Power and limitations of learning quantum dynamics incoherently

Sofiene Jerbi,[1,2,3,*] Joe Gibbs,[4,5] Manuel S. Rudolph,[6] Matthias C. Caro,[7,3] Patrick J. Coles,[1,8] Hsin-Yuan Huang,[7,9] and Zoë Holmes[1,6]

[1]*Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico, USA*
[2]*Institute for Theoretical Physics, University of Innsbruck, Innsbruck, Austria*
[3]*Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, Berlin, Germany*
[4]*Department of Physics, University of Surrey, Guildford GU2 7XH, United Kingdom*
[5]*AWE, Aldermaston, Reading RG7 4PR, United Kingdom*
[6]*Institute of Physics, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland*
[7]*Institute for Quantum Information and Matter, Caltech, Pasadena, California, USA*
[8]*Normal Computing Corporation, New York, New York, USA*
[9]*Department of Computing and Mathematical Sciences, Caltech, Pasadena, California, USA*

(Received 29 June 2023; accepted 5 February 2026; published 11 May 2026)

Quantum process learning is emerging as an important tool to study quantum systems. While studied extensively in coherent frameworks, where the target and model system can share quantum information, less attention has been paid to whether the dynamics of quantum systems can be learned without the system and target directly interacting. Such incoherent frameworks are practically appealing since they open up methods of transpiling quantum processes between the different physical platforms without the need for technically challenging hybrid-entanglement schemes. Here, we provide bounds on the sample complexity of learning unitary processes incoherently by analyzing the number of measurements that are required to emulate well-established coherent learning strategies. We prove that if arbitrary measurements are allowed, then any efficiently representable unitary can be efficiently learned within the incoherent framework; however, when restricted to shallow-depth measurements, logarithmic depth unitaries can be learned but there exist linear-depth unitaries that are unlearnable. We demonstrate our incoherent learning algorithm for low-entangling unitaries by successfully learning a 16-qubit unitary on `ibmq_kolkata`, and further demonstrate the scalability of our proposed algorithm through extensive numerical experiments.

DOI: 10.1103/s5xy-lk91

I. INTRODUCTION

Classical computing power and classical machine learning are routinely used to process data from quantum experiments. However, as these are *quantum* experiments, it would perhaps be more natural to use a *quantum* computer to process their data. Within this line of thought, a particularly promising application of quantum hardware is *quantum process learning* [1–16]. At its simplest, the process to be learned will be the unitary dynamics of an experimental system that one wishes to study, and the machine learning model will be a quantum circuit representation or a classical description of the target process. In this manner, quantum process learning is a means of digitizing an analog quantum process and uploading it to a quantum or classical computer for further study.

Quantum process learning has thus far been predominantly investigated in the *coherent* setting. That is, it is assumed that the target and model system can coherently interact and quantum information can be shared between them. While the coherent setting is theoretically powerful (in the sense that one can learn efficiently implementable unitaries using only a polynomial number of "easy-to-prepare" product states [13]), engineering the coherent interaction of the target and model system can be experimentally challenging. In particular, if the target and model system are made up of physically different constituents, then hybrid-entanglement engineering [17–22] is required (see Fig. 1). More fundamentally, given the need for the target and model system to coherently interact, the model system cannot be purely classical.

The *incoherent* setting, where the training is performed using classical data (i.e., numerical values collected from measurement outcomes; see Fig. 1), is considerably more practical. Crucially, as there is no need for the target system and model system to physically interact in this case, the target system and model system can be radically different. Moreover, if the target unitary is classically simulable, the trained unitary could be a fully classical model such as those enabled by tensor network methods [23,24]. The appeal of this flexibility is reflected in the growing body of literature studying so-called "distributed learning" benchmarking protocols [25], as well as various proposed algorithms for Hamiltonian

*Contact author: sofiene.jerbi@fu-berlin.de

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

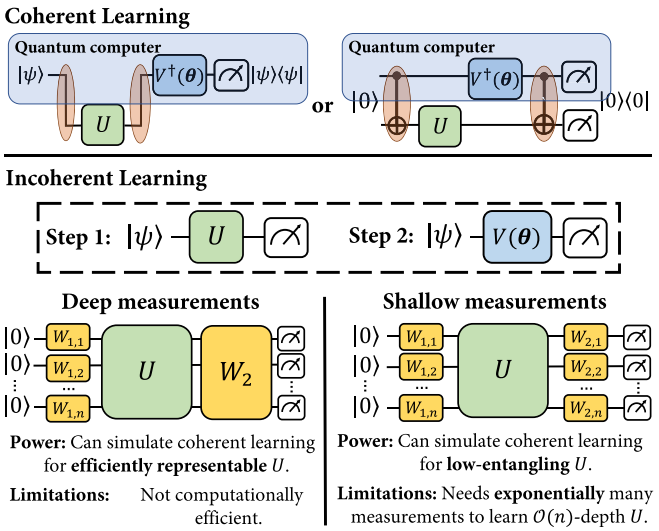


FIG. 1. Coherent vs incoherent unitary learning. We consider the task of learning a target unitary U using a variational ansatz $V(\boldsymbol{\theta})$ in an incoherent setting where quantum information is not shared between the quantum computer running $V(\boldsymbol{\theta})$ and the system implementing U. This setting is divided into two phases: (1) a measurement phase where information is collected about U and (2) a training phase where this information is used to train $V(\boldsymbol{\theta})$.

learning/verification that implicitly assume an incoherent setting [26–36].

Recently, the power of the incoherent setting has begun to be investigated from a learning theory perspective. For example, Ref. [14] develops an algorithm for predicting local properties of the output states of an unknown process using only local measurements, and Ref. [37] investigates the power of the incoherent setting in the context of hypothesis testing. Here, we take a different approach and investigate the extent to which the full coherent setting can be emulated within the incoherent setting in the context of learning unitary processes.

We start by showing that if arbitrary measurements can be performed on the system, then, using results from random measurement theory [38,39] and covering net arguments [12], it is possible to incoherently learn arbitrary efficiently representable unitaries using Clifford shadows and an efficient number of samples. However, there are two crucial caveats to this first result. First, Clifford shadows require implementing circuits of linear depth and so are impractical on many near-term platforms. More fundamentally, there are good reasons to believe (due to the globality of the cost and difficulties surrounding computing it efficiently on quantum or classical hardware) that the algorithm we use for our sampling complexity proof is not computationally efficient in general.

Thus, the core of this paper investigates the more practical situation where the experimentalist is limited to performing *shallow* measurements, i.e., those that can be implemented with shallow circuits. In this case, we prove that low-entangling unitaries can be efficiently learned using only product measurements. However, arbitrary unitaries cannot be learned in this setting even if they are linear depth. This provides a fundamental limit to quantum process learning in the incoherent paradigm. In parallel, our positive result for low-entangling unitaries opens up a family of tools for studying the short-time dynamics of locally interacting systems. We demonstrate the scalability, and plausible computational efficiency, of these tools via numerical simulations and a 16-qubit implementation on a quantum computer (`ibmq_kolkata`).

II. BACKGROUND

In this work, we consider the task of learning to optimize the parameters $\boldsymbol{\theta}$ of $V(\boldsymbol{\theta})$, an n-qubit parametrized unitary, such that, for the optimized parameters $\boldsymbol{\theta}_{\text{opt}}$, $V(\boldsymbol{\theta}_{\text{opt}})$ well approximates an unknown n-qubit unitary $U \in \mathcal{U}(\mathbb{C}^{2^n})$. The Hilbert-Schmidt inner product between U and $V(\boldsymbol{\theta})$ provides a natural measure of the success of the training. Thus, our goal is to minimize the Hilbert-Schmidt test cost [4],

$$C_{\text{HST}}(\boldsymbol{\theta}) := 1 - \frac{1}{d^2}|\text{Tr}[U^\dagger V(\boldsymbol{\theta})]|^2. \tag{1}$$

This cost is operationally meaningful in virtue of its equivalence to the average fidelity between a state $|\psi\rangle$ evolved under U and a state $|\psi\rangle$ evolved under $V(\boldsymbol{\theta})$. That is,

$$C_{\text{HST}}(\boldsymbol{\theta}) = \frac{d+1}{d}\mathbb{E}_{|\psi\rangle\sim\text{Haar}_n}[1 - |\langle\psi|U^\dagger V(\boldsymbol{\theta})|\psi\rangle|^2], \tag{2}$$

where $|\psi\rangle$ is drawn from the (Haar) uniform distribution of states. This suggests an alternative approach to training, where the learner has access to a training data set consisting of input-output pairs of pure n-qubit states,

$$\mathcal{D}_Q(N) = \{(|\psi^{(j)}\rangle, U|\psi^{(j)}\rangle)\}_{j=1}^N, \tag{3}$$

where the N input states $|\psi^{(j)}\rangle$ are drawn independently from the distribution Q. In particular, taking Q to be the Haar distribution, $V(\boldsymbol{\theta})$ can be trained using the loss

$$C_{\mathcal{D}_Q(N)}(\boldsymbol{\theta}) := 1 - \frac{1}{N}\sum_{i=1}^N |\langle\psi^{(i)}|U^\dagger V(\boldsymbol{\theta})|\psi^{(i)}\rangle|^2. \tag{4}$$

Prior results on *in-distribution* generalization give bounds on the number of training pairs N needed to ensure that this training loss is faithful to the Hilbert-Schmidt cost. In particular, for a variational circuit $V(\boldsymbol{\theta})$, with T trainable gates, $N \sim T\log(T)$ training pairs suffice [12].

In fact, more recent results on *out-of-distribution* generalization [13] imply that simple products of single-qubit stabilizer states suffice to generalize well on globally Haar-random states. That is, the training ensemble Q can be taken to be the uniform distribution over the set of stabilizer states $\text{Stab}_1^{\otimes n} := \{|0\rangle, |1\rangle, |+\rangle, |-\rangle, |i\rangle, |-i\rangle\}^{\otimes n}$. This is practically useful because, in contrast to Haar-random states that require deep circuits and a complete gate set, random stabilizer product states are classically simulatable and can be prepared using a single layer of single-qubit rotations.

In practice, to avoid cost-function-dependent barren plateaus [40], it is necessary to use a local variant of the training loss. A natural choice using product state training data is

$$C_N^\ell(\boldsymbol{\theta}) = 1 - \frac{1}{N}\sum_{j=1}^N \text{Tr}[|\psi^{(j)}\rangle\langle\psi^{(j)}|U^\dagger V(\boldsymbol{\theta})H_\ell^{(j)}V^\dagger(\boldsymbol{\theta})U]. \tag{5}$$

Here, $|\psi^{(j)}\rangle = \bigotimes_{i=1}^{n} |\psi_i^{(j)}\rangle \sim \mathrm{Stab}_1^{\otimes n}$ for all j and $H_\ell^{(j)} = \frac{1}{n}\sum_{i=1}^{n} |\psi_i^{(j)}\rangle\langle\psi_i^{(j)}| \otimes \mathbb{1}_{\bar{i}}$ is a 1-local measurement (i.e., each term acts nontrivially on at most one qubit). This local cost is faithful to the global product state training cost in the sense that it vanishes under the same conditions [4], but enjoys trainability guarantees for shallow-depth circuits [40]. More concretely, in the coherent setting, the following bound holds:

$$C_{\mathrm{HST}}(\boldsymbol{\theta}) \leqslant \frac{2n(d+1)}{d} C_N^{\ell}(\boldsymbol{\theta}) + \mathcal{O}\left(n\sqrt{\frac{T\log(T)}{N}}\right),$$

with high probability over the choice of product state training data of size N [13]. Thus, the local-product state cost in Eq. (5) can be used to indirectly train the Hilbert-Schmidt test cost and thereby learn the target unitary U.

