



This is the accepted manuscript made available via CHORUS. The article has been published as:

## Barren Plateaus Preclude Learning Scramblers

Zoë Holmes, Andrew Arrasmith, Bin Yan, Patrick J. Coles, Andreas Albrecht, and Andrew T. Sornborger

Phys. Rev. Lett. **126**, 190501 — Published 12 May 2021

DOI: [10.1103/PhysRevLett.126.190501](https://doi.org/10.1103/PhysRevLett.126.190501)

# Barren plateaus preclude learning scramblers

Zoë Holmes,<sup>1,\*</sup> Andrew Arrasmith,<sup>2,\*</sup> Bin Yan,<sup>3,2,\*</sup>

Patrick J. Coles,<sup>2</sup> Andreas Albrecht,<sup>4</sup> and Andrew T. Sornborger<sup>1</sup>

<sup>1</sup>*Information Sciences, Los Alamos National Laboratory, Los Alamos, NM USA.*

<sup>2</sup>*Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM USA.*

<sup>3</sup>*Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, NM USA.*

<sup>4</sup>*Center for Quantum Mathematics and Physics and Department of Physics and Astronomy  
UC Davis, One Shields Ave, Davis CA.*

(Dated: February 23, 2021)

Scrambling processes, which rapidly spread entanglement through many-body quantum systems, are difficult to investigate using standard techniques, but are relevant to quantum chaos and thermalization. In this Letter, we ask if quantum machine learning (QML) could be used to investigate such processes. We prove a no-go theorem for learning an unknown scrambling process with QML, showing that *any variational ansatz* is highly probable to have a barren plateau landscape, i.e., cost gradients that vanish exponentially in the system size. This implies that the required resources scale exponentially even when strategies to avoid such scaling (e.g., from ansatz-based barren plateaus or No-Free-Lunch theorems) are employed. Furthermore, we numerically and analytically extend our results to approximate scramblers. Hence, our work places generic limits on the learnability of unitaries when lacking prior information.

**Introduction.** The growth of entanglement in many body quantum systems can rapidly distribute the information contained in the initial conditions of the system throughout a large number of degrees of freedom. This process is known as *scrambling* [1–3]. In recent years scrambling has proven central not only to understanding quantum chaos but also to the study of the dynamics of quantum information [4–6], thermalization phenomena [7, 8], the black hole information paradox [1, 9–11], holography [12, 13], random circuits [14–16], fluctuation relations [17, 18] and entropic uncertainty relations [19]. However, the complexity of strongly-interacting many-body quantum systems makes scrambling rather challenging to study analytically. Furthermore, experimental studies of scramblers are demanding given the difficulties of distinguishing scrambling from decoherence and other experimental imperfections [10, 20].

Quantum computers have recently been used as a testbed for the study of scramblers [20–23]. A possible further use of quantum computing would be to investigate scrambling using quantum machine learning (QML) methods [24]. Here we define QML [25–28] as any method that optimizes a parameterized quantum circuit by minimizing a problem-specific cost function. This includes variational quantum algorithms (VQAs) [29–45], which are used for numerous applications. It has recently been shown that entanglement provides a resource to exponentially reduce the number of training states required to learn a quantum process [28]. Thus, one might hope that QML could prove an effective tool to study quantum scrambling. For example, Figure 1 shows a schematic of how this could be done for the Hayden-Preskill thought experiment [1].




FIG. 1. **Learning a scrambling unitary.** Panel (a) shows the setup of the classic Hayden-Preskill thought experiment where someone attempts to retrieve information (shown as a book) thrown into a black hole (a scrambler). If the scrambling unitary  $U$  is known, then information can be retrieved. Panel (b) shows the process of attempting to learn  $U$ . This requires a time that is exponential in the number of quantum degrees of freedom (qubits) due to an exponentially vanishing cost gradient, see Letter below. This precludes the information retrieval shown in (a).

However, despite the high expectations placed on QML, there remain fundamental questions concerning its scalability and breadth of applicability. Of particular concern is the growing body of literature on the existence of *barren plateaus*, i.e., regions in parameter space where cost gradients vanish exponentially as the size of the system studied increases. This phenomenon, which exponentially increases the resources required to train large scale quantum neural networks, has been demonstrated in a number of proposed architectures and classes of cost function [46–51].

In this paper we present a no-go theorem for the use of QML to study quantum scrambling. Namely, we show that any QML approach used to learn the unitary dynamics implemented by a typical scrambler will exhibit a barren plateau and thus be untrainable in the absence

\* The first three authors contributed equally to this work.

of further prior knowledge.

