 0$  and its variance vanishes exponentially with the number of qubits  $n$  as*

$$\text{Var}_{\boldsymbol{\theta}}[\partial_\mu C(\boldsymbol{\theta})] \leq F(n), \quad \text{with } F(n) \in \mathcal{O}\left(\frac{1}{b^n}\right). \quad (4)$$

for some  $b > 1$ . As indicated, the expectation values are taken over the parameters  $\boldsymbol{\theta}$ .

We remark here that, as shown in Definition 1, the barren plateau phenomenon is a probabilistic statement. In fact, from Chebyshev's inequality we know that  $\text{Var}_{\boldsymbol{\theta}}[\partial_\mu C(\boldsymbol{\theta})]$  bounds the probability that the cost function partial derivative deviates from its mean of zero as

$$P(|\partial_\mu C(\boldsymbol{\theta})| \geq c) \leq \frac{\text{Var}_{\boldsymbol{\theta}}[\partial_\mu C(\boldsymbol{\theta})]}{c^2}, \quad (5)$$

for any  $c > 0$ . In practice this means that by randomly initializing the parameters  $\boldsymbol{\theta}$ , there is a high probability that one ends up in a flat region of the landscape where the gradients are exponentially suppressed.

Let us now discuss different mechanisms that can lead to barren plateaus in the cost function landscape. As shown in the seminal work of Ref. [19], deep random unstructured circuits which form 2-designs on  $n$  qubits will exhibit barren plateaus. Here we use the term deep when the depth of the ansatz is in  $\mathcal{O}(\text{poly}(n))$ . For instance, as shown in [51–53] local circuits will form 2-designs when their depth is in  $\mathcal{O}(\text{poly}(n))$ .

The barren plateau phenomenon was extended in [20] to a type of shallow depth ansatz known as the layered hardware efficient ansatz, where random local gates act on alternating pairs of neighboring qubits in a brick-like structure. Here it was shown that the locality of the cost function can be linked to its trainability. Specifically, global cost functions (those where one compares operators living in exponentially large Hilbert spaces) exhibit barren plateaus for any circuit depth. On the other hand, it was shown that local cost functions (where one compares operators on an individual qubit level) are trainable when the ansatz depth is in  $\mathcal{O}(\log(n))$ , as here their gradients vanish at worst polynomially (rather than exponentially) with the system size.

Barren plateaus have also been shown to arise in more general QNN architectures [21, 28]. In perceptron-based QNNs with hidden and visible layers, connecting a large number of qubits in different layers with random global perceptrons (and hence highly entangling them) can lead to exponentially vanishing gradients. These results have shown that the barren plateau phenomenon is a

generic problem that can arise in multiple architectures for quantum machine learning.

Finally, in [22] a noise-induced barren plateau mechanism was found. Here it was proven that the presence of noise acting before and after each unitary layer in a parametrized quantum circuit leads to exponentially vanishing gradients for circuits with linear or super-linear depth. When the cost exhibits a noise-induced barren plateau we have  $|\partial_\mu C(\boldsymbol{\theta})| \leq \widehat{F}(n)$  with  $\widehat{F}(n) \in \mathcal{O}(1/\widehat{b}^n)$  for some  $\widehat{b} > 1$ . The underlying mechanism here is that the state gets corrupted due to noise, leading to a flattening of the whole cost landscape. This phenomenon is conceptually different from the previous barren plateaus as here one does not average over the parameters  $\boldsymbol{\theta}$ . Nevertheless, the noise-induced barren plateau still satisfies Definition 1, which is a weaker condition.

### 3 Main Results

In this section we first present our main analytical results in the form of Proposition 1 and Corollary 1. We then discuss the implications for employing gradient-free optimizers in a barren plateau.

#### 3.1 Exponentially suppressed cost differences

Here we consider two relevant scenarios where we analyze, on average, how large the difference  $\Delta C = C(\boldsymbol{\theta}_B) - C(\boldsymbol{\theta}_A)$  between two points in the landscape can be. First we consider the case when  $\boldsymbol{\theta}_A$  and  $\boldsymbol{\theta}_B$  are not independent, but rather  $\boldsymbol{\theta}_B$  can be obtained from  $\boldsymbol{\theta}_A$  through a given translation in parameter space. We then analyze the case when  $\boldsymbol{\theta}_A$  and  $\boldsymbol{\theta}_B$  are independent.