It is natural to compute the product state training cost, and its local variants, in the coherent framework. This can be done via a Loschmidt echo-type circuit as sketched in Fig. 1 (the top left circuit). While such circuits are relatively straightforward to run on a single quantum device, they are much harder to implement if the system implementing the target unitary and the platform implementing the trained unitary are distinct. In such cases, experimentally challenging hybrid entanglement is required to pass the quantum information between the target and learning systems. This motivates considering the incoherent learning paradigm where the target and model systems do not need to interact. The key question tackled in this paper is the following: Is it possible to emulate the coherent learning setting using only incoherent learning protocols? More precisely, is it possible to rigorously estimate the training losses used coherently, without the target and test system directly interacting?

III. RESULTS

A. Deep measurements

Let us start by considering what can be computed in the incoherent learning paradigm if we suppose that arbitrary measurements can be performed on the system. In this case, we can use random Clifford measurements to generate the Clifford shadow [38] of each of the N output states in our training ensemble. This allows us to compute the fidelities of our target output states with an exponential number of guess output states (i.e., parameter settings for the parametrized circuit) after collecting a polynomial-sized shadow. Combining this observation with a covering set argument of all efficiently representable unitaries, we show it is possible to train a variational unitary $V(\boldsymbol{\theta})$ to achieve $C_{\mathrm{HST}}(\boldsymbol{\theta}) \leqslant \epsilon$ for any target accuracy ϵ with a polynomial sampling complexity. As proven and stated more formally in Appendix B 2 a, this is captured by the following theorem [41].

Theorem 1 (Power of incoherent learning with deep measurements (informal)). For any efficiently representable n-qubit unitary U, that is, a unitary that can be implemented using a quantum circuit of size $\mathrm{poly}(n)$, there exists a parametrized unitary $V(\boldsymbol{\theta})$ such that at most $\mathrm{poly}(n)$ calls to U and deep measurements are required to incoherently train $V(\boldsymbol{\theta})$ to achieve $1/\mathrm{poly}(n)$ Hilbert-Schmidt cost between $V(\boldsymbol{\theta})$ and U.

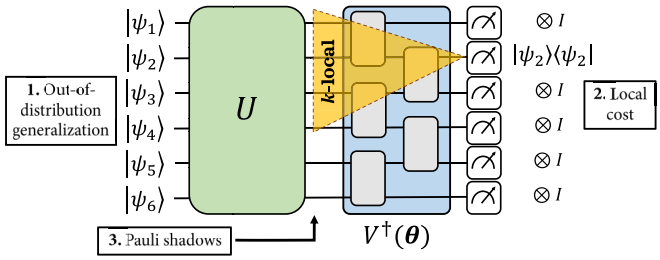


FIG. 2. Learning low-entangling unitaries using Pauli shadows. Our learning protocol for low-entangling unitaries makes use of three results. (1) Out-of-distribution generalization allows to consider locally scrambled training states with generalization guarantees to global Haar-random states. (2) Cost concentration results allow us to consider a local version of the training loss. (3) For shallow-depth ansätze $V(\boldsymbol{\theta})$ [e.g., depth $\mathcal{O}(\log(n))$], the observables appearing in the training loss remain local [$k \in \mathcal{O}(\log(n))$], and therefore Pauli shadow can efficiently simulate the loss incoherently.

It follows that if the target system allows for a Clifford gate set and long coherence times (all n-qubit Clifford operations can be implemented in depth $\mathcal{O}(n)$ [42]), then coherent learning protocols are not strictly needed for quantum process learning. However, there are two important caveats here. First, Clifford shadows are only well suited to compute global fidelities and therefore a global cost [of the form of Eq. (4)]. However, global costs exhibit barren plateaus. Thus, using generic ansätze, the time complexity of this approach will typically be exponential [40, 43]. Second, while Clifford shadows are guaranteed information theoretically to allow the computation of fidelities with other quantum states, this computation is by no means guaranteed to be efficient. When computing fidelities using Clifford shadows, we are indeed considering observables of the form $V(\boldsymbol{\theta})|\psi^{(i)}\rangle\langle\psi^{(i)}|V^\dagger(\boldsymbol{\theta})$, which cannot be efficiently represented/computed classically for deep circuits $V(\boldsymbol{\theta})$. Even when we consider access to a quantum computer to compute the overlap between these operators, we face the problem that classical shadows are completely unphysical operators that cannot be efficiently encoded into quantum states (see Appendix B 3 for an in-depth discussion).

These two constraints limit the breadth of applicability of Theorem 1. Hence, we will now turn our attention to the focus of this work: the power and limitation of the incoherent framework if we limit ourselves to shallow measurements.

B. Shallow measurements

We can follow a similar approach as with Clifford shadows, but this time restrict ourselves to Pauli shadows. In this case, since the sample complexity of Pauli shadows scales exponentially with the locality of the operators measured, we can only efficiently compute the effect of measuring operators whose locality scales at most logarithmically with the system size. Thus, we can only emulate the coherent setting and estimate the local version of the loss function $C_N^{\ell}(\boldsymbol{\theta})$ [see Eq. (5)] for low-entangling unitaries that transform 1-local operators into $O(\log(n))$-local operators, as sketched in Fig. 2 and further detailed in Appendix B 2 b.

Theorem 2 (Power of incoherent learning with shallow measurements (informal)). For any *low-entangling* n-qubit

unitary U, that is, a unitary that transforms 1-local operators into $O(\log(n))$-local operators, there exists a parameterized unitary $V(\boldsymbol{\theta})$ such that at most poly(n) calls to U and shallow measurements are required to incoherently train $V(\boldsymbol{\theta})$ to achieve $1/\text{poly}(n)$ Hilbert-Schmidt cost between $V(\boldsymbol{\theta})$ and U.

The core intuition for this result is the observation, also seen elsewhere [44–46], that local tomography schemes become efficient for characterizing low-entangling unitaries. More formally, our proof strategy for Theorem 2 again uses a covering net argument to argue that an exhaustive search strategy requires only polynomial samples. However, as opposed to the deep-measurement protocol, we expect that the shallow-measurement protocol is also computationally efficient in simulating the coherent training loss $C_N^{\ell}(\boldsymbol{\theta})$. Since the observables $V(\boldsymbol{\theta})H_{\ell}^{(j)}V^{\dagger}(\boldsymbol{\theta})$ are $\mathcal{O}(\log(n))$-local, they can be represented and computed efficiently classically. This allows one to train $V(\boldsymbol{\theta})$ classically (see Appendix B 3 for an in-depth discussion) and thereby learn a classical representation of the target process. If $V(\boldsymbol{\theta})$ is parametrized such that it can be implemented on a different quantum system, this approach further provides a way of transpiling the quantum process between quantum platforms.

This learning algorithm breaks down if the target unitary is too entangling as the sample complexity of our method scales exponentially with the locality of the operator $V(\boldsymbol{\theta})H_{\ell}V(\boldsymbol{\theta})^{\dagger}$. It turns out that this is not a limitation of our algorithm but rather a fundamental limitation of incoherent learning with shallow measurements. In particular, we prove that there exist efficiently implementable unitaries that require exponentially many samples to be learned in the incoherent framework with only product measurements. This is captured by the following theorem.

Theorem 3 (Limit of incoherent learning with shallow measurements (informal)). There exist efficiently representable unitaries U such that when restricted to performing only shallow measurements, at least $2^{\Omega(n)}$ calls to U are required to achieve a Hilbert-Schmidt cost $C_{\text{HST}}(\boldsymbol{\theta}) \leqslant 1/4$ in the incoherent framework.

In Appendix C, we prove this theorem by reduction of the task of distinguishing between orthogonal unitaries that can be used to prepare Greenberger-Horne-Zeilinger (GHZ)-like states. We show that the distinguishing task requires exponential sample complexity for shallow measurements. Moreover, the unitaries we consider are implementable using a linear-depth circuit (see Fig. 7).

As any unitary can be implemented via the evolution under some Hamiltonian H for some time t, and arbitrary sparse Hamiltonians are learnable in the incoherent setting using only Pauli measurements [47], this no-go theorem further highlights a key difference between the process learning and Hamiltonian learning frameworks. Namely, access to the continual time dynamics of a system is substantially more powerful than access to the dynamics at a single time. For more discussion on this point, see the end of Appendix C 1.

C. Implementations

Quantum processing unit (QPU) implementation. We demonstrate our shallow-measurement incoherent learning algorithm by learning a 16-qubit low-entangling unitary on

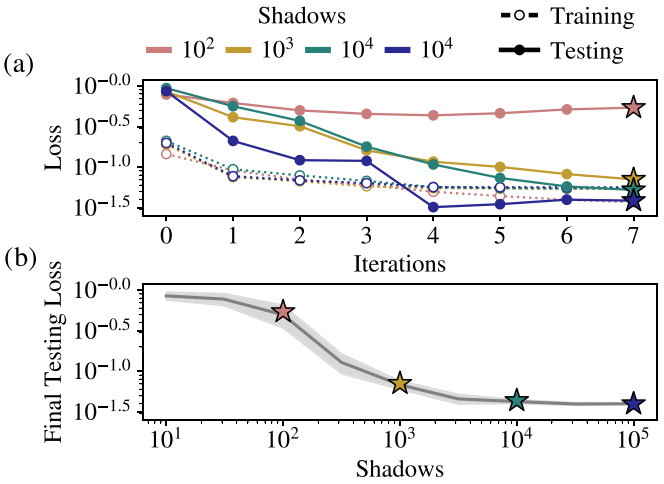


FIG. 3. Incoherent learning of 16-qubit unitary on `ibmq_kolkata`. A short-time Trotteriation of a 16-qubit transverse-field Ising model is incoherently compiled using shallow measurements. The PQC's rotation angles are initialized from the uniform distribution $\theta_i \sim \mathcal{U}(-1, 1)$. Panel (a) shows typical training and testing curves. In panel (b), the testing error upon convergence is shown to decrease for an increasing number of shadows trained on. The mean and standard deviation shown here are produced from ten repeats of training on independent sets of shadows. The star markers emphasize that the testing losses plotted in panel (b) are produced from the final converged parameters shown in panel (a).

`ibmq_kolkata`. The target unitary is a first-order short-time ($\Delta t = 0.1$) Trotterization for the transverse-field Ising model,

$$H_{\text{Ising}} = \sum_{i=0}^{14} Z_i Z_{i+1} + \sum_{i=0}^{15} \alpha_i X_i, \tag{6}$$

where α_i are randomly drawn from the normal distribution $\mathcal{N}(\mu = 0, \sigma = 0.5)$. As described in detail in Appendix B 2 b, a Pauli classical shadow of size 280 000 is created for only two training states $\{U(\Delta t)|\psi_i\rangle\}_{i=0}^{1}$, where $|\psi_i\rangle$ is a Haar-random product state [48]. These shadows are used to compute the local training cost, which is minimized classically by training a simulated parametrized quantum circuit (PQC) with the same gate structure as the target Trotterized unitary. The testing error is measured by an approximation to the Hilbert-Schmidt test, computed using product state overlaps [13]. We compute the testing loss $C_{\mathcal{D}_{\text{Haar}_1^{\otimes n}(N)}}(\boldsymbol{\theta})$ for a set of $N = 100$ random product states. The different sized shadows used in Fig. 3 are created from randomly chosen subsets of the full 280 000 sample shadow originally collected on the quantum computer.

As shown in Fig. 3, for small shadow sizes the final testing loss decreases (signifying improved generalization) as the shadow size increases. The relatively large training error in panel (a) can be attributed to hardware noise. For shadows above $\sim 10^4$, the improvement to the final testing error plateaus, indicating that a shadow of $\sim 10^4$ suffices to well reproduce the true underlying output states.