In contrast to the barren plateau phenomenon established in Ref. [46], which is a consequence of the ansatz structure and parameter initialization strategy, our barren plateau result holds *for any choice of ansatz and any initialization of parameters*. Thus, previously proposed strategies to avoid barren plateaus [52–54], such as correlating parameters, do not address the issue raised in our work. Our result is conceptually distinct from previous barren plateau results, and additionally provides an alternative perspective on trainability issues in QML.

Given the close connection between chaos and randomness, our no-go theorem also applies to learning random and pseudo-random unitaries. As such, our result implies that to efficiently learn an unknown unitary process using QML, prior information about that process is required. For this reason, our work constrains the use of QML in the study of complex, arbitrary physical processes.

*Preliminaries.* To illustrate the machine learning task we consider, let us start by recalling the famous Hayden-Preskill thought experiment [1] (Fig. 1). Suppose Alice attempts to destroy a secret, encoded in a quantum state, by throwing it into Nature’s fastest scrambler, a black hole. How safe is Alice’s secret? Hayden and Preskill argued that if Bob knows the unitary dynamics,  $U$ , implemented by the black hole, and shares a maximally entangled state with the black hole, it is possible to decode Alice’s secret by collecting a few additional photons emitted from the black hole. However, this prompts a second question, how might Bob learn the scrambling unitary in the first place? Here we investigate whether QML can be used to learn the scrambling unitary,  $U$ .

To address this, we first motivate our notion of a scrambler. A diagnostic for information scrambling that has attracted considerable recent attention is the out-of-time-ordered correlator (OTOC) [3, 13, 55], a four-point correlator with unusual time ordering,

$$f_{\text{OTOC}} \equiv \langle \tilde{X} Y \tilde{X}^\dagger Y^\dagger \rangle. \quad (1)$$

Here  $X$  and  $Y$  are local operators on different subsystems,  $\tilde{X} = U^\dagger X(0)U$  is the Heisenberg evolved initial operator  $X(0)$  and the average is taken over an infinite temperature state  $\rho \propto \mathbb{1}$ .

For scrambling dynamics, this quantity decays rapidly and persists at a small value. This behavior can be made more transparent by noting that if  $X$  and  $Y$  are both Hermitian and unitary then the OTOC can be written as

$$f_{\text{OTOC}} = 1 - \frac{1}{2} \langle [\tilde{X}, Y][\tilde{X}, Y]^\dagger \rangle. \quad (2)$$

Since  $X$  and  $Y$  act on different subsystems, their un-evolved commutator vanishes. However, scrambling dynamics evolve these local operators into global ones, inducing a growth of the commutator.

In this work, we are interested in learning the evolution unitary at late times, when the dynamics become sufficiently complex for universal structures to form. For instance, the unitary  $U = \exp(-iHt)$  of a chaotic Hamiltonian  $H$  appears increasingly random with time. Specifically, after a time scale called the scrambling time, the OTOC of a chaotic system tends to a minimal value that is equivalent to taking its average over a random distribution of unitaries [14].

The link between scrambling and randomness can be made more precise by introducing the concept of unitary designs. An ensemble of unitaries with distribution  $\mu$  is a unitary  $k$ -design if its statistics agree with those of the Haar random distribution up to the  $k$ -th moment, i.e., if for any  $X$ ,

$$\int_{\mu} dU U^{\otimes k}(X) U^{\dagger \otimes k} = \int_{\text{Haar}} dU U^{\otimes k}(X) U^{\dagger \otimes k}. \quad (3)$$

Since  $f_{\text{OTOC}}$  in (1) only involves the second moment of the unitary, its asymptotic smallness can be attributed to the fact that the scrambling unitary appears to be a typical element of a 2-design [14]. Hence, a scrambler can be modelled as a unitary that is drawn from a distribution that forms at least a unitary 2-design [1, 2]. We remark that the dynamics of chaotic systems before reaching the scrambling time provide a model for approximate scrambling behavior [56], which we study in our numerics below.

*Main results.* Having formalized our notion of scrambling, we are now in a position to present our main result on the learnability of scramblers. The aim of QML is to minimize a problem-specific cost function that is evaluated on a quantum computer. In the context of learning unitaries, one considers an *ansatz* (i.e., parameterized quantum circuit)  $U(\theta)$  and a target unitary  $V$ . It is then natural to define the product  $W(\theta) = V^\dagger U(\theta)$ , which would be proportional to the identity in the case of perfect training, i.e., when  $U(\theta)$  matches the target  $V$ . To quantify the quality of the training, one can employ a generic cost function of the form

$$C(\theta, V) = \langle \psi | W(\theta)^\dagger H W(\theta) | \psi \rangle, \quad (4)$$