The following proposition constitutes the main result of our work. The proof is presented in the Appendix.

**Proposition 1.** *Consider the cost function of Eq. (1). Let  $\boldsymbol{\theta}_A$  be a randomly chosen point in parameter space. Let  $\boldsymbol{\theta}_B = \boldsymbol{\theta}_A + L\hat{\boldsymbol{\ell}}$  be a point at a distance  $L = \|\boldsymbol{\theta}_B - \boldsymbol{\theta}_A\|$  from  $\boldsymbol{\theta}_A$  in parameter space, for some unit vector  $\hat{\boldsymbol{\ell}}$ . If the cost exhibits a barren plateau according to Definition 1, then the expectation value of the difference  $\Delta C = C(\boldsymbol{\theta}_B) - C(\boldsymbol{\theta}_A)$  is*

$$E_{\boldsymbol{\theta}_A}[\Delta C] = 0, \quad (6)$$

and the variance is exponentially vanishing with  $n$  as

$$\text{Var}_{\boldsymbol{\theta}_A}[\Delta C] \leq G(n), \quad (7)$$

with

$$G(n) = m^2 L^2 F(n), \quad \text{and} \quad G(n) \in \widetilde{\mathcal{O}}\left(\frac{1}{b^n}\right), \quad (8)$$

for some  $b > 1$ . Here  $m$  is the dimension of the parameter space, and  $F(n)$  was defined in (4).

Let us here recall that we have assumed that  $m \in \mathcal{O}(\text{poly}(n))$ . Similarly, we have that  $\theta_\mu$  parametrizes a unitary generated by a Hermitian operator  $H_\mu$  with two distinct non-zero eigenvalues. From the latter it then follows that  $L$  is always in  $\mathcal{O}(\text{poly}(n))$ , and hence that  $G(n) \in \widetilde{\mathcal{O}}(1/b^n)$ .

From the previous results one can readily evaluate the case when  $\boldsymbol{\theta}_B$  and  $\boldsymbol{\theta}_A$  are independent. This case is of relevance to global optimizers, such as Bayesian approaches, where initial points on the landscape are chosen independently. This scenario can be analyzed by computing the expectation value  $E_{\boldsymbol{\theta}_A, \boldsymbol{\theta}_B}[\Delta C] = E_{\boldsymbol{\theta}_B}[E_{\boldsymbol{\theta}_A}[\Delta C]]$ . From Proposition 1, we can derive the following corollary.

**Corollary 1.** *Consider the cost function of Eq. (1). Let  $\boldsymbol{\theta}_A$  and  $\boldsymbol{\theta}_B$  be two randomly chosen points in parameter space. Without loss of generality we assume that  $\boldsymbol{\theta}_B = \boldsymbol{\theta}_A + L\hat{\boldsymbol{\ell}}$  for random  $L$  and  $\hat{\boldsymbol{\ell}}$  so that  $E_{\boldsymbol{\theta}_A, \boldsymbol{\theta}_B}[\dots] = E_{\boldsymbol{\theta}_A, L, \hat{\boldsymbol{\ell}}}[\dots]$ . If the cost exhibits a barren plateau according to Definition 1, then the expectation value of the difference  $\Delta C = C(\boldsymbol{\theta}_B) - C(\boldsymbol{\theta}_A)$  is  $E_{\boldsymbol{\theta}_A, L, \hat{\boldsymbol{\ell}}}[\Delta C] = 0$ , and the variance is exponentially vanishing with  $n$  as*

$$\text{Var}_{\boldsymbol{\theta}_A, L, \hat{\boldsymbol{\ell}}}[\Delta C] \leq \widehat{G}(n), \quad (9)$$

with

$$\widehat{G}(n) = m^2 \overline{L}^2 F(n), \quad \text{and} \quad \widehat{G}(n) \in \widetilde{\mathcal{O}}\left(\frac{1}{b^n}\right), \quad (10)$$

for some  $b > 1$ . Here  $m$  is the dimension of the parameter space,  $F(n)$  was defined in (4), and

$$\overline{L} = E_{L, \hat{\boldsymbol{\ell}}}[L] \quad (11)$$

is the average distance between any two points in parameter space.