Numerical simulations. To investigate its scalability, we classically simulated our incoherent learning algorithm for the

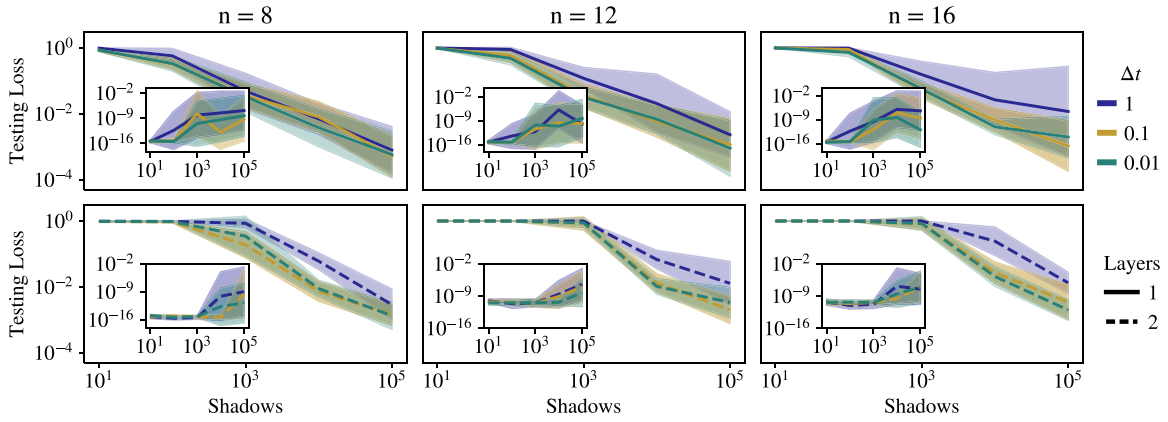


FIG. 4. Scalability of incoherent learning. For a range of system sizes (n), shadow sizes (M), Trotter time steps (Δt), and number of ansatz layers, we perform the incoherent learning algorithm using shallow measurements. For each configuration, we save the final testing loss (a product state overlap approximation of the Hilbert-Schmidt test), and ten repeats generate the mean and standard deviation plotted. The insets show the corresponding training loss upon termination of the optimization.

evolution induced by the Heisenberg model,

$$H_{\text{Heis}} = \sum_{i=0}^{n-1} X_i X_{i+1} + Y_i Y_{i+1} + Z_i Z_{i+1}. \tag{7}$$

We learn a first-order (one-layer) and second-order (two-layer) Trotterization of H_{Heis} using an ansatz of the same structure and using two product training states $\{U(\Delta t)|\psi_i\rangle\}_{i=0}^{1}$. Again, the training loss function is minimized with gradient-based methods by the Limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) optimizer. For these classical simulations, the gradient calculation is accelerated by automatic differentiation provided by the JAX library [49]. Figure 4 compares the final testing loss found for a range of system sizes, shadow sizes, and number of layers in the ansatz.

We find that for increasing system sizes, there is no significant decline in the testing performance. Conversely, as the degree of entanglement in the target unitary and ansatz increases, by increasing Δt and the number of layers, the testing performance degrades. This can be explained by the growth of the light cone of the local measurement operators under $V^\dagger(\boldsymbol{\theta})$ as sketched in Fig. 2. Interestingly, the shadow size required depends more strongly on Δt than on the number of layers, indicating that it is the effective size (rather than absolute size) of the light cone that matters. We discuss this further in Appendix B 2 c.

IV. DISCUSSION

Theorem 3 establishes a fundamental limitation on our ability to incoherently learn a target process. Crucially, our success metric is an average fidelity between the true and learned output states. This is to be contrasted with Ref. [38], which presents a time- and sample-efficient incoherent learning algorithm for predicting local measurement outcomes on the outputs of arbitrary processes. Thus, our no-go theorem highlights that predicting output states is substantially harder than learning to predict local measurement outcomes.

This makes it especially noteworthy that it is possible to fully emulate the coherent setting in the case of *low-entangling* unitaries. Crucially, in this case our algorithm uses only product training states and simple Pauli measurements, making it Noisy Intermediate-Scale Quantum (NISQ) friendly [50]. We have proven (Theorem 2) that the sample complexity of this approach scales polynomially with the size of the target unitary. While the computational complexity of this algorithm is unknown, we have good reasons to expect it to be efficient since it uses a trainable local cost [40]. This optimism is supported by our 16-qubit implementation on real quantum hardware. In contrast, the significance of our polynomial sample complexity lower bound for general measurements (Theorem 1) is more subtle as the time complexity is not expected to be efficient in this case (see Appendix B 3 for more details).

Here, as is common elsewhere [4,7,13,14,51–53], we tacitly assume that one wants to learn the output of a quantum process over random input states (or, equivalently, maximize the Hilbert-Schmidt inner product). However, in practice one may be interested in predicting the outputs for a given set of physically interesting input states [15]. Generalizing our results to this setting is an important direction for future work.

ACKNOWLEDGMENTS

S.J. was supported by the U.S. Department of Energy (DOE) through a quantum computing program sponsored by the Los Alamos National Laboratory (LANL) Information Science and Technology Institute. S.J. acknowledges support from the Austrian Science Fund (FWF) through the projects Grant No. DK-ALM:W1259-N27 and SFB BeyondC F7102. S.J. also acknowledges the Austrian Academy of Sciences as a recipient of the DOC Fellowship. M.C.C. was supported by a DAAD PRIME fellowship. P.J.C. acknowledges initial support from the Los Alamos National Laboratory (LANL) ASC Beyond Moore's Law project and subsequent support by the U.S. DOE, Office of Science, Office of Advanced Scientific Computing Research, under the Accelerated Research in Quantum Computing (ARQC) program. H.-Y.H. was supported by a Google Ph.D. Fellowship. Z.H. acknowledges initial support from the LANL Mark Kac Fellowship and subsequent support from the Sandoz Family Foundation-Monique de Meuron Program for Academic Promotion.

DATA AVAILABILITY

No data were created or analyzed in this study.

APPENDIX A: PRELIMINARIES

We consider a unitary compilation scenario where the training data about the target unitary U are not given coherently as $\{|\psi_{\text{in}}\rangle, U|\psi_{\text{in}}\rangle\}_{|\psi_{\text{in}}\rangle}$, but rather as incoherent (classical) data $\{|\psi_{\text{in}}\rangle, m\}_{(|\psi_{\text{in}}\rangle,F)}$, where m is the (random) outcome of measuring $|\psi_{\text{in}}\rangle$ using the positive operator-valued measure (POVM) $F = \{w_m 2^n |\phi_m\rangle\langle\phi_m|\}_m$, and where $|\psi_{\text{in}}\rangle$ is given by an efficient classical description. We investigate the power and limitations of different choices of input state distributions and POVMs.

APPENDIX B: THE POWER OF INCOHERENT LEARNING

We propose algorithms for learning an n-qubit unitary U using incoherent data. In this section, we describe how to generate this incoherent data, how to use it to train a parametrized ansatz $V(\boldsymbol{\theta})$, and study the sample complexity of our proposed algorithms in order to achieve a certain learning performance.

1. Summary of the algorithms

In order to learn a target unitary U, the parametrized ansatz $V(\boldsymbol{\theta})$ is trained via optimization of a training loss that is representative of the overlap between U and $V(\boldsymbol{\theta})$. For reasons explained below, we consider different training losses for the incoherent setting using deep measurements and the setting using shallow measurements. In the deep measurement setting, we consider the training loss:

$$C_{\mathcal{D}(N)}(\boldsymbol{\theta}) = 1 - \frac{1}{N}\sum_{i=1}^{N} |\langle\psi^{(i)}|U^\dagger V(\boldsymbol{\theta})|\psi^{(i)}\rangle|^2, \tag{B1}$$

where, for all $j = 1, \ldots, N$, $|\psi^{(j)}\rangle = \bigotimes_{i=1}^n |\psi_i^{(j)}\rangle \sim \text{Stab}_1^{\otimes n}$ is a tensor product of random single-qubit stabilizer states. As for the shallow-measurement setting, we look into the local version of this training loss:

$$C_N^\ell(\boldsymbol{\theta}) = 1 - \frac{1}{N}\sum_{j=1}^{N} \text{Tr}[|\psi^{(j)}\rangle\langle\psi^{(j)}|U^\dagger V(\boldsymbol{\theta}) H_\ell^{(j)} V(\boldsymbol{\theta})^\dagger U], \tag{B2}$$

where $H_\ell^{(j)} = \frac{1}{n}\sum_{i=1}^n |\psi_i^{(j)}\rangle\langle\psi_i^{(j)}| \otimes \mathbb{1}_{\bar{i}}$.

The goal of our incoherent learning algorithm is to gather incoherent data that can be used to "simulate" a coherent learning procedure. More precisely, we want to incoherently estimate the training losses above to a desirable precision $\varepsilon > 0$, for all parameter assignments $\boldsymbol{\theta}$. This then allows us to train the parameters $\boldsymbol{\theta}$ similarly to a coherent setting.

Deep-measurement algorithm. For each of the N training states $U|\psi^{(j)}\rangle$, we first gather an M-shot Clifford shadow. That is, we use M copies of the state $U|\psi^{(j)}\rangle$ and, on each copy indexed by $m = 1, \ldots, M$, we apply a randomly chosen Clifford unitary $W_m^{(j)}$ and measure in the computational basis. On doing so, we obtain a basis state $|b_m^{(j)}\rangle$ with probability $|\langle b_m^{(j)}|W_m^{(j)}U|\psi^{(j)}\rangle|^2$. The M-shot Clifford shadow [38] of the state $U|\psi^{(j)}\rangle$ is then given by the collection $\{\rho_m^{(j)}\}_{m=1}^M$, for

$$\rho_m^{(j)} = (2^n + 1) W_m^\dagger |b_m^{(j)}\rangle\langle b_m^{(j)}| W_m - \mathbb{1}. \tag{B3}$$

From here, no more interaction with the target system is required. Clifford shadows are particularly suited to estimate global observables with support on a small subspace (i.e., observables O such that $\text{Tr}[O^2]$ is small). To see how this makes them interesting for estimating the training loss $C_{\mathcal{D}(N)}(\boldsymbol{\theta})$ defined in Eq. (B1), we rewrite it as

$$C_{\mathcal{D}(N)}(\boldsymbol{\theta}) = 1 - \frac{1}{N}\sum_{j=1}^{N} \text{Tr}[U|\psi^{(j)}\rangle\langle\psi^{(j)}|U^\dagger\, O^{(j)}(\boldsymbol{\theta})], \tag{B4}$$

where we have absorbed the parametrized unitary $V(\boldsymbol{\theta})$ into the observable

$$O^{(j)}(\boldsymbol{\theta}) = V(\boldsymbol{\theta})|\psi^{(j)}\rangle\langle\psi^{(j)}|V^\dagger(\boldsymbol{\theta}). \tag{B5}$$

Since $\text{Tr}[O^{(j)}(\boldsymbol{\theta})^2] = 1$, we can therefore use the Clifford shadows of the states $U|\psi^{(j)}\rangle$ to estimate $C_{\mathcal{D}(N)}(\boldsymbol{\theta})$.