where  $|\psi\rangle$  is some state and  $H$  is some Hermitian operator. While in the main text we focus on the cost in (4), in Appendix C we extend our results to a more general cost of the form

$$C_{\text{gen}}(\theta, V) = \sum_i p_i \text{Tr}[H_i (\mathcal{E}_\theta \mathcal{E}_V \otimes I_R) (|\psi_i\rangle\langle\psi_i|)], \quad (5)$$

where  $\mathcal{E}_V(X) = V X V^\dagger$  and the weighted sum with  $\sum_i p_i = 1$  accounts for multiple training states  $|\psi_i\rangle$  and measurements  $H_i$ . Here the ansatz is modeled by a general channel  $\mathcal{E}_\theta$ , and thus can incorporate dissipative effects such as those explored in Refs. [26, 48]. We remark that standard cost functions for variational compiling of

unitaries [33, 37] fall under the framework of our general cost in (5), as discussed in Appendix E.

Crucially, the training data  $\{|\psi_i\rangle, H_i\}$  in Eq. (5) now acts not just on the scrambling system  $S$  but can also be entangled with a reference system  $R$ . This is important as the No-Free-Lunch theorem introduced in Ref. [27] entails that in the absence of entanglement an exponential number of training states are required to learn an unknown unitary. Conversely, the subsequent entanglement-enhanced No-Free-Lunch theorem showed this exponential scaling can be avoided using entangled training data [28]. In fact, a single, appropriately chosen, entangled training state may be sufficient to learn any unitary operation. This suggests it could be possible to use QML to study scramblers.

However, both the original and the entanglement-enhanced No-Free-Lunch theorems only concern the size of the training set needed to learn the action of a unitary over all states, assuming perfect training. They say nothing about the *trainability* of a VQA, that is whether the cost landscape is sufficiently featured to enable the action of the unitary to be learnt on any training data. Here we address the latter.

Barren plateaus have been proven [46] for a wide class of variational quantum algorithms, including those that aim to learn unitaries, whenever the training ansatz  $U(\theta)$  is sufficiently random. This result is a consequence of the random nature of both the ansatz structure and the initialization of parameters, and it makes no reference to the form of a potential target unitary to be learned. Here, we prove a complementary result. Namely, if the target unitary  $V$  is drawn from a sufficiently random ensemble (i.e., an ensemble that forms at least a 2-design), one also encounters a barren plateau, *irrespective of the training ansatz or the initialization of parameters*. Thus, regardless of the QML strategy that is employed, a typical scrambler will manifest a barren plateau.

Suppose one wants to learn an unknown target unitary  $V$  where all that is known is that it is drawn from an ensemble of scramblers  $\mathbb{V}$ . This corresponds to  $\mathbb{V}$  forming a 2-design as noted above. Consider learning  $V$  by variationally minimizing a cost  $C(\theta, V)$  of the general form in (4). The following proposition, which we prove in Appendix B of the Supplementary Material, establishes that the average gradient of the cost is zero.

**Proposition 1.** *The average partial derivative of  $C(\theta, V)$ , with respect to any parameter  $\theta_k$ , for an ensemble of target unitaries  $\mathbb{V}$  that form a 2-design, is given by*

$$\langle \partial_{\theta_k} C(\theta, V) \rangle_{\mathbb{V}} = 0. \quad (6)$$

Proposition 1 establishes that the gradient is unbiased, but this alone does not preclude the possibility of large variations in the gradient, and thus is insufficient to assess trainability. However, Chebyshev's inequality bounds the probability that the partial derivative of the cost function deviates from its mean value

$$P(|\partial_{\theta_k} C| \geq |x|) \leq \frac{\text{Var}_{\mathbb{V}}[\partial_{\theta_k} C]}{x^2} \quad (7)$$

in terms of the variance of the cost partial derivative for a typical target unitary,

$$\text{Var}_{\mathbb{V}}[\partial_{\theta_k} C] = \left\langle (\partial_{\theta_k} C(\theta, V))^2 \right\rangle_{\mathbb{V}} - \langle \partial_{\theta_k} C(\theta, V) \rangle_{\mathbb{V}}^2. \quad (8)$$

As a result, a vanishingly small  $\text{Var}_{\mathbb{V}}[\partial_{\theta_k} C]$  (combined with a vanishing average gradient) for all  $\theta_k$  would imply that the probability that the cost partial derivative is non-zero is vanishingly small for all parameters, i.e., the cost landscape forms a barren plateau.

Indeed, this behavior is precisely what we find here. As shown in Appendix B, we prove the following.