The proof of Corollary 1 readily follows from Proposition 1 by additionally computing the expectation value over  $L$  and  $\hat{\boldsymbol{\ell}}$ . Moreover, here we



sider the toy-model problem where  $U = \mathbb{1}$  and the goal is to train the parameters in  $V(\boldsymbol{\theta})$  such that  $V(\boldsymbol{\theta})|\mathbf{0}\rangle = |\mathbf{0}\rangle$ , with  $|\mathbf{0}\rangle = |0\rangle^{\otimes n}$  the all-zero state. As shown in [7, 12], the following local cost function is faithful

$$C(\boldsymbol{\theta}) = \text{Tr}[O_L V(\boldsymbol{\theta})|\mathbf{0}\rangle\langle\mathbf{0}|V^\dagger(\boldsymbol{\theta})], \quad (14)$$

where

$$O_L = \mathbb{1} - \frac{1}{n} \sum_{i=1}^n |0_i\rangle\langle 0_i|, \quad (15)$$

in that can verify that  $C(\boldsymbol{\theta}) \in [0, 1]$ , with  $C(\boldsymbol{\theta}) = 0$  iff  $V(\boldsymbol{\theta})|\mathbf{0}\rangle = |\mathbf{0}\rangle$  (up to global phase). We remark that there is an efficient quantum circuit to compute  $C(\boldsymbol{\theta})$  [12].

For  $V(\boldsymbol{\theta})$  we employ a layered hardware efficient ansatz as shown in Fig. 2. Moreover, we recall that Ref. [19] showed that a cost function such as (14) with a layered hardware efficient ansatz that is randomly initialized will exhibit barren plateaus when the depth of  $V(\boldsymbol{\theta})$  scales at least linearly in  $n$ .

In our numerics we simulated the aforementioned quantum compilation task for different numbers of qubits  $n = 5, 6, \dots, 11$ . Letting  $p$  be the number of layers in the ansatz in Fig. 2, we choose  $p = n$ , so that the depth grows linearly in  $n$ . This corresponds to the barren plateau scenario in Ref. [19]. For each value of  $n$ , we solved the optimization of Eq. (2) by employing the Nelder-Mead, Powell, and COBYLA methods. These simulations were performed using MATLAB (Nelder-Mead) and SciPy (Powell and COBYLA). In all cases we randomly initialized the parameters  $\boldsymbol{\theta}$  in the ansatz and we ran the optimizer until a cost function value  $C = 0.4$  was achieved or until a maximal total number of shots used throughout the optimization was surpassed. For simplicity we use the default values for hyperparameters not related to optimization termination. We note that we chose a relatively large value ( $C = 0.4$ ) for the cost threshold because we are interested in the initial stages of the training process, i.e., the question of whether one can get past the barren plateau. With this choice, the computational expense of reaching this threshold does not take into account the difficulty of finding a minimum, it only reflects the burden of the barren plateau.

Since the goal is to heuristically determine the precision (i.e., the number of shots  $N$ ) needed to minimize the cost, we first ran simulations with

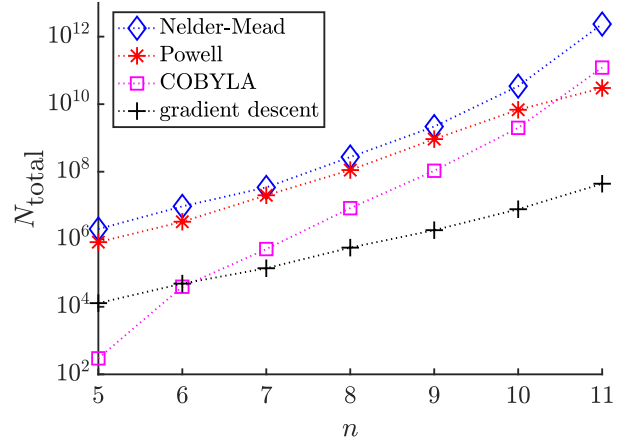


Figure 3: Median value of the total shot number  $N_{\text{total}}$  necessary to reach  $C = 0.4$  plotted versus the number of qubits  $n$ . We show results for Nelder-Mead (blue diamonds), Powell (red asterisks), and COBYLA (pink squares) implementations, as well as results for a gradient-descent implementation (black crosses) for reference. Each data point was obtained from a sample of 20 runs initialized by random initial angles  $\boldsymbol{\theta}$ .

different values of  $N$  allocated per cost-function evaluation. For each  $N$  we simulated 20 optimization instances (runs) with different initial points and we kept track of  $N_{\text{total}}$  used throughout the optimization. The next step was to determine the value of  $N$  for which a cost of  $C = 0.4$  could be reached and which minimizes the median total number of shots  $N_{\text{total}}$  computed over different runs. We analyze the scaling of this median value as a function of  $n$  below.