Shallow-measurement algorithm. In the shallow-measurement setting, we instead gather M-shot Pauli shadows of the training states. That is, we measure each copy in a random Pauli basis by applying randomly chosen single-qubit Clifford gates on each qubit $W_m^{(j)} = W_{m,1}^{(j)} \otimes \cdots \otimes W_{m,n}^{(j)}$ and measuring in the computational basis. The M-shot Pauli shadow [38] of the state $U|\psi^{(j)}\rangle$ is now given by the collection $\{\rho_m^{(j)}\}_{m=1}^M$, for

$$\rho_m^{(j)} = \bigotimes_{i=1}^n \left(3W_{m,i}^\dagger |b_m^{(j)}\rangle\langle b_m^{(j)}| W_{m,i} - \mathbb{1}\right). \tag{B6}$$

From here, no more interaction with the target system is required. Pauli shadows are particularly suited to estimate local observables (i.e., that act nontrivially on a small number of qubits). To see how this makes them interesting for estimating the training loss $C_N^\ell(\boldsymbol{\theta})$ defined in Eq. (B2), we rewrite it as

$$C_N^\ell(\boldsymbol{\theta}) = 1 - \frac{1}{nN}\sum_{j=1}^{N}\sum_{i=1}^{n} \text{Tr}[U|\psi^{(j)}\rangle\langle\psi^{(j)}|U^\dagger\, O_i^{(j)}(\boldsymbol{\theta})], \tag{B7}$$

where we have absorbed the parametrized unitary $V(\boldsymbol{\theta})$ into the observable

$$O_i^{(j)}(\boldsymbol{\theta}) = V(\boldsymbol{\theta})\left(|\psi_i^{(j)}\rangle\langle\psi_i^{(j)}| \otimes \mathbb{1}_{\bar{i}}\right)V^\dagger(\boldsymbol{\theta}). \tag{B8}$$

The size of the Pauli shadow M required to provide precise estimates for these expectation values scales exponentially with the locality of the observables $O_i^{(j)}(\boldsymbol{\theta})$. Therefore, for parametrized unitaries that are low entangling, in the sense that they transform 1-local operators into $\mathcal{O}(\log(n))$-local operators, $O_i^{(j)}(\boldsymbol{\theta})$ is then a $\mathcal{O}(\log(n))$-local observable, and the number of copies M required should remain polynomial in n.

In the following, we analyze in more detail the sample complexity of both these algorithms.

2. Sample complexity analysis

We divide the analysis of the sample complexity of our incoherent learning protocols into two stages: (1) We analyze the sample complexity to incoherently estimate the training

loss $C(\boldsymbol{\theta})$, via the algorithms described above, for any sequence of parameter settings $\boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(L)}$, and (2) we show how this ability to estimate the training loss can be leveraged to learn an arbitrary target unitary U. Interestingly, we find in stage (1) that, for both the shallow- and deep-measurement settings, this sample complexity scales logarithmically with L, the number of parameter settings to be evaluated. This allows us in stage (2) to perform an exhaustive search for the parameter setting $\boldsymbol{\theta}_{\text{opt}}$ that approximately minimizes the Hilbert-Schmidt cost between U and $V(\boldsymbol{\theta})$. However, more generally, the result (1) can be seen as a guarantee that the sample complexity needed to simulate a coherent learning approach, e.g., based on gradient-based optimization of $\boldsymbol{\theta}$, depends at most logarithmically on the number of optimization steps.

Along with the theoretical guarantees of classical shadow estimation [38], the most important tool that we make use of in our sample complexity analysis is the notion of covering net of a parametrized quantum circuit, and an upper bound on its size [12]. These tools are presented formally in the next definitions and lemmas.

Definition B 1 (Covering net of a parametrized quantum circuit). Let $V(\boldsymbol{\theta})$ be a parametrized n-qubit quantum circuit. Let $\epsilon > 0$ and $D(U, V)$ a distance measure between unitaries. We call $\mathcal{N}_\epsilon$ an ϵ-covering net of the variational family $\{V(\boldsymbol{\theta})\}_{\boldsymbol{\theta}}$ an ensemble of unitaries $\{V(\boldsymbol{\theta}^{(i)})\}_{i=1}^{|\mathcal{N}_\epsilon|}$, such that for any $V(\boldsymbol{\theta})$ accessible to the parametrized quantum circuit,

$$\exists i \in \{1, \ldots, |\mathcal{N}_\epsilon|\} \text{ such that } D(V(\boldsymbol{\theta}^{(i)}), V(\boldsymbol{\theta})) \leqslant \epsilon. \quad \text{(B9)}$$

For our purposes, we only require a covering net for arbitrary two-qubit unitaries.

Lemma B 1 (Covering net for 2-qubit unitaries (Lemma C.1 in [12])). For any $\epsilon \in (0, 1]$, the set of two-qubit unitaries admits an ϵ-covering net $\mathcal{N}_\epsilon = \{U_i\}_{i=1}^{|\mathcal{N}_\epsilon|}$ with respect to the spectral distance $||U - V||_\infty$ with a size $|\mathcal{N}_\epsilon|$ that satisfies

$$|\mathcal{N}_\epsilon| \leqslant \left(\frac{6}{\epsilon}\right)^{32}. \quad \text{(B10)}$$

The sample complexities of Clifford and Pauli shadows are given by the following.

Lemma B 2 (Sample complexity of Clifford shadows (Theorem S1, Example 1 in [38])). Let $O_1, \ldots, O_M$ be a collection of M Hermitian operators and $\epsilon, \delta' \in [0, 1]$ be accuracy parameters. Clifford shadows using N' copies of a quantum state ρ allow to compute estimates $\hat{o}_1, \ldots, \hat{o}_M$ that satisfy

$$|\hat{o}_i - \text{Tr}[\rho O_i]| \leqslant \epsilon, \text{ for all } 1 \leqslant i \leqslant M \quad \text{(B11)}$$

with probability at least $1 - \delta'$ over the randomness of the measurements of ρ, assuming that

$$N' \geqslant \mathcal{O}\left(\frac{\max_i \text{Tr}\left[O_i^2\right] \log(M/\delta')}{\epsilon^2}\right). \quad \text{(B12)}$$

Lemma B 3 (Sample complexity of Pauli shadows (Theorem S1, Example 2 in [38])). Let $O_1, \ldots, O_M$ be a collection of M Hermitian operators such that $\|O_i\|_\infty \leqslant 1$, $\forall i$, and $\epsilon, \delta \in [0, 1]$ be accuracy parameters. Pauli shadows using N' copies of a quantum state ρ allow to compute estimates $\hat{o}_1, \ldots, \hat{o}_M$ that satisfy

$$|\hat{o}_i - \text{Tr}[\rho O_i]| \leqslant \epsilon, \text{ for all } 1 \leqslant i \leqslant M \quad \text{(B13)}$$

with probability at least $1 - \delta'$ over the randomness of the measurements of ρ, assuming that

$$N' \geqslant \mathcal{O}\left(\frac{\max_i 4^{\text{locality}(O_i)} \log(M/\delta')}{\epsilon^2}\right). \quad \text{(B14)}$$

a. Deep measurements

We start by analyzing the sample complexity of simulating the coherent training loss for L parameter settings as follows.

Lemma B 4 (Simulation of the coherent training loss using deep measurements). Let U be any n-qubit unitary and $V(\boldsymbol{\theta})$ a parametrized n-qubit unitary. Let $\epsilon, \delta > 0$ and $N, L \in \mathbb{N}^*$. The global training cost $C_{\mathcal{D}(N)}(\boldsymbol{\theta})$ defined in Eq. (B1) can be estimated using Clifford shadows for an arbitrary sequence of parameter settings $\boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(L)}$, up to additive error ϵ and with probability $1 - \delta$ over the randomness of the shadowing procedure, using

$$\widetilde{\mathcal{O}}\left(\frac{N \log(NL/\delta)}{\epsilon^2}\right) \quad \text{(B15)}$$

calls to the unitary U.

Proof. For each training state $|\psi^{(j)}\rangle$, $1 \leqslant j \leqslant N$, used to compute $C_{\mathcal{D}(N)}(\boldsymbol{\theta})$ call

$$O_{j,1}, \ldots, O_{j,L} = V(\boldsymbol{\theta}^{(1)})|\psi^{(j)}\rangle\langle\psi^{(j)}|V^\dagger(\boldsymbol{\theta}^{(1)}), \ldots, V(\boldsymbol{\theta}^{(L)})|\psi^{(j)}\rangle\langle\psi^{(j)}|V^\dagger(\boldsymbol{\theta}^{(L)}). \quad \text{(B16)}$$

By applying Lemma B 2 for $\rho = U|\psi^{(j)}\rangle\langle\psi^{(j)}|U^\dagger$, the $M = L$ observables above, and a probability of failure $\delta' = \delta/N$, we get that

$$N' \in \mathcal{O}\left(\frac{\log(LN/\delta)}{\epsilon^2}\right) \quad \text{(B17)}$$

calls to U are sufficient to guarantee that the term $\text{Tr}[U|\psi^{(j)}\rangle\langle\psi^{(j)}|U^\dagger\, O^{(j)}(\boldsymbol{\theta})]$ is estimated to additive error ϵ for all L parameter settings, with probability at least $1 - \delta/N$.

From the union bound, we have

$$P\left(\bigvee_{j=1}^{N} \text{“estimation } j \text{ fails”}\right) \leqslant \sum_{j=1}^{N} P(\text{“estimation } j \text{ fails”}) \leqslant \delta,$$

which implies that all estimations succeed with probability at least $1 - \delta$. Therefore,

$$C_{\mathcal{D}(N)}(\boldsymbol{\theta}) = 1 - \frac{1}{N}\sum_{i=1}^{N} \text{Tr}[U|\psi^{(j)}\rangle\langle\psi^{(j)}|U^\dagger \times V(\boldsymbol{\theta})|\psi^{(j)}\rangle\langle\psi^{(j)}|V^\dagger(\boldsymbol{\theta})] \quad \text{(B18)}$$

is computed successfully to additive error ϵ for all L parameter settings with probability $1 - \delta$, using the total claimed sample complexity. ■

We can now move to the proof of Theorem 1 in the main text, exposed formally in the following theorem.

Theorem 1 (Learning efficiently representable unitaries using deep measurements). Call an n-qubit unitary efficiently representable if it can be implemented using a quantum circuit of size poly(n). Let $\epsilon > 0$. For every polynomial circuit size, there exists a parametrized n-qubit unitary $V(\boldsymbol{\theta})$ that can be

trained to achieve a Hilbert-Schmidt cost $C_{\mathrm{HST}}(\boldsymbol{\theta}_{\mathrm{opt}}) \leqslant \epsilon$ for any efficiently representable U of that size, with high probability over the randomness of the shadowing procedure and the choice of training states, using

$$\widetilde{\mathcal{O}}\left(\frac{\mathrm{poly}(n)}{\epsilon^4}\right) \tag{B19}$$

calls to the unitary U.

Proof. We first show how it is possible to construct an ansatz that captures all efficiently representable n-qubit unitaries of a given size. Assume that we are interested in unitaries that can be implemented using a quantum circuit of at most n^c two-qubit gates with $c > 0$. Then, our ansatz is composed of n^c applications of the following gadget G: a universal two-qubit gate (see, e.g., Ref. [54]) applied on the first two qubits, sandwiched by $2n$ discretely parametrized gates SWAP^b, $b \in \{0, 1\}$, that can swap the first two qubits with any of the other qubits. This gadget allows one to apply an arbitrary two-qubit gate to any pair of qubits in the circuit, and therefore, the entire circuit captures all unitaries that can be implemented using at most n^c two-qubit gates.