**Theorem 2.** *Consider a generic cost function  $C(\theta, V)$ , Eq. (4), to learn an  $n$ -qubit target unitary  $V$  using an arbitrary ansatz  $U(\theta)$ . The variance of the partial derivative of  $C(\theta, V)$ , with respect to any parameter  $\theta_k$ , for an ensemble of target unitaries  $\mathbb{V}$  that form a 2-design, is given by*

$$\text{Var}_{\mathbb{V}}[\partial_{\theta_k} C] = \left[ \frac{2 \text{Tr}[H^2]}{2^{2n} - 1} - \frac{2(\text{Tr}[H])^2}{2^n(2^{2n} - 1)} \right] \text{Var}_{\chi}[-iU\partial_{\theta_k}U^\dagger], \quad (9)$$

where  $\text{Var}_{\mathbb{V}}$  denotes the variance over the ensemble  $\mathbb{V}$ , and  $\text{Var}_{\chi}$  denotes the quantum-mechanical variance with respect to the ansatz-evolved state  $|\chi(\theta)\rangle = U(\theta)|\psi\rangle$ .

From Theorem 2, we derive the following corollary on the scaling of  $\text{Var}_{\mathbb{V}}[\partial_{\theta_k} C]$ .

**Corollary 3.** *Consider a generic cost function  $C(\theta, V)$ , Eq. (4), to learn an  $n$ -qubit target unitary  $V$ . Without loss of generality, the ansatz can be written in the form*

$$U(\theta) = \prod_{i=1}^N U_i(\theta_i) W_i, \quad (10)$$

where  $\{W_i\}$  is a chosen set of fixed unitaries and  $U_i(\theta_i) = \exp(-i\theta_i G_i)$  with  $G_i$  an Hermitian operator. If  $\text{Tr}[H^2] \in \mathcal{O}(2^n)$  and  $\|G_k^2\|_{\infty} \in \mathcal{O}(1)$ , then

$$\text{Var}_{\mathbb{V}}[\partial_{\theta_k} C] \in \mathcal{O}(2^{-n}). \quad (11)$$

We note that for practical cost functions the condition  $\text{Tr}(H^2) \in \mathcal{O}(2^n)$  holds. Similarly, standard ansätze use normalized generators and therefore it is reasonable to assume that  $\|G_k^2\|_{\infty} \in \mathcal{O}(1)$ ,  $\forall k$ . As such, we conclude that the variance in the gradient will in general vanish exponentially with  $n$ , the number of qubits in the system. We note that Appendix C extends this exponential scaling result to the generalized cost function in (5). Hence, these results establish that QML approaches to learn a typical target scrambling unitary, that is a unitary drawn from a 2-design, will exhibit a barren plateau.

*Numerical implementation for approximate scramblers.* Here we extend our results by numerically studying approximate scramblers. For concreteness, we now take a dynamical perspective and model a scrambler




FIG. 2. **Numerical Simulations of Approximate Scramblers.** a) A random cut of the landscape of the LHST cost function  $C_{\text{LHST}}(U, V)$  (defined in Appendix E) where  $V(g, t)$  is a randomly generated 9 qubit scrambler ( $n = 9$ ) modelled via Eq. (12) and  $U(g, t)$  is an ansatz of the same form. Here  $\epsilon$  is a noise parameter that determines the deviation of the ansatz parameters,  $\theta$ , from the target scrambler's parameters,  $\theta^{\text{target}}$ . Specifically, we set  $\theta_k = \theta_k^{\text{target}} + \epsilon R$ , where  $R$  is a random number between -1 and 1. The landscape for a weak scrambler with  $t = 1$  and  $g = 0.1$  ( $g = 5$ ) is plotted in yellow (red). Stronger scramblers with  $t = 15$  and  $g = 0.1$  ( $g = 5$ ) are plotted in green (blue). b) The variance in the gradient of a single local term of the LHST cost function,  $C_{\text{LHST}}^0(U, V)$  (defined in Appendix E), where  $V(g, t)$  is a scrambler modelled via Eq. (12) with  $g = 0.5$  (dashed) and  $g = 1$  (solid) and  $U(g, t)$  is an ansatz of the same form. In the first (second) row the variance is calculated over an ensemble of random target (ansatz) unitaries. In the first (second) column the variance in the gradient is plotted as a function of  $t$  ( $n$ ). As indicated by the dotted line, for large  $g$  and  $t$  the variance vanishes exponentially as  $\text{Var}[\partial_{\theta_k} C] \propto 2^{-2n}$ .