In Fig. 3 we present our numerical results. Here we can see that the total number of shots scales exponentially with  $n$  for the Powell method, and scales super-exponentially for the Nelder-Mead optimizer. For the COBYLA method, the  $N_{\text{total}}$  behavior as a function of  $n$  is not very regular but it is consistent with at least an exponential increase. As a reference point we also show in Fig. 3 results obtained with a custom gradient-descent optimizer. As expected, the total number of shots also scales exponentially in this case.

## 5 Discussion

With a wide range of applications spanning chemistry, optimization, and big data analysis, training parameterized quantum circuits is arguably the leading paradigm for near-term quantum computing. Yet barren plateaus in the training

landscape remains an obstacle to making these paradigms scalable. Hence, one of the most important lines of research in this field is developing methods to avoid barren plateaus.

In this work, we consider the question of whether the choice of optimizer could be a potential strategy in avoiding barren plateaus. We focus on gradient-free optimizers, since there has been recent debate in the community about whether barren plateaus affect such optimizers.

Our main result is an analytical argument suggesting that gradient-free optimizers will, indeed, be impacted by barren plateaus. Proposition 1 is relevant to gradient-free optimizers that search through the landscape starting from a (random) initial point. For example, this includes simplex-based optimizers like Nelder-Mead. This proposition asserts that the variance of cost function differences is exponentially suppressed in a barren plateau. This implies that such optimizers will need to expend exponentially large resources in order to make decisions about where to move in the landscape.

Corollary 1 considers a slightly different scenario, where both points are randomly and independently chosen. This is relevant to global gradient-free optimizers, such as Bayesian methods, which initially choose multiple random points on the landscape, and then proceed from this initial set of points. This corollary implies that these optimizers will also need to utilize exponentially large resources in order to make progress.

We also numerically attempt to train in a barren plateau scenario using several gradient-free optimizers. We ask how many shots are required to begin to train the cost function. In all cases, we find that the required number of shots grows exponentially in the number of qubits. This is consistent with our main result, and demonstrates that barren plateaus can lead to exponential scaling even for gradient-free optimizers. We note that this exponential scaling is a lower bound on the asymptotics. For the Nelder-Mead we find super-exponential scaling, likely due to the chances of prematurely shrinking the simplex when it is hard to order the cost values [46]. For the case of COBYLA, there may be a similar effect from prematurely shrinking the radius of the trust region, though this is not clearly demonstrated in our data. Finally, the Powell method

appears to show exponential scaling. It is likely that the reason Powell shows better scaling is that, unlike the other optimizers, statistical noise does not have a cumulative effect on the state of the optimizer.

Our work casts doubt on the notion that the choice of optimizer could provide a strategy to avoid barren plateaus. While the asymptotically exponential scaling cannot be avoided, we note that the size limits of trainable problems may be extended by a careful choice of optimization strategy. For example, techniques using neural networks [34] or natural evolutionary strategies [55] may improve the constants multiplying the exponential scaling. However, we emphasize that all such strategies at minimum require the comparison of cost function values at different points and thus are subject to our scaling analysis.

This result highlights the difficult challenge posed by barren plateaus. Future work certainly should continue to develop strategies to avoid them. Additionally, in future work, we hope to develop a unified treatment that covers the impact of barren plateaus on various types of optimizers, gradient-based and gradient-free.

## Acknowledgements

AA and LC were initially supported by LDRD program of LANL under project number 20190065DR. MC acknowledges support from the Center for Nonlinear Studies at Los Alamos National Laboratory (LANL). Piotr C. was supported by the Laboratory Directed Research and Development (LDRD) program of Los Alamos National Laboratory (LANL) under project number 20190659PRD4. PJC acknowledges initial support from the LANL ASC Beyond Moore's Law project. This work was supported by the U.S. DOE, Office of Science, Office of Advanced Scientific Computing Research, under the Accelerated Research in Quantum Computing (ARQC) program.

## References

- [1] 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*, *Nature Reviews Physics* **1**, 19 (2021).