We now construct the covering net for this ansatz with respect to the diamond distance $\|\mathcal{U} - \mathcal{V}\|_\diamond$ (between the unitary channels $\mathcal{U}, \mathcal{V}$ corresponding to the unitaries U and V). We start from the covering net $\widetilde{\mathcal{N}}_{\widetilde{\epsilon}}$ of an arbitrary two-qubit unitary with respect to the spectral norm $\|U - V\|_\infty$, for $\widetilde{\epsilon} = \frac{\epsilon}{2n^c}$. From Lemma B 1, we can take a covering net of size $|\widetilde{\mathcal{N}}_{\widetilde{\epsilon}}| \leqslant (12n^c/\epsilon)^{32}$. From here, we construct the set of unitaries corresponding to the gadget G above simply by counting all possible choices of pairs of qubits this unitary can be applied to. This set, corresponding to one application of the gadget G, has size $|\widetilde{\mathcal{N}}_{\widetilde{\epsilon}}|\binom{n}{2}$. Applying this gadget n^c times then results in a set of size $|\mathcal{N}_\epsilon| = (|\widetilde{\mathcal{N}}_{\widetilde{\epsilon}}|\binom{n}{2})^{n^c} \leqslant ((12n^c/\epsilon)^{32}n^2)^{n^c}$. It follows that

$$\log(|\mathcal{N}_\epsilon|) \leqslant \widetilde{\mathcal{O}}(n^c \log(n^c/\varepsilon)). \tag{B20}$$

We now proceed to show that this $|\mathcal{N}_\epsilon|$ sized set of unitaries corresponding to n^c applications of the gadget G is indeed an ϵ-covering net of the ansatz with respect to the diamond distance. To do so, we first note that the unitaries in the set take the form $\widetilde{V} = U_{n_c}\widetilde{V}_{n^c}U_{n^c-1}\cdots U_1\widetilde{V}_1U_0$, where U_i are composed of swap gates and $\widetilde{V}_i$ are two-qubit unitaries, and they cover elements of the form $V(\boldsymbol{\theta}) = U_{n_c}V_{n^c}(\boldsymbol{\theta})U_{n^c-1}\cdots U_1V_1(\boldsymbol{\theta})U_0$. Now, consider such a unitary $V(\boldsymbol{\theta})$ and its corresponding $\widetilde{V}$ in the covering net obtained by replacing $V_1(\boldsymbol{\theta}), \ldots, V_{n^c}(\boldsymbol{\theta})$ with the $\widetilde{\epsilon}$-close unitaries $\widetilde{V}_1, \ldots, \widetilde{V}_{n^c}$ in $\widetilde{\mathcal{N}}_{\widetilde{\epsilon}}$. We use the subadditivity of the diamond distance and the relation between diamond distance and spectral norm distance (see Eq. (C8) in Ref. [12]) to show

$$\begin{aligned}\left\|\widetilde{\mathcal{V}} - \mathcal{V}(\boldsymbol{\theta})\right\|_\diamond &\leqslant \sum_{i=1}^{n^c}\left\|\widetilde{\mathcal{V}}_i - \mathcal{V}_i(\boldsymbol{\theta})\right\|_\diamond + \sum_{i=0}^{n^c}\|\mathcal{U}_i - \mathcal{U}_i\|_\diamond \\ &\leqslant 2\sum_{i=1}^{n^c}\left\|\widetilde{V}_i - V_i(\boldsymbol{\theta})\right\|_\infty \leqslant \epsilon.\end{aligned} \tag{B21}$$

From here, we consider the parameter settings $\boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(L)}$ for $L = |\mathcal{N}_\epsilon|$ such that $V(\boldsymbol{\theta}^{(1)}), \ldots, V(\boldsymbol{\theta}^{(L)})$ correspond to the elements of the covering net $\mathcal{N}_\epsilon$. Since the target unitary U is guaranteed to correspond to some assignment $\boldsymbol{\theta}_U$, then by definition of the covering net, there exist $1 \leqslant i \leqslant L$ such that $\|\mathcal{V}(\boldsymbol{\theta}^{(i)}) - \mathcal{U}\|_\diamond \leqslant \epsilon$. We make use of the definition of the diamond norm and its relation to the fidelity between pure states for unitary channels [55] to show

$$\begin{aligned}&||\mathcal{V}(\boldsymbol{\theta}^{(i)}) - \mathcal{U}||_\diamond \\ &\quad:= \max_\psi ||V(\boldsymbol{\theta}^{(i)})|\psi\rangle\langle\psi|V^\dagger(\boldsymbol{\theta}^{(i)}) - U|\psi\rangle\langle\psi|U^\dagger||_1 \\ &\quad= \max_\psi 2\sqrt{1 - |\langle\psi|V^\dagger(\boldsymbol{\theta}^{(i)})U|\psi\rangle|^2} \\ &\quad\geqslant \max_\psi(1 - |\langle\psi|V^\dagger(\boldsymbol{\theta}^{(i)})U|\psi\rangle|^2) \\ &\quad\geqslant 1 - \frac{1}{N}\sum_{j=1}^{N}|\langle\psi^{(j)}|V^\dagger(\boldsymbol{\theta}^{(i)})U|\psi^{(j)}\rangle|^2 \\ &\quad:= C_{\mathcal{D}(N)}(\boldsymbol{\theta}^{(i)}).\end{aligned} \tag{B22}$$

And therefore, there exists a parameter setting $\boldsymbol{\theta}^{(i)}$ in the covering net that achieves a training error $C_{\mathcal{D}(N)}(\boldsymbol{\theta}^{(i)}) \leqslant \epsilon$. We also make use of the generalization bounds for this ansatz [12,13]. With high probability,

$$C_{\mathrm{HST}}(\boldsymbol{\theta}^{(i)}) \leqslant 2\frac{d+1}{d}\epsilon + \mathcal{O}\left(\sqrt{\frac{T\log(T)}{N}}\right) \tag{B23}$$

for $T = 2n^{c+1} + 15n^c$ its number of parameters (given that a universal two-qubit gate has 15 parameters [54], and that we use $2n$ parametrized swap gates per gadget G). This means that we can guarantee $C_{\mathrm{HST}}(\boldsymbol{\theta}^{(i)}) \leqslant \mathcal{O}(\epsilon)$ with high probability using

$$N \in \widetilde{\mathcal{O}}\left(\frac{n^{c+1}}{\epsilon^2}\right) \tag{B24}$$

training states. Then, by applying Lemma B 4 for $\boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(L)}$ ($L = |\mathcal{N}_\epsilon|$) defined above, and using the bounds in Eqs. (B20) and (B24), we get a total sample complexity in

$$\widetilde{\mathcal{O}}\left(\frac{n^{2c+1}}{\epsilon^4}\right). \tag{B25}$$

■

b. Shallow measurements

We also prove analogous results for the case of shallow measurements.

Lemma B 5 (Simulation of the coherent training loss using shallow measurements). Let U be any n-qubit unitary and $V(\boldsymbol{\theta})$ a parametrized n-qubit unitary that uses T trainable gates and is low entangling; i.e., it transforms 1-local operators into $\mathcal{O}(\log(n))$-local operators. Let $\epsilon, \delta > 0$ and $N, L \in \mathbb{N}^*$. The local training cost $C_N^\ell(\boldsymbol{\theta})$ defined in Eq. (B2) can be estimated using Pauli shadows for an arbitrary sequence of parameter settings $\boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(L)}$, up to additive error ϵ and with probability $1 - \delta$ over the randomness of the shadowing procedure, using

$$\widetilde{\mathcal{O}}\left(\frac{N\log(NLn/\delta)}{\epsilon^2}\right) \tag{B26}$$

calls to the unitary U.

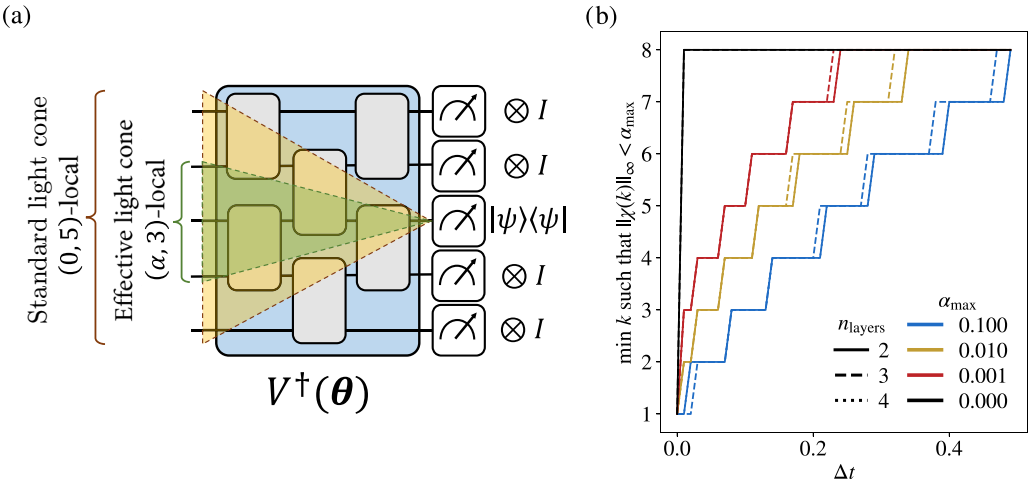


FIG. 5. Standard and approximate locality. (a) Cartoon representation of the difference between standard and approximate locality. (b) Here, we compare the growth of standard locality, i.e., $(0, k)$-locality, to approximate locality, i.e., (α, k)-locality. We focus on the locality of the operator $V(\Delta t)^\dagger(I_{\frac{n}{2}-1} \otimes |\psi\rangle\langle\psi| \otimes I_{\frac{n}{2}})V(\Delta t)$, where $V(\Delta t) = U(\Delta t/n_{\text{layers}})^{n_{\text{layers}}}$ and $U(\Delta t)$ is the first-order Trotter operator for the eight-qubit one-dimensional Heisenberg Hamiltonian with open boundary conditions, and $|\psi\rangle$ is a single-qubit Haar-random state. While standard locality only depends on the number of layers n_{layers} and not the evolution time Δt, the approximate locality is strongly dependent on the overall evolution time and not dependent on the total number of layers. The upper limit of k is the number of qubits in the system, which is 8 in our example.

Proof. The proof only differs from that of Lemma B 4 in the definition of the observables:

$$O_{i,j,l} = V(\boldsymbol{\theta})\big(|\psi_i^{(j)}\rangle\langle\psi_i^{(j)}| \otimes \mathbb{1}_{\bar{i}}\big)V^\dagger(\boldsymbol{\theta}),$$
$$1 \leqslant i \leqslant n, 1 \leqslant j \leqslant N, 1 \leqslant l \leqslant L, \tag{B27}$$

where we have now $M = nL$ observables per training state, and the application of Lemma B 3 instead of Lemma B 2.

The following theorem is a formal exposition of Theorem 2 in the main text. ■

Theorem 2 (Learning low-entangling unitaries using shallow measurements). Call an n-qubit unitary low entangling if it transforms 1-local operators into $\mathcal{O}(\log(n))$-local operators. Let $\epsilon > 0$. There exists a parametrized n-qubit unitary $V(\boldsymbol{\theta})$ that can be trained to achieve a Hilbert-Schmidt cost $C_{\text{HST}}(\boldsymbol{\theta}_{\text{opt}}) \leqslant \epsilon$ for any low-entangling U with high probability over the randomness of the shadowing procedure and the choice of training states, using

$$\widetilde{\mathcal{O}}\left(\frac{\text{poly}(n)}{\epsilon^4}\right) \tag{B28}$$

calls to the unitary U.