using a variant of the minimal model introduced in Ref. [56], where the scrambling unitary

$$V_S(g, t) := (V_2(g)V_1)^t \quad (12)$$

consists of alternating random gates and entangling layers. Specifically, the first layer is composed of a series of random single qubit rotations

$$V_1 = \prod_i R_x^i(\theta_x^i) R_y^i(\theta_y^i) R_z^i(\theta_z^i), \quad (13)$$

where  $R_k^i(\theta_k^i)$  is a rotation about the  $k = x, y, z$  axis of the  $i_{\text{th}}$  qubit and the  $\{\theta_k^i\}$  are randomly chosen angles between 0 and  $2\pi$ . The second layer is a global entangling gate

$$V_2(g) = \prod_{i < j} e^{-\frac{igZ_iZ_j}{\sqrt{n}}}, \quad (14)$$

where  $Z_k$  is a Pauli operator on the  $k_{\text{th}}$  qubit and  $n$  is the number of qubits. The degree to which  $V_S(t)$  is scrambling increases with time  $t$  (i.e. circuit depth), with the rate of increase determined by the entangling rate  $g$ .

We consider learning a scrambler modelled by  $V_S(g, t)$  using an ansatz of precisely the same structure. That is we generate a target scrambler  $V_S^{\text{target}}$  by randomly generating a set of single qubit rotation angles  $\theta_{\text{target}}$  and attempt to learn the angles using an ansatz of the form  $U(\theta) = V_S(g, t)$ . For concreteness, we suppose the local Hilbert-Schmidt cost [33], detailed in Appendix E, is used to learn  $V_S$ .

In Fig. 2(a) we plot a cross-section of the cost landscape in the region around the true parameters  $\theta_{\text{target}}$  that minimize the cost. As we increase the degree to which the target unitary is scrambling, by increasing the duration of evolution  $t$  and the strength of the entangling gates  $g$ , the variance in the cost visibly decreases. In the case of a highly scrambling unitary (blue) the majority of the landscape forms a *barren plateau* with only a *narrow gorge* where the cost dips down to its minimum. In contrast for weaker scramblers (yellow) the valley around the minimum is wider and the plateau more featured.

In Fig. 2(b) we plot the variance of the cost partial derivative as a function of scrambling time  $t$  (left) and system size  $n$  (right) for varying entangling rate  $g$ . For completeness and for comparison, the variance is calculated both over an ensemble  $\mathbb{V}$  of target unitaries (top), denoted  $\text{Var}_{\mathbb{V}}[\partial_{\theta_k} C]$ , and over an ensemble  $\mathbb{U}$  of random parameterized ansätze (bottom), denoted  $\text{Var}_{\mathbb{U}}[\partial_{\theta_k} C]$ . For sufficiently large  $t$  and  $g$  the target unitary is a perfect scrambler, and the variance of the partial derivative vanishes exponentially as  $\text{Var}_{\mathbb{V}}[\partial_{\theta_k} C] \propto 2^{-2n}$ . We note that this in fact decays faster than the minimal scaling predicted by Corollary 3, which is expected for typical scrambling ansätze (as we discuss in Appendix D). Once the scrambling time is sufficient for perfect scrambling,  $\text{Var}_{\mathbb{V}}[\partial_{\theta_k} C]$  saturates. For weaker scramblers  $\text{Var}_{\mathbb{V}}[\partial_{\theta_k} C]$  similarly decreases with system size but at a slower rate. The same behavior is seen for  $\text{Var}_{\mathbb{U}}[\partial_{\theta_k} C]$ , demonstrating a duality between averaging over targets and ansätze.

In Appendix F, we complement these numerical results for approximate scramblers with an analytic argument.



Specifically, we show that the exponential suppression of the variance of the gradient is preserved, as long as the target ensemble is sufficiently close to a unitary 2-design.

*Discussion.* Bob’s ability to decode Alice’s secret from observing relatively few emitted photons in the Hayden-Preskill thought experiment relies on the scrambling nature of the black hole’s unitary dynamics. This ensures that the information contained in Alice’s state is quickly distributed across the black hole. Our work establishes that the same scrambling nature inhibits Bob’s ability to learn  $U$  in the first place. Therefore, perhaps Alice’s secret is safer than previously thought.

We find that a barren plateau will be encountered when learning a typical random (or pseudo-random) unitary with any ansatz. This no-go result provides a fundamental limit on the learnability of random processes. Related complexity theoretic arguments precluding learning black hole unitaries depend on Alice’s restricted access to the full unitary [57]. Our work shows that, even with complete access, the unitary is unlearnable.

Thankfully, most physically interesting processes do not resemble a typical random (or pseudo-random) unitary. For example, the short time evolution of even chaotic systems form only a 1-design and consequently, as supported by our numerical results, may be learnable. Therefore the no-go theorem does not condemn QML but rather shows limits to its domain of applicability.