- [2] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwok, and A. Aspuru-Guzik, *Noisy intermediate-scale quantum (nisq) algorithms*, [arXiv preprint arXiv:2101.08448](#) (2021).
- [3] 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*, *Nature communications* **5**, 1 (2014).
- [4] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, *The theory of variational hybrid quantum-classical algorithms*, *New Journal of Physics* **18**, 023023 (2016).
- [5] E. Farhi, J. Goldstone, and S. Gutmann, *A quantum approximate optimization algorithm*, [arXiv preprint arXiv:1411.4028](#) (2014).
- [6] J. Romero, J. P. Olson, and A. Aspuru-Guzik, *Quantum autoencoders for efficient compression of quantum data*, *Quantum Science and Technology* **2**, 045001 (2017).
- [7] S. Khatri, R. LaRose, A. Poremba, L. Cincio, A. T. Sornborger, and P. J. Coles, *Quantum-assisted quantum compiling*, *Quantum* **3**, 140 (2019).
- [8] R. LaRose, A. Tikku, É. O’Neel-Judy, L. Cincio, and P. J. Coles, *Variational quantum state diagonalization*, *npj Quantum Information* **5**, 1 (2019).
- [9] A. Arrasmith, L. Cincio, A. T. Sornborger, W. H. Zurek, and P. J. Coles, *Variational consistent histories as a hybrid algorithm for quantum foundations*, *Nature communications* **10**, 1 (2019).
- [10] M. Cerezo, A. Poremba, L. Cincio, and P. J. Coles, *Variational quantum fidelity estimation*, *Quantum* **4**, 248 (2020).
- [11] C. Cirstoiu, Z. Holmes, J. Iosue, L. Cincio, P. J. Coles, and A. Sornborger, *Variational fast forwarding for quantum simulation beyond the coherence time*, *npj Quantum Information* **6**, 1 (2020).
- [12] K. Sharma, S. Khatri, M. Cerezo, and P. J. Coles, *Noise resilience of variational quantum compiling*, *New Journal of Physics* **22**, 043006 (2020).
- [13] C. Bravo-Prieto, R. LaRose, M. Cerezo, Y. Subasi, L. Cincio, and P. Coles, *Variational quantum linear solver*, [arXiv preprint arXiv:1909.05820](#) (2019).
- [14] M. Cerezo, K. Sharma, A. Arrasmith, and P. J. Coles, *Variational quantum state eigensolver*, [arXiv preprint arXiv:2004.01372](#) (2020).
- [15] M. Schuld, I. Sinayskiy, and F. Petruccione, *The quest for a quantum neural network*, *Quantum Information Processing* **13**, 2567 (2014).
- [16] I. Cong, S. Choi, and M. D. Lukin, *Quantum convolutional neural networks*, *Nature Physics* **15**, 1273 (2019).
- [17] K. Beer, D. Bondarenko, T. Farrelly, T. J. Osborne, R. Salzmann, D. Scheiermann, and R. Wolf, *Training deep quantum neural networks*, *Nature Communications* **11**, 808 (2020).
- [18] G. Verdon, J. Pye, and M. Broughton, *A universal training algorithm for quantum deep learning*, [arXiv preprint arXiv:1806.09729](#) (2018).
- [19] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, *Barren plateaus in quantum neural network training landscapes*, *Nature communications* **9**, 1 (2018).
- [20] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, *Cost function dependent barren plateaus in shallow parametrized quantum circuits*, *Nature Communications* **12**, 1791 (2021).
- [21] K. Sharma, M. Cerezo, L. Cincio, and P. J. Coles, *Trainability of dissipative perceptron-based quantum neural networks*, [arXiv preprint arXiv:2005.12458](#) (2020).
- [22] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, *Noise-induced barren plateaus in variational quantum algorithms*, [arXiv preprint arXiv:2007.14384](#) (2020).
- [23] M. Cerezo and P. J. Coles, *Higher order derivatives of quantum neural networks with barren plateaus*, *Quantum Science and Technology* **6**, 035006 (2021).
- [24] Z. Holmes, A. Arrasmith, B. Yan, P. J. Coles, A. Albrecht, and A. T. Sornborger, *Barren plateaus preclude learning scramblers*, *Physical Review Letters* **126**, 190501 (2021).