Proof. We apply the same reasoning as in Theorem 1 but where we keep only low-entangling unitaries in the covering net, which results, using Lemma B 5, in a similar sample complexity. Note that we additionally need to make use of $C_{\mathcal{D}(N)}(\boldsymbol{\theta}^{(i)}) \geqslant C_N^{\ell}(\boldsymbol{\theta}^{(i)})$ in the last line of Eq. (B22) and $C_{\text{HST}}(\boldsymbol{\theta}^{(i)}) \leqslant nC_{\text{LHST}}(\boldsymbol{\theta}^{(i)})$ (i.e., the local version of the Hilbert-Schmidt test) in Eq. (B23) in order to navigate back and forth between local and global cost functions at this stage of the proof. This only contributes with an $\widetilde{\mathcal{O}}(n^2)$ factor to the sample complexity. ■

c. Approximate locality

In this section, we argue that our proposed shallow-measurement incoherent learning algorithm works not only for unitaries that are low entangling in the sense of transforming 1-local operators into $\mathcal{O}(\log(n))$-local operators but also ones that transform 1-local operators into approximately $\mathcal{O}(\log(n))$-local operators (see Fig. 5). That is, we pin down the intuition, supported by our numerics in Fig. 4, that terms in the observables O_i that are highly nonlocal but have a small weight do not contribute substantially to the scaling of our algorithm.

To see this, we first formally define what we mean by approximately k-local.

Definition B 2 (Approximate locality- (α, k)-locality). Let $O_i = \tilde{O}_i(k) + \chi_i(k)$, where $\max_i$ locality $(O_i(k)) = k$ and $\min_i$ locality $(\chi_i(k)) = k + 1$. Then, we say an operator O_i is (α, k)-local if $||\chi_i(k)||_\infty = \alpha$.

Equipped with this definition, we can rewrite Theorem 2 as follows.

Theorem 3 (Learning low-entangling unitaries using shallow measurements: (α, k)-local version). Call an n-qubit unitary low-entangling n-qubit unitary if it transforms 1-local operators into (α, k)-local operators with $k \in \mathcal{O}(\log(n))$. Let $\epsilon > 0$. There exists a parametrized n-qubit unitary $V(\boldsymbol{\theta})$ that can be trained to achieve a Hilbert-Schmidt cost $C_{\text{HST}}(\boldsymbol{\theta}_{\text{opt}}) \leqslant \epsilon + \alpha$ for any low-entangling U with high probability over the randomness of the shadowing procedure and the choice of training states, using

$$\widetilde{\mathcal{O}}\left(\frac{\text{poly}(n)}{(\epsilon + \alpha)^4}\right) \tag{B29}$$

calls to the unitary U.

Proof. The core observation is that

$$\begin{aligned}|\hat{o}_i - \mathrm{Tr}[\rho O_i]| &= |\hat{o}_i - \mathrm{Tr}[\rho \tilde{O}_i(k)] - \mathrm{Tr}[\rho\, \chi_i(k)]| \\ &\leqslant |\hat{o}_i - \mathrm{Tr}[\rho \tilde{O}_i(k)]| + |\mathrm{Tr}[\rho\, \chi_i(k)]| \\ &\leqslant |\hat{o}_i - \mathrm{Tr}[\rho \tilde{O}_i(k)]| + \alpha, \end{aligned} \tag{B30}$$

where we use the triangle inequality and $|\mathrm{Tr}[\rho\, \chi_i(k)]| \leqslant ||\chi_i(k)||_\infty$. The remainder of the proof is identical to that of Theorem 2.

That indeed (α, k)-locality, and not standard locality (i.e., $\alpha = 0$ locality), determines the ultimate scaling of our algorithm (as claimed in Theorem 3) is seen in the numerics shown in Fig. 4. There, we study the performance of our learning algorithm as a function of shadow size for different evolution times of the target unitary (Δt) and different target unitary depths. If standard locality were what mattered, then we would expect the complexity of the algorithm to depend directly on the depth of the guess/target unitary. This is because the width of the light cone of the ansatz grows linearly with depth of the guess unitary, and so the locality of the parametrized observables [Eq. (B8)] similarly increases linearly with depth. However, in Fig. 4 the learning performance hardly changes as we increase from a one-layer target to a two-layer target. Instead, the performance depends on Δt. This makes sense because Δt determines how entangling the guess/target unitary is and hence the effective light cone of the guess/target unitary once small tails have been disregarded. In other words, Δt determines the (α, k)-locality of the parametrized observables and thus the complexity of the algorithm.

3. Remarks on the computational complexity

So far, we have only discussed the sample complexity of our incoherent learning protocols. In this section, we point out the caveats that appear when looking at their computational complexity.

The most obvious caveat of both protocols has to do with the exhaustive search used in both Theorems 1 and 2. While the scaling of the sample complexity is logarithmic in terms of the number L of parameter settings considered, the scaling remains linear in L for the computational complexity. Since L is exponential in n, this becomes quickly intractable.

However, we would like to consider here other computational aspects of our protocols, when used to train parametrized circuits in a standard fashion (e.g., using gradient-based optimization). As mentioned in the main text, the deep-measurement setting suffers from several pitfalls. While Lemma B 2 guarantees that the expectation values

$$\mathrm{Tr}[U|\psi^{(j)}\rangle\langle\psi^{(j)}|U^\dagger\, V(\boldsymbol{\theta})|\psi^{(j)}\rangle\langle\psi^{(j)}|V^\dagger(\boldsymbol{\theta})] \tag{B31}$$

can be estimated efficiently in terms of sample complexity, the estimation algorithm in Ref. [38] still needs to compute the observables $O^{(j)}(\boldsymbol{\theta}) = V(\boldsymbol{\theta})|\psi^{(j)}\rangle\langle\psi^{(j)}|V^\dagger(\boldsymbol{\theta})$ classically in order to compute their inner product with the classical shadows $\rho_m^{(j)}$ in Eq. (B3).[1] For a general $V(\boldsymbol{\theta})$, these observables can represent arbitrary rank-1 density matrices, and therefore are neither easy to represent nor to compute classically. One could imagine here making use of Clifford shadows again to prepare states $O^{(j)}(\boldsymbol{\theta})$ on a quantum computer, collect their shadows, and estimate their inner products with $\rho_m^{(j)}$. However, since $\rho_m^{(j)}$ is always a high-rank observable $[\mathrm{Tr}[(\rho_m^{(j)})^2] = 2^{n+1} - 1]$, then Clifford shadows have no guarantees of estimating this inner product efficiently. More generally, any protocol that would only use measurement outcomes of the density matrices $O^{(j)}(\boldsymbol{\theta})$ without knowledge of how these states were generated would be at risk of facing the $\Omega(\sqrt{2^n}/\epsilon)$ lower bound in sample complexity of incoherently estimating the overlap of two unknown input quantum states to ϵ precision (see Theorem 5 in Ref. [25]). Similarly, if one would attempt preparing the classical shadows $\rho_m^{(j)}$ on a quantum computer, evolve them under $V^\dagger(\boldsymbol{\theta})$, and measure the expectation value of $|\psi^{(j)}\rangle\langle\psi^{(j)}|$, i.e.,

$$\mathrm{Tr}[V^\dagger(\boldsymbol{\theta})\rho_m^{(j)}\, V(\boldsymbol{\theta})|\psi^{(j)}\rangle\langle\psi^{(j)}|], \tag{B32}$$

which is equivalent to the expression in Eq. (B31), one would face the problem that the classical shadows $\rho_m^{(j)}$ are completely unphysical states: while they have unit trace, they are not positive semidefinite (PSD), and have trace norm $||\rho_m^{(j)}||_1 = 2^{n+1} - 1$.

The recent work [56] further supports our arguments about the infeasibility of learning in the case of efficiently representable unitaries. Namely, Ref. [56] showed that under certain cryptographic assumptions, tomography for quantum circuits with superlogarithmically many two-qubit gates cannot be achieved computationally efficiently. In particular, general efficiently representable unitaries cannot be learned computationally efficiently. While phrased for a different notion of learning, these results are consistent with the pitfalls of the deep-measurement setting highlighted above.

As for our shallow-measurement protocol, aside from the exhaustive search used in Theorem 2, the simulation of the coherent training loss in Lemma B 5 is computationally efficient. Indeed, as one can visualize in Fig. 2, the observables $O_i^{(j)}(\boldsymbol{\theta}) = V(\boldsymbol{\theta})(|\psi_i^{(j)}\rangle\langle\psi_i^{(j)}| \otimes \mathbb{1}_{\bar{i}})V^\dagger(\boldsymbol{\theta})$ are $\mathcal{O}(\log(n))$-local, and the classical shadows in Eq. (B6) have a tensor-product form. Then, we can simply classically simulate the circuit $V(\boldsymbol{\theta})$ on an $\mathcal{O}(\log(n))$-qubit region [i.e., an $\mathcal{O}(\mathrm{poly}(n))$ Hilbert space], and efficiently compute the expectation values of $O_i^{(j)}(\boldsymbol{\theta})$. Moreover, since the entire training loss C_N^ℓ [Eq. (B2)] can therefore be computed classically, it can also be trained classically. We note as well that the training loss that we use in this protocol only contains local terms, making it less prone to barren plateaus [40].

[1] Note that the operators $O^{(j)}(\boldsymbol{\theta}) = V(\boldsymbol{\theta})|\psi^{(j)}\rangle\langle\psi^{(j)}|V^\dagger(\boldsymbol{\theta})$ can be viewed as observables but also as quantum states. In this paragraph, we make use of both of these views, whenever they are more convenient.

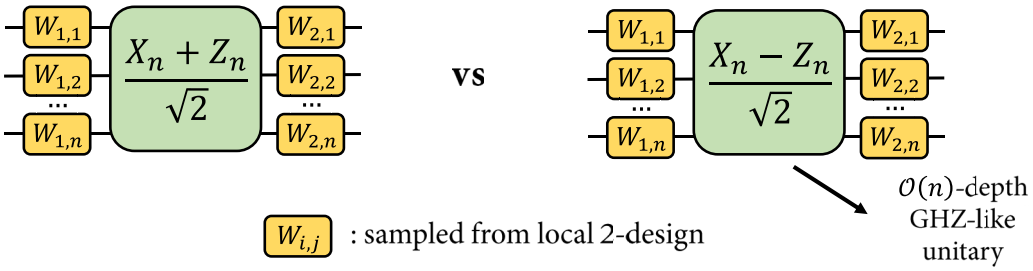


FIG. 6. Learning task used to show the limitations of shallow measurements. We show the hardness of learning polynomial-depth unitaries using shallow measurements by reduction of the problem of distinguishing between the unitaries $U_+ = (X^{\otimes n} + Z^{\otimes n})/\sqrt{2}$ and $U_- = (X^{\otimes n} - Z^{\otimes n})/\sqrt{2}$ (for n odd). We first consider the warm-up case where these target unitaries are taken as is (Lemma C 1), followed by the harder case where they are sandwiched with products of unitaries sampled from a local 2-design (Theorem C 1). We call U_+ and U_- "GHZ-like" unitaries because they prepare orthogonal GHZ states $(|0\rangle^{\otimes n} \pm |1\rangle^{\otimes n}))/\sqrt{2}$ when they are applied on a $|0\rangle^{\otimes n}$ state.

APPENDIX C: SHALLOW MEASUREMENTS FAIL IN LEARNING LINEAR-DEPTH CIRCUITS

1. Warm-up example

We start by studying the simpler case where the shallow measurements considered are limited to be random Pauli measurements, as those used in Pauli shadows.