Crucially, Bob’s situation in the Hayden-Preskill thought experiment differs from other machine learning tasks. Elsewhere, he might be able to learn more information about the target unitary, and perhaps avoid a barren plateau. When attempting to learn the dynamics

of a black hole, Bob only knows that he needs to learn a scrambler and lacks a way to peek behind the event horizon to the black hole formation time, so this additional information is inaccessible.

Methods applicable when further information is attainable remain to be studied. For example, one might try adapting the ansatz based on partial information extracted about the target. Alternatively, for certain applications, one might instead learn a target by breaking the evolution down into shorter, more tractable, time steps. It may also be worth investigating whether barren plateaus can be avoided by only seeking partial information about the target unitary. More generally, while previous results [46–48] show the need to design clever ansätze, our results highlight the need to carefully consider the properties of the target unitary.

## ACKNOWLEDGMENTS

We thank Marco Cerezo for helpful discussions. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics QuantISED program under under Contract Nos. DE-AC52-06NA25396 and KA2401032 (ZH, AA, PJC, AA, ATS). BY acknowledges support of the U.S. Department of Energy, Office of Science, Basic Energy Sciences, Materials Sciences and Engineering Division, Condensed Matter Theory Program, and partial support from the Center for Nonlinear Studies. PJC and ATS acknowledge initial support from the Los Alamos National Laboratory (LANL) ASC Beyond Moore’s Law project.

- 
- [1] Patrick Hayden and John Preskill, “Black holes as mirrors: quantum information in random subsystems,” *Journal of High Energy Physics* **2007**, 120–120 (2007).
  - [2] Yasuhiro Sekino and L Susskind, “Fast scramblers,” *Journal of High Energy Physics* **2008**, 065–065 (2008).
  - [3] Juan Maldacena, Stephen H. Shenker, and Douglas Stanford, “A bound on chaos,” *Journal of High Energy Physics* **2016**, 106 (2016).
  - [4] Pavan Hosur, Xiao-Liang Qi, Daniel A. Roberts, and Beni Yoshida, “Chaos in quantum channels,” *Journal of High Energy Physics* **2016**, 4 (2016).
  - [5] Adam Nahum, Jonathan Ruhman, Sagar Vijay, and Jeongwan Haah, “Quantum entanglement growth under random unitary dynamics,” *Phys. Rev. X* **7**, 031016 (2017).
  - [6] Winton Brown and Omar Fawzi, “Decoupling with random quantum circuits,” *Communications in Mathematical Physics* **340**, 867–900 (2015).
  - [7] R. J. Lewis-Swan, A. Safavi-Naini, J. J. Bollinger, and A. M. Rey, “Unifying scrambling, thermalization and entanglement through measurement of fidelity out-of-time-order correlators in the dicke model,” *Nature Communications* **10**, 1581 (2019).
  - [8] Markus J. Klug, Mathias S. Scheurer, and Jörg Schmalian, “Hierarchy of information scrambling, thermalization, and hydrodynamic flow in graphene,” *Phys. Rev. B* **98**, 045102 (2018).
  - [9] Beni Yoshida and Alexei Kitaev, “Efficient decoding for the hayden-preskill protocol,” *arXiv preprint arXiv:1710.03363* (2017).
  - [10] Beni Yoshida and Norman Y. Yao, “Disentangling scrambling and decoherence via quantum teleportation,” *Phys. Rev. X* **9**, 011006 (2019).
  - [11] K. A. Landsman, C. Figgatt, T. Schuster, N. M. Linke, B. Yoshida, N. Y. Yao, and C. Monroe, “Verified quantum information scrambling,” *Nature* **567**, 61–65 (2019).
  - [12] Stephen H. Shenker and Douglas Stanford, “Black holes and the butterfly effect,” *Journal of High Energy Physics* **2014**, 67 (2014).
  - [13] A. Kitaev, “A simple model of quantum holography,” (2015), proceedings of the KITP Program: Entanglement in Strongly-Correlated Quantum Matter (Kavli Institute for Theoretical Physics, Santa Barbara).
  - [14] Daniel A. Roberts and Beni Yoshida, “Chaos and com-