- [25] A. Pesah, M. Cerezo, S. Wang, T. Volkoff, A. T. Sornborger, and P. J. Coles, *Absence of barren plateaus in quantum convolutional neural networks*, [arXiv preprint arXiv:2011.02966](#) (2020).
- [26] K. Zhang, M.-H. Hsieh, L. Liu, and D. Tao, *Toward trainability of quantum neural networks*, [arXiv preprint arXiv:2011.06258](#) (2020).
- [27] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, *The power of quantum neural networks*, *Nature Computational Science* **1**, 403 (2021).
- [28] C. O. Marrero, M. Kieferová, and N. Wiebe, *Entanglement induced barren plateaus*, [arXiv preprint arXiv:2010.15968](#) (2020).
- [29] T. L. Patti, K. Najafi, X. Gao, and S. F. Yelin, *Entanglement devised barren plateau mitigation*, *Physical Review Research* **3**, 033090 (2021).
- [30] A. Uvarov and J. D. Biamonte, *On barren plateaus and cost function locality in variational quantum algorithms*, *Journal of Physics A: Mathematical and Theoretical* **54**, 245301 (2021).
- [31] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, *Connecting ansatz expressibility to gradient magnitudes and barren plateaus*, [arXiv preprint arXiv:2101.02138](#) (2021).
- [32] Y. Du, M.-H. Hsieh, T. Liu, S. You, and D. Tao, *On the learnability of quantum neural networks*, [arXiv preprint arXiv:2007.12369](#) (2020).
- [33] A. Arrasmith, Z. Holmes, M. Cerezo, and P. J. Coles, *Equivalence of quantum barren plateaus to cost concentration and narrow gorges*, [arXiv preprint arXiv:2104.05868](#) (2021).
- [34] 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*, [arXiv preprint arXiv:1907.05415](#) (2019).
- [35] T. Volkoff and P. J. Coles, *Large gradients via correlation in random parameterized quantum circuits*, *Quantum Science and Technology* **6**, 025008 (2021).
- [36] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, and M. Leib, *Layer-wise learning for quantum neural networks*, *Quantum Machine Intelligence* **3**, 1 (2021).
- [37] E. Grant, L. Wossnig, M. Ostaszewski, and M. Benedetti, *An initialization strategy for addressing barren plateaus in parametrized quantum circuits*, *Quantum* **3**, 214 (2019).
- [38] E. Campos, A. Nasrallah, and J. Biamonte, *Abrupt transitions in variational quantum circuit training*, *Physical Review A* **103**, 032607 (2021).
- [39] E. Farhi and H. Neven, *Classification with quantum neural networks on near term processors*, [arXiv preprint arXiv:1802.06002](#) (2018).
- [40] N. Killoran, T. R. Bromley, J. M. Arrazola, M. Schuld, N. Quesada, and S. Lloyd, *Continuous-variable quantum neural networks*, *Physical Review Research* **1**, 033063 (2019).
- [41] J. M. Kübler, A. Arrasmith, L. Cincio, and P. J. Coles, *An adaptive optimizer for measurement-frugal variational algorithms*, *Quantum* **4**, 263 (2020).
- [42] R. Sweke, F. Wilde, J. J. Meyer, M. Schuld, P. K. Fährmann, B. Meynard-Piganeau, and J. Eisert, *Stochastic gradient descent for hybrid quantum-classical optimization*, *Quantum* **4**, 314 (2020).
- [43] A. Arrasmith, L. Cincio, R. D. Somma, and P. J. Coles, *Operator sampling for shot-frugal optimization in variational algorithms*, [arXiv preprint arXiv:2004.06252](#) (2020).
- [44] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, *Quantum natural gradient*, *Quantum* **4**, 269 (2020).
- [45] J. A. Nelder and R. Mead, *A simplex method for function minimization*, *The computer journal* **7**, 308 (1965).
- [46] R. R. Barton and J. S. Ivey Jr, *Modifications of the nelder-mead simplex method for stochastic simulation response optimization*, 1991 Winter Simulation Conference Proceedings (1991), 10.1109/WSC.1991.185709.
- [47] X. Huang, *Robust simplex algorithm for online optimization*, *Physical Review Accelerators and Beams* **21**, 104601 (2018).
- [48] M. J. Powell, *An efficient method for finding the minimum of a function of several vari-*