Lemma C 1. Consider the product POVM $F = \frac{1}{3^n}\{|0\rangle\langle 0|, |1\rangle\langle 1|, |+\rangle\langle +|, |-\rangle\langle -|, |i\rangle\langle i|, |-i\rangle\langle -i|\}^{\otimes n}$, which corresponds to measuring in the eigenbasis of a random n-qubit Pauli observable in $\{X, Y, Z\}^{\otimes n}$, and a batch $\mathcal{D}$ of T quantum states sampled uniformly from $\{|0\rangle, |1\rangle, |+\rangle, |-\rangle, |i\rangle, |-i\rangle\}^{\otimes n}$. Suppose there exists an algorithm that, for any unitary U that can be implemented in $\mathcal{O}(n)$ depth, returns (the description of) a unitary V such that $1 - \frac{1}{4^n}|\mathrm{Tr}[UV^\dagger]|^2 \leqslant 1/4$ with high probability, using, for each quantum state $|\psi_{\mathrm{in}}\rangle$ in $\mathcal{D}$, the measurement outcomes of the POVM F on M copies of $U|\psi_{\mathrm{in}}\rangle$. Then, necessarily, the total number of measurements that any such algorithm must use is

$$MT \geqslant 2^{\Omega(n)}. \tag{C1}$$

Proof. We prove this result by reduction of the task of distinguishing between the unitaries (see Figs. 6 and 7):

$$\begin{cases} \mathrm{U}_+ = (X^{\otimes n} + Z^{\otimes n})/\sqrt{2} \quad \text{and} \quad U_- = (X^{\otimes n} - Z^{\otimes n})/\sqrt{2}, \\ \mathrm{U}_+ = (X^{\otimes n} + iZ^{\otimes n})/\sqrt{2} \quad \text{and} \quad U_- = (X^{\otimes n} - iZ^{\otimes n})/\sqrt{2}. \end{cases} \tag{C2}$$

for n odd and even, respectively, which we show requires $MT \geqslant 2^{\Omega(n)}$ measurements.

a. Reduction

Assume that an algorithm can achieve $\frac{1}{4^n}|\mathrm{Tr}[U_+V^\dagger]|^2 \geqslant 3/4$, when U_+ is the target unitary. We first note that

$$\frac{1}{4^n}|\mathrm{Tr}[U_+U_-^\dagger]|^2 = 0. \tag{C3}$$

Therefore, we can decompose V in a $\{U_+, U_-, \ldots\}$ basis as

$$V = \frac{1}{2^n}\mathrm{Tr}[VU_+^\dagger]U_+ + \frac{1}{2^n}\mathrm{Tr}[VU_-^\dagger]U_- + \cdots \tag{C4}$$

and deduce (from Parseval's identity) that

$$1 = \frac{1}{4^n}|\mathrm{Tr}[VV^\dagger]|^2 \tag{C5}$$

$$= \left(\frac{1}{4^n}|\mathrm{Tr}[U_+V^\dagger]|^2\right)^2 + \left(\frac{1}{4^n}|\mathrm{Tr}[U_-V^\dagger]|^2\right)^2 + c, \tag{C6}$$

where $c \geqslant 0$. From here, we have

$$\frac{1}{4^n}|\mathrm{Tr}[U_-V^\dagger]|^2 \leqslant \sqrt{1 - \left(\frac{1}{4^n}|\mathrm{Tr}[U_+V^\dagger]|^2\right)^2} \tag{C7}$$

$$\leqslant \sqrt{1 - \left(\frac{3}{4}\right)^2} < 0.67. \tag{C8}$$

We find $\frac{1}{4^n}|\mathrm{Tr}[U_-V^\dagger]|^2 < \frac{1}{4^n}|\mathrm{Tr}[U_+V^\dagger]|^2$ when U_+ is the target unitary, and through a similar argument $\frac{1}{4^n}|\mathrm{Tr}[U_+V^\dagger]|^2 < \frac{1}{4^n}|\mathrm{Tr}[U_-V^\dagger]|^2$ when U_- is the target unitary. Therefore, this algorithm can be used as a subroutine to distinguish between U_+ and U_-, simply by learning V, computing $\frac{1}{4^n}|\mathrm{Tr}[U_+V^\dagger]|^2$ and $\frac{1}{4^n}|\mathrm{Tr}[U_-V^\dagger]|^2$, and outputting the unitary with the largest overlap.

b. Hardness of distinguishing U_+ and U_-

For $U_\pm \in \{U_+, U_-\}$, consider the joint probability distribution

$$p_\pm(\psi_1, \ldots, \psi_T, \phi_1^1, \ldots, \phi_T^M) = p_\pm(\psi_{1:T}, \phi_{1:T}^{1:M}) \tag{C9}$$

of sampling the input states

$$|\psi_1\rangle, \ldots, |\psi_T\rangle \in (\{|0\rangle, |1\rangle, |+\rangle, |-\rangle, |i\rangle, |-i\rangle\}^{\otimes n})^T \tag{C10}$$

when generating the batch $\mathcal{D}$, and of measuring the projections

$$\left|\phi_1^1\right\rangle\!\left\langle\phi_1^1\right|, \ldots, \left|\phi_T^M\right\rangle\!\left\langle\phi_T^M\right| \in (\{|0\rangle\langle 0|, |1\rangle\langle 1|, |+\rangle\langle +|, |-\rangle\langle -|, |i\rangle\langle i|, |-i\rangle\langle -i|\}^{\otimes n})^{MT}, \tag{C11}$$

respectively, on the M-copy blocks of these states, after being all evolved under $U_\pm$.

LeCam's method (see, e.g., Lemma 1 in Ref. [57]) tells us that the probability of success of any algorithm that attempts to distinguish between U_+ and U_- with the type of access

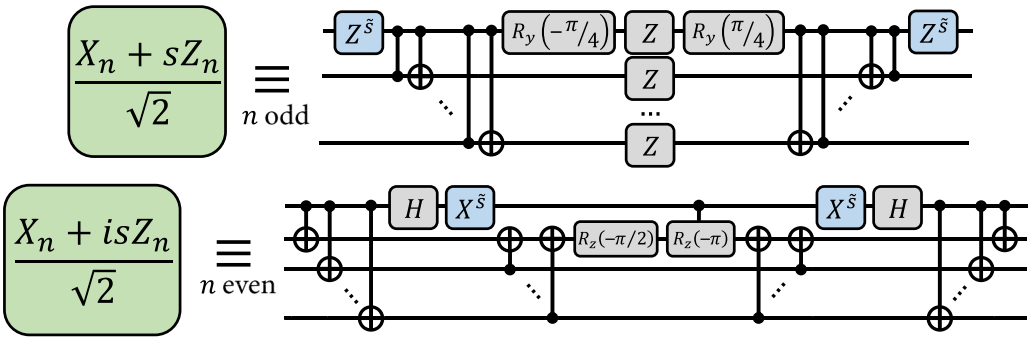


FIG. 7. Circuit decomposition of the GHZ-like unitaries. We give an $\mathcal{O}(n)$-depth circuit decomposition of the unitaries U_+ and U_- defined in Eq. (C2). We use the notation $s = \pm 1$ and $\tilde{s} = \frac{s+1}{2} = 0/1$. The resulting unitaries are equivalent up to a global phase s.

defined in the statement of this lemma is upper bounded by

$$\mathrm{P}[\text{success}] \leqslant \tfrac{1}{2} + \tfrac{1}{2}\mathrm{TV}(p_+, p_-), \tag{C12}$$

where

$$\mathrm{TV}(p_+, p_-) = \frac{1}{2}\sum_{\psi_{1:T}, \phi_{1:T}^{1:M}} \left|p_+\left(\psi_{1:T}, \phi_{1:T}^{1:M}\right) - p_-\left(\psi_{1:T}, \phi_{1:T}^{1:M}\right)\right|. \tag{C13}$$

Therefore, to show that $MT \geqslant 2^{\Omega(n)}$ measurements are necessary to distinguish the two unitaries with high probability, our goal is to show that this total variation distance between the two joint probability distributions is of the order $MT/2^{\mathcal{O}(n)}$.

We establish this result using a very handy property from probability theory. Namely, if we take p_i, q_i to be probability distributions on a finite space $\Omega_i, i = 1, \ldots, N$, then, for the product distributions $p = \prod_{i=1}^N p_i$ and $q = \prod_{i=1}^N q_i$ on the finite space $\Omega = \prod_{i=1}^N \Omega_i$, we have [58]

$$\mathrm{TV}(p, q) \leqslant \sum_{i=1}^N \mathrm{TV}(p_i, q_i). \tag{C14}$$

From here, we note that all of the T M-copy blocks of measurements are independent and identically distributed (i.i.d.), and that inside each of these blocks, the M measurements are also i.i.d., conditioned on the choice of $|\psi_i\rangle$. Therefore, we can write

$$\begin{aligned}
&\mathrm{TV}(p_+, p_-) \\
&\quad= \frac{1}{2}\sum_{\psi_{1:T}, \phi_{1:T}^{1:M}} \left|p_+\left(\psi_{1:T}, \phi_{1:T}^{1:M}\right) - p_-\left(\psi_{1:T}, \phi_{1:T}^{1:M}\right)\right| \\
&\quad\leqslant \frac{T}{2}\sum_{\psi_1, \phi_1^{1:M}} \left|p_+\left(\psi_1, \phi_1^{1:M}\right) - p_-\left(\psi_1, \phi_1^{1:M}\right)\right| \\
&\quad= \frac{T}{2}\sum_{\psi_1, \phi_1^{1:M}} p(\psi_1)\left|p_+\left(\phi_1^{1:M}|\psi_1\right) - p_-\left(\phi_1^{1:M}|\psi_1\right)\right| \\
&\quad\leqslant \frac{MT}{2}\sum_{\psi_1, \phi_1} p(\psi_1)\left|p_+\left(\phi_1^{1}|\psi_1\right) - p_-\left(\phi_1^{1}|\psi_1\right)\right|,
\end{aligned} \tag{C15}$$

which means we only need to focus on the total variation distance of a single measurement. Here, we made use of Eq. (C14) in the second and fourth lines, and of the decomposition $p_\pm(\psi_1, \phi_1^{1:M}) = p(\psi_1)p_\pm(\phi_1^{1:M}|\psi_1)$ in the third line.

For n odd, the total variation distance of a single measurement is given by

$$\begin{aligned}
\mathrm{TV}(p_+^1, p_-^1) &= \frac{1}{4}\sum_{\psi,\phi}\frac{p(\psi)}{3^n}\left|\mathrm{Tr}[(X_n+Z_n)|\psi\rangle\langle\psi|(X_n+Z_n)|\phi\rangle\langle\phi|] - \mathrm{Tr}[(X_n-Z_n)|\psi\rangle\langle\psi|(X_n-Z_n)|\phi\rangle\langle\phi|]\right| \\
&= \frac{1}{2}\sum_{\psi,\phi}\frac{p(\psi)}{3^n}\left|\mathrm{Tr}[X_n|\psi\rangle\langle\psi|Z_n|\phi\rangle\langle\phi|] + \mathrm{Tr}[Z_n|\psi\rangle\langle\psi|X_n|\phi\rangle\langle\phi|]\right| \\
&= \sum_{\psi,\phi}\frac{p(\psi)}{3^n}\left|\mathrm{Re}[\langle\phi|X_n|\psi\rangle\langle\psi|Z_n|\phi\rangle]\right|,
\end{aligned} \tag{C16}$$

where we have abbreviated $X^{\otimes n}$ and $Z^{\otimes n}$ by X_n and Z_n, and have dropped the sub- and superscripts of $|\psi_1\rangle$ and $|\phi_1^1\rangle$. Similarly, for n even, we have

$$\mathrm{TV}(p_+^1, p_-^1) = \sum_{\psi,\phi}\frac{p(\psi)}{3^n}\left|\mathrm{Im}[\langle\phi|X_n|\psi\rangle\langle\psi|Z_n|\phi\rangle]\right|. \tag{C17}$$

For all n, we then have

$$\mathrm{TV}(p_+^1, p_-^1) \leqslant \sum_{\psi,\phi}\frac{p(\psi)}{3^n}\left|\langle\phi|X_n|\psi\rangle\langle\psi|Z_n|\phi\rangle\right|. \tag{C18}$$