- plexity by design,” *Journal of High Energy Physics* **2017**, 121 (2017).
- [15] Jordan Cotler, Nicholas Hunter-Jones, Junyu Liu, and Beni Yoshida, “Chaos, complexity, and random matrices,” *Journal of High Energy Physics* **2017**, 48 (2017).
- [16] Bruno Bertini and Lorenzo Piroli, “Scrambling in random unitary circuits: Exact results,” *Phys. Rev. B* **102**, 064305 (2020).
- [17] A. Chenu, I. L. Egusquiza, J. Molina-Vilaplana, and A. del Campo, “Quantum work statistics, loschmidt echo and information scrambling,” *Scientific Reports* **8**, 12634 (2018).
- [18] Nicole Yunger Halpern, “Jarzynski-like equality for the out-of-time-ordered correlator,” *Phys. Rev. A* **95**, 012120 (2017).
- [19] Nicole Yunger Halpern, Anthony Bartolotta, and Jason Pollack, “Entropic uncertainty relations for quantum information scrambling,” *Communications Physics* **2**, 92 (2019).
- [20] K. A. Landsman, C. Figgatt, T. Schuster, N. M. Linke, B. Yoshida, N. Y. Yao, and C. Monroe, “Verified quantum information scrambling,” *Nature* **567**, 61–65 (2019).
- [21] Bin Yan and Nikolai A Sinitsyn, “Recovery of damaged information and the out-of-time-ordered correlators,” *Physical Review Letters* **125**, 040605 (2020).
- [22] Qingling Zhu, Zheng-Hang Sun, Ming Gong, Fusheng Chen, Yu-Ran Zhang, Yulin Wu, Yangsen Ye, Chen Zha, Shaowei Li, Shaojun Guo, *et al.*, “Observation of thermalization and information scrambling in a superconducting quantum processor,” *arXiv preprint arXiv:2101.08031* (2021).
- [23] Xiao Mi, Pedram Roushan, Chris Quintana, Salvatore Mandra, Jeffrey Marshall, Charles Neill, Frank Arute, Kunal Arya, Juan Atalaya, Ryan Babbush, *et al.*, “Information scrambling in computationally complex quantum circuits,” *arXiv preprint arXiv:2101.08870* (2021).
- [24] John Preskill, “Quantum Computing in the NISQ era and beyond,” *Quantum* **2**, 79 (2018).
- [25] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd, “Quantum machine learning,” *Nature* **549**, 195–202 (2017).
- [26] Kerstin Beer, Dmytro Bondarenko, Terry Farrelly, Tobias J. Osborne, Robert Salzmänn, Daniel Scheiermann, and Ramona Wolf, “Training deep quantum neural networks,” *Nature Communications* **11**, 808 (2020).
- [27] Kyle Poland, Kerstin Beer, and Tobias J Osborne, “No free lunch for quantum machine learning,” *arXiv preprint arXiv:2003.14103* (2020).
- [28] Kunal Sharma, M Cerezo, Zoë Holmes, Lukasz Cincio, Andrew Sornborger, and Patrick J Coles, “Reformulation of the no-free-lunch theorem for entangled data sets,” *arXiv preprint arXiv:2007.04900* (2020).
- [29] 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**, 4213 (2014).
- [30] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik, “The theory of variational hybrid quantum-classical algorithms,” *New Journal of Physics* **18**, 023023 (2016).
- [31] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann, “A quantum approximate optimization algorithm,” *arXiv preprint arXiv:1411.4028* (2014).
- [32] J. Romero, J. P. Olson, and A. Aspuru-Guzik, “Quantum autoencoders for efficient compression of quantum data,” *Quantum Science and Technology* **2**, 045001 (2017).
- [33] Sumeet Khatri, Ryan LaRose, Alexander Poremba, Lukasz Cincio, Andrew T Sornborger, and Patrick J Coles, “Quantum-assisted quantum compiling,” *Quantum* **3**, 140 (2019).
- [34] R. LaRose, A. Tikku, É. O’Neel-Judy, L. Cincio, and P. J. Coles, “Variational quantum state diagonalization,” *npj Quantum Information* **5**, 1–10 (2018).
- [35] Andrew Arrasmith, Lukasz Cincio, Andrew T Sornborger, Wojciech H Zurek, and Patrick J Coles, “Variational consistent histories as a hybrid algorithm for quantum foundations,” *Nature communications* **10**, 1–7 (2019).
- [36] Marco Cerezo, Alexander Poremba, Lukasz Cincio, and Patrick J Coles, “Variational quantum fidelity estimation,” *Quantum* **4**, 248 (2020).
- [37] Kunal Sharma, Sumeet Khatri, Marco Cerezo, and Patrick J Coles, “Noise resilience of variational quantum compiling,” *New Journal of Physics* **22**, 043006 (2020).
- [38] Carlos Bravo-Prieto, Ryan LaRose, M. Cerezo, Yigit Subasi, Lukasz Cincio, and Patrick J. Coles, “Variational quantum linear solver: A hybrid algorithm for linear systems,” *arXiv:1909.05820* (2019).
- [39] M Cerezo, Kunal Sharma, Andrew Arrasmith, and Patrick J Coles, “Variational quantum state eigensolver,” *arXiv preprint arXiv:2004.01372* (2020).
- [40] Kentaro Heya, Ken M Nakanishi, Kosuke Mitarai, and Keisuke Fujii, “Subspace variational quantum simulator,” *arXiv preprint arXiv:1904.08566* (2019).
- [41] Cristina Cirstoiu, Zoe Holmes, Joseph Iosue, Lukasz Cincio, Patrick J Coles, and Andrew Sornborger, “Variational fast forwarding for quantum simulation beyond the coherence time,” *npj Quantum Information* **6**, 1–10 (2020).
- [42] Benjamin Commeau, Marco Cerezo, Zoë Holmes, Lukasz Cincio, Patrick J Coles, and Andrew Sornborger, “Variational hamiltonian diagonalization for dynamical quantum simulation,” *arXiv preprint arXiv:2009.02559* (2020).
- [43] Ying Li and Simon C Benjamin, “Efficient variational quantum simulator incorporating active error minimization,” *Physical Review X* **7**, 021050 (2017).
- [44] Suguru Endo, Jinzhao Sun, Ying Li, Simon C. Benjamin, and Xiao Yuan, “Variational quantum simulation of general processes,” *Phys. Rev. Lett.* **125**, 010501 (2020).
- [45] Xiao Yuan, Suguru Endo, Qi Zhao, Ying Li, and Simon C Benjamin, “Theory of variational quantum simulation,” *Quantum* **3**, 191 (2019).
- [46] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven, “Barren plateaus in quantum neural network training landscapes,” *Nature Communications* **9**, 4812 (2018).
- [47] Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles, “Cost-function-dependent barren plateaus in shallow quantum neural networks,” *arXiv preprint arXiv:2001.00550* (2020).
- [48] Kunal Sharma, Marco Cerezo, Lukasz Cincio, and Patrick J Coles, “Trainability of dissipative perceptron-based quantum neural networks,” *arXiv preprint arXiv:2005.12458* (2020).
- [49] Samson Wang, Enrico Fontana, M Cerezo, Kunal