- ables without calculating derivatives, *The computer journal* **7**, 155 (1964).
- [49] R. P. Brent, *Algorithms for minimization without derivatives* (Courier Corporation, 2013).
- [50] M. J. Powell, in *Advances in optimization and numerical analysis* (Springer, 1994) pp. 51–67.
- [51] A. W. Harrow and R. A. Low, *Random quantum circuits are approximate 2-designs*, *Communications in Mathematical Physics* **291**, 257 (2009).
- [52] F. G. Brandao, A. W. Harrow, and M. Horodecki, *Local random quantum circuits are approximate polynomial-designs*, *Communications in Mathematical Physics* **346**, 397 (2016).
- [53] A. Harrow and S. Mehraban, *Approximate unitary  $t$ -designs by short random quantum circuits using nearest-neighbor and long-range gates*, arXiv preprint arXiv:1809.06957 (2018).
- [54] J. Moćkus, in *Optimization techniques IFIP technical conference* (Springer, 1975) pp. 400–404.
- [55] A. Anand, M. Degroote, and A. Aspuru-Guzik, *Natural evolutionary strategies for variational quantum computation*, *Machine Learning: Science and Technology* (2021), 10.1088/2632-2153/abf3ac.

## Appendix

### A Proof of Proposition 1

We now prove our main result. We restate the proposition here for convenience.

**Proposition 1.** *Consider the cost function of Eq. (1). Let  $\theta_A$  be a randomly chosen point in parameter space. Let  $\theta_B = \theta_A + L\hat{\ell}$  be a point at a distance  $L = \|\theta_B - \theta_A\|$  from  $\theta_A$  in parameter space, for some unit vector  $\hat{\ell}$ . If the cost exhibits a barren plateau according to Definition 1, then the expectation value of the difference  $\Delta C = C(\theta_B) - C(\theta_A)$  is*

$$E_{\theta_A}[\Delta C] = 0, \quad (16)$$

and the variance is exponentially vanishing with  $n$  as

$$\text{Var}_{\theta_A}[\Delta C] \leq G(n), \quad (17)$$

with

$$G(n) = m^2 L^2 F(n), \quad \text{and} \quad G(n) \in \tilde{\mathcal{O}}\left(\frac{1}{b^n}\right), \quad (18)$$

for some  $b > 1$ . Here  $m$  is the dimension of the parameter space, and  $F(n)$  was defined in (4).

*Proof.* We begin by noting that finite differences in cost values can be expressed as:

$$\begin{aligned} \Delta C &= C(\theta_B) - C(\theta_A) \\ &= \int_{\theta_A}^{\theta_B} \nabla C(\theta) \cdot d\theta. \end{aligned} \quad (19)$$

Note that as we are integrating over gradients (which form a conservative vector field by definition) this integral is path independent. We can therefore choose to integrate along the line segment between  $\theta_B$  and  $\theta_A$ , setting  $\theta = \theta_A + \ell\hat{\ell}$ , with  $\hat{\ell}$  being the unit vector along  $\theta_B - \theta_A$  and  $\ell \in [0, L]$  where  $L = \|\theta_B - \theta_A\|$ . We then have  $d\theta = \hat{\ell}d\ell$ . To show that the mean difference is zero, we simply take the expectation value of this integral with respect to  $\theta_A$

$$\begin{aligned} E_{\theta_A}[\Delta C] &= \int_0^L E_{\theta_A}[\nabla C(\theta_A + \ell\hat{\ell}) \cdot \hat{\ell}]d\ell \\ &= 0. \end{aligned} \quad (20)$$

This follows as averaging  $\boldsymbol{\theta}_A$  over parameter space is equivalent to averaging  $\boldsymbol{\theta}_A + \ell\hat{\boldsymbol{\ell}}$  over the same space (for any fixed  $\ell\hat{\boldsymbol{\ell}}$ ), and the average gradient over the parameter space is zero from the definition of a barren plateau (see Definition 1).

Similarly, we can compute the magnitude of the second moment of this difference as

$$\begin{aligned} |\mathbb{E}_{\boldsymbol{\theta}_A}[(\Delta C)^2]| &= \left| \int_0^L \int_0^L \mathbb{E}_{\boldsymbol{\theta}_A}[(\nabla C(\boldsymbol{\theta}_A + \ell\hat{\boldsymbol{\ell}}) \cdot \hat{\boldsymbol{\ell}})(\nabla C(\boldsymbol{\theta}_A + \ell'\hat{\boldsymbol{\ell}}) \cdot \hat{\boldsymbol{\ell}})] d\ell d\ell' \right| \\ &= \left| \int_0^L \int_0^L \text{Cov}_{\boldsymbol{\theta}_A} \left( (\nabla C(\boldsymbol{\theta}_A + \ell\hat{\boldsymbol{\ell}}) \cdot \hat{\boldsymbol{\ell}}), (\nabla C(\boldsymbol{\theta}_A + \ell'\hat{\boldsymbol{\ell}}) \cdot \hat{\boldsymbol{\ell}}) \right) d\ell d\ell' \right|. \end{aligned} \quad (21)$$