- Sharma, Akira Sone, Lukasz Cincio, and Patrick J Coles, “Noise-induced barren plateaus in variational quantum algorithms,” [arXiv preprint arXiv:2007.14384](#) (2020).
- [50] M Cerezo and Patrick J Coles, “Impact of barren plateaus on the hessian and higher order derivatives,” [arXiv preprint arXiv:2008.07454](#) (2020).
- [51] Andrew Arrasmith, M. Cerezo, Piotr Czarnik, Lukasz Cincio, and Patrick J Coles, “Effect of barren plateaus on gradient-free optimization,” [arXiv preprint arXiv:2011.12245](#) (2020).
- [52] Edward Grant, Leonard Wossnig, Mateusz Ostaszewski, and Marcello Benedetti, “An initialization strategy for addressing barren plateaus in parametrized quantum circuits,” *Quantum* **3**, 214 (2019).
- [53] Guillaume Verdon, Michael Broughton, Jarrod R McClean, Kevin J Sung, Ryan Babbush, Zhang Jiang, Hartmut Neven, and Masoud Mohseni, “Learning to learn with quantum neural networks via classical neural networks,” [arXiv preprint arXiv:1907.05415](#) (2019).
- [54] Tyler Volkoff and Patrick J Coles, “Large gradients via correlation in random parameterized quantum circuits,” *Quantum Science and Technology* **6**, 025008 (2021).
- [55] A. Larkin and Y. N. Ovchinnikov, “Quasiclassical method in the theory of superconductivity,” *Sov. Phys. JETP* **28**, 1200 (1969).
- [56] Ron Belyansky, Przemyslaw Bienias, Yaroslav A. Kharkov, Alexey V. Gorshkov, and Brian Swingle, “Minimal model for fast scrambling,” *Phys. Rev. Lett.* **125**, 130601 (2020).
- [57] Daniel Harlow and Patrick Hayden, “Quantum computation vs. firewalls,” *Journal of High Energy Physics* **2013**, 85 (2013).
- [58] Zbigniew Puchała and Jarosław Adam Miszczyk, “Symbolic integration with respect to the haar measure on the unitary groups,” *Bulletin of the Polish Academy of Sciences Technical Sciences* **65**, 21–27 (2017).
- [59] Bin Yan, Lukasz Cincio, and Wojciech H Zurek, “Information scrambling and loschmidt echo,” *Physical Review Letters* **124**, 160603 (2020).
- [60] Fernando G S L Brandão, Aram W Harrow, and Michał Horodecki, “Local random quantum circuits are approximate Polynomial-Designs,” *Communications in Mathematical Physics* **346**, 397–434 (2016).