Note that we have dropped the square of the mean as the mean is zero, so the second moment is the covariance. Next, we upper bound this covariance using the product of the variances using the Cauchy-Schwartz inequality

$$|\mathbb{E}_{\boldsymbol{\theta}_A}[(\Delta C)^2]| \leq \int_0^L \int_0^L \left( \text{Var}_{\boldsymbol{\theta}_A}(\nabla C(\boldsymbol{\theta}_A + \ell\hat{\boldsymbol{\ell}}) \cdot \hat{\boldsymbol{\ell}}) \text{Var}_{\boldsymbol{\theta}_A}(\nabla C(\boldsymbol{\theta}_A + \ell'\hat{\boldsymbol{\ell}}) \cdot \hat{\boldsymbol{\ell}}) \right)^{\frac{1}{2}} d\ell d\ell'. \quad (22)$$

Next, we split these variances into the covariances between individual components of the gradient

$$\begin{aligned} |\mathbb{E}_{\boldsymbol{\theta}_A}[(\Delta C)^2]| &\leq \int_0^L \int_0^L \left( \sum_{i=1}^m \sum_{j=1}^m \sum_{p=1}^m \sum_{q=1}^m \text{Cov}_{\boldsymbol{\theta}_A} \left( \partial_i C(\boldsymbol{\theta}_A + \ell\hat{\boldsymbol{\ell}}) \hat{\ell}_i, \partial_j C(\boldsymbol{\theta}_A + \ell'\hat{\boldsymbol{\ell}}) \hat{\ell}_j \right) \right. \\ &\quad \left. \text{Cov}_{\boldsymbol{\theta}_A}(\partial_p C(\boldsymbol{\theta}_A + \ell\hat{\boldsymbol{\ell}}) \hat{\ell}_p, \partial_q C(\boldsymbol{\theta}_A + \ell'\hat{\boldsymbol{\ell}}) \hat{\ell}_q) \right)^{\frac{1}{2}} d\ell d\ell'. \end{aligned} \quad (23)$$

As before, these covariances can be bounded by variances using the Cauchy-Schwartz inequality

$$\begin{aligned} |\mathbb{E}_{\boldsymbol{\theta}_A}[(\Delta C)^2]| &\leq \int_0^L \int_0^L \left( \sum_{i=1}^m \sum_{j=1}^m \sum_{p=1}^m \sum_{q=1}^m \left( \text{Var}_{\boldsymbol{\theta}_A} \left( \partial_i C(\boldsymbol{\theta}_A + \ell\hat{\boldsymbol{\ell}}) \right) \text{Var}_{\boldsymbol{\theta}_A} \left( \partial_j C(\boldsymbol{\theta}_A + \ell'\hat{\boldsymbol{\ell}}) \right) \right)^{\frac{1}{2}} \right. \\ &\quad \left. \left( \text{Var}_{\boldsymbol{\theta}_A} \left( \partial_p C(\boldsymbol{\theta}_A + \ell\hat{\boldsymbol{\ell}}) \right) \text{Var}_{\boldsymbol{\theta}_A} \left( \partial_q C(\boldsymbol{\theta}_A + \ell'\hat{\boldsymbol{\ell}}) \right) \right)^{\frac{1}{2}} \right)^{\frac{1}{2}} d\ell d\ell'. \end{aligned} \quad (24)$$

If each variance can be bounded by a function  $F(n)$ , we can then write

$$\begin{aligned} |\mathbb{E}_{\boldsymbol{\theta}_A}[(\Delta C)^2]| &\leq m^2 \int_0^L \int_0^L F(n) d\ell d\ell' \\ &= m^2 L^2 F(n) \\ &= G(n). \end{aligned} \quad (25)$$

If the cost exhibits a barren plateau, then by Definition 1,  $F(n) \in \mathcal{O}\left(\frac{1}{b^n}\right)$  for some  $b > 1$ . It then follows that if  $m \in \mathcal{O}(\text{poly}(n))$  and  $L \in \mathcal{O}(\text{poly}(n))$ , then  $G(n) \in \mathcal{O}\left(\frac{1}{b^n}\right)$  and Proposition 1 holds.