To show that this quantity is exponentially small in n, we exploit the tensor product structure of $|\psi\rangle$ and $|\phi\rangle$ as well as a property of the stabilizer states $\mathcal{S} = \{|0\rangle, |1\rangle, |+\rangle, |-\rangle, |i\rangle, |-i\rangle\}$ that is easy to check (numerically):

$$\begin{aligned}
\mathbb{E}_{|x\rangle,|y\rangle\in\mathcal{S}}\left[\left|\langle x|X|y\rangle\langle y|Z|x\rangle\right|\right] &= \frac{1}{36}\sum_{|x\rangle,|y\rangle\in\mathcal{S}}\left|\langle x|X|y\rangle\langle y|Z|x\rangle\right| \\
&= \frac{7}{18}.
\end{aligned} \tag{C19}$$

This gives us

$$\begin{aligned}\mathrm{TV}(p_+^1, p_-^1) &\leqslant \sum_{\psi,\phi} \frac{p(\psi)}{3^n} |\langle\phi|X_n|\psi\rangle\langle\psi|Z_n|\phi\rangle| \\ &= \sum_{\psi,\phi} \frac{1}{3^n} \prod_{i=1}^{n} p(\psi^{(i)}) |\langle\phi^{(i)}|X|\psi^{(i)}\rangle\langle\psi^{(i)}|Z|\phi^{(i)}\rangle| \\ &= \sum_{\psi,\phi} \frac{1}{18^n} \prod_{i=1}^{n} |\langle\phi^{(i)}|X|\psi^{(i)}\rangle\langle\psi^{(i)}|Z|\phi^{(i)}\rangle| \\ &= \frac{1}{18^n} \prod_{i=1}^{n} \sum_{\psi^{(i)},\phi^{(i)}\in\mathcal{S}} |\langle\phi^{(i)}|X|\psi^{(i)}\rangle\langle\psi^{(i)}|Z|\phi^{(i)}\rangle| \\ &= \frac{36^n}{18^n}\left(\frac{7}{18}\right)^n = \left(\frac{14}{18}\right)^n. \end{aligned} \tag{C20}$$

Combining this last inequality with Eq. (C15), we get

$$\mathrm{TV}(p_+, p_-) \leqslant MT\left(\frac{14}{18}\right)^n < MT/2^{0.36n}, \tag{C21}$$

and therefore MT must be scaling as $2^{\Omega(n)}$ for the algorithm to have high probability of success. ■

Note, however, that when we step away from randomly chosen input states and measurement bases, distinguishing efficiently between U_+ and U_- using product input states and product measurements becomes possible. In the case where n is odd, for instance, it is sufficient to consider the input state $|\psi\rangle = |+\rangle^{\otimes n}$, and to measure $U|\psi\rangle$ in the computational basis. One can then check that the states $U_+|\psi\rangle$ and $U_-|\psi\rangle$ have completely disjoint supports, which allows to distinguish between the two unitaries using only one measurement. A similar observation can be made for the case when n is even, by taking, for instance, $|\psi\rangle = |+\rangle^{\otimes n-1} \otimes |i\rangle$. This justifies why we look into a more general case (and a harder learning task) in the next section.

Comparison with Hamiltonian learning. It is perhaps worth highlighting at this point that this "warm up" no-go theorem on the difficulty of learning efficiently representable processes with Pauli measurements in the incoherent framework highlights a key difference between the process learning and Hamiltonian learning frameworks. Namely, access to the continual time dynamics of a system is substantially more powerful than access to the dynamics at a single time. To see this, we note that Result 2 in Ref. [47] shows that it *is possible* to incoherently learn arbitrary sparse Hamiltonians with only local stabilizer inputs and Pauli measurements. However, we have just shown that it *is not possible* to learn a unitary of the form $U = X^{\otimes n} + Z^{\otimes n}$ (for n odd), which corresponds to evolution under the Hamiltonian $H = \frac{\pi}{2t}\frac{X^{\otimes n}+Z^{\otimes n}}{\sqrt{2}}$ for time t, using local stabilizer inputs and Pauli measurements. This apparent contradiction is reconciled via the observation that the learning model in Ref. [47] is more powerful in virtue of allowing access to the evolution of the system at multiple times, whereas we only allow access to the single unitary U corresponding to the evolution under H at a single time.

2. General case

We now move to the general setting of arbitrary product measurements, which are moreover allowed to be adaptively chosen (i.e., the next input state and measurement basis can be chosen as a function of previous measurement outcomes).

Theorem C 1. Call $F_1, \ldots, F_T$ an arbitrary sequence of n-qubit product POVMs, i.e., $F_i = \{w_{i,j}2^n|\phi_{i,j}\rangle\langle\phi_{i,j}|\}_j$, where $|\phi_{i,j}\rangle = |\phi_{i,j}^{(1)}\rangle \otimes \cdots \otimes |\phi_{i,j}^{(n)}\rangle$ and $\sum_j w_{i,j} = 1$, and $|\psi_1\rangle, \ldots, |\psi_T\rangle$ an arbitrary sequence of product n-qubit quantum states, i.e., $|\psi_i\rangle = |\psi_i^{(1)}\rangle \otimes \cdots \otimes |\psi_i^{(n)}\rangle$, that can both be chosen adaptively (i.e., a choice of quantum state $|\psi_i\rangle$ and POVM F_i can depend on the previous measurement outcomes resulting from the choices $|\psi_1\rangle, \ldots, |\psi_{i-1}\rangle$ and $F_1, \ldots, F_{i-1}$). Suppose there exists an algorithm that, for any unitary U that can be implemented in $\mathcal{O}(n)$ depth, returns (the description of) a unitary V such that $1 - \frac{1}{4^n}|\mathrm{Tr}[UV^\dagger]|^2 \leqslant 1/4$ with high probability, using the measurement outcomes of $F_1, \ldots, F_T$ on $U|\psi_1\rangle, \ldots, U|\psi_T\rangle$, respectively. Then, necessarily, the total number of measurements any such algorithm must use is

$$T \geqslant 2^{\Omega(n)}. \tag{C22}$$

Proof. Consider the "GHZ-like" unitaries (Fig. 6 and 7):

$$\begin{cases} U_+ = (X^{\otimes n}+Z^{\otimes n})/\sqrt{2} \quad \text{and} \quad U_-=(X^{\otimes n} - Z^{\otimes n})/\sqrt{2}, \\ U_+ = (X^{\otimes n}+iZ^{\otimes n})/\sqrt{2} \quad \text{and} \quad U_-=(X^{\otimes n} - iZ^{\otimes n})/\sqrt{2} \end{cases} \tag{C23}$$

for n odd and even, respectively.

We prove the result by reduction of the task of distinguishing between the unitaries

$$W_2U_+W_1 \quad \text{and} \quad W_2U_-W_1 \tag{C24}$$

for U_+ and U_- given by Eq. (C23), and

$$W_1 = \bigotimes_{i=1}^{n} W_1^{(i)}, \qquad W_2 = \bigotimes_{i=1}^{n} W_2^{(i)} \tag{C25}$$

for some choice of single-qubit unitaries $W_{1/2}^{(1)}, \ldots, W_{1/2}^{(n)}$ taken from the single-qubit Clifford group $\mathcal{C}_1$, i.e., $(W_1, W_2) \in \mathcal{W} = \mathcal{C}_1^{\otimes n} \times \mathcal{C}_1^{\otimes n}$, which are revealed to the learner *after* the T measurements.

a. Reduction

The reduction is similar to that of our previous lemma, by noting that, for any unitaries W_1, W_2, we still have

$$\tfrac{1}{4^n}|\mathrm{Tr}[W_2U_+W_1W_1^\dagger U_-^\dagger W_2^\dagger]|^2 = \tfrac{1}{4^n}|\mathrm{Tr}[U_+U_-^\dagger]|^2 = 0, \tag{C26}$$

and the rest of the derivation in distinguishing $W_2U_+W_1$ and $W_2U_-W_1$ simply follows. Since this holds for any W_1, W_2, it means that the algorithm can also distinguish *any* $W_2U_+W_1$ and $W_2U_-W_1$ for any choice of W_1, W_2 in $\mathcal{W}$ when defining the target unitaries. However, the distinguishing task we consider here must have an additional layer of complexity. Indeed, when W_1 and W_2 are given to the learner *before* learning V, they can then be taken into account during the learning procedure, which makes the task equivalent to distinguishing between U_+ and U_- directly. We therefore consider a distinguishing task where W_1 and W_2 are revealed to the learner

after learning V. Then, the learner can use the learned V to compute its overlap with $W_2U_+W_1$ and $W_2U_-W_1$ and distinguish between the two unitaries, but must learn V without the knowledge of W_1, W_2.

b. Hardness of distinguishing $W_2U_+W_1$ and $W_2U_-W_1$

For some choice of $(W_1, W_2) \in \mathcal{W}$ to be defined later, we want to bound the total variation distance between p_+ and p_-, where

$$p_\pm(\psi_{1:T}, \phi_{1:T}) \tag{C27}$$

is the joint probability distribution of measuring the projections

$$|\phi_1\rangle\langle\phi_1|, \ldots, |\phi_T\rangle\langle\phi_T| \in \{|\phi_{1,j}\rangle\langle\phi_{1,j}|\}_j \times \cdots \times \{|\phi_{T,j}\rangle\langle\phi_{T,j}|\}_j \tag{C28}$$

on the states

$$W_2U_\pm W_1|\psi_1\rangle, \ldots, W_2U_\pm W_1|\psi_T\rangle, \tag{C29}$$

respectively, where the states $|\psi_i\rangle$ and the POVMs $F_i = \{w_{i,j}2^n|\phi_{i,j}\rangle\langle\phi_{i,j}|\}_j$ have been chosen adaptively[2].

We start by considering the total variation distance of individual measurements (where $t = 1, \ldots, T$ indexes the measurement):

$$\mathrm{TV}(p_+^t, p_-^t) = \frac{1}{2}\sum_j |p_+(\psi_t, \phi_{t,j}) - p_-(\psi_t, \phi_{t,j})| \tag{C30}$$

and we want to show that for an overwhelming fraction of $(W_1, W_2) \in \mathcal{W}$, this total variation distance is exponentially small in n.

From a similar derivation to that of our previous lemma, we get

$$\mathrm{TV}(p_+^t, p_-^t) = \begin{cases} \sum_j w_{t,j}2^n \left|\mathrm{Re}[\langle\phi_{t,j}|W_2^\dagger X_n W_1|\psi_t\rangle\langle\psi_t|W_1^\dagger Z_n W_2|\phi_{t,j}\rangle]\right| \\ \sum_j w_{t,j}2^n \left|\mathrm{Im}[\langle\phi_{t,j}|W_2^\dagger X_n W_1|\psi_t\rangle\langle\psi_t|W_1^\dagger Z_n W_2|\phi_{t,j}\rangle]\right| \end{cases}. \tag{C31}$$

for n odd and even, respectively.

We now average this total variation distance over the random choice of W_1, W_2 (sampled uniformly from $\mathcal{W}$), in the case where n is odd:

$$\mathbb{E}_{W_1,W_2}[\mathrm{TV}(p_+^t, p_-^t)] = \sum_j w_{t,j}2^n \mathbb{E}_{W_1,W_2}[|\mathrm{Re}[\langle\phi_{t,j}|W_2^\dagger X_n W_1|\psi_t\rangle\langle\psi_t|W_1^\dagger Z_n W_2|\phi_{t,j}\rangle]|] \tag{C32}$$

and want to show that the expected values in this sum are all exponentially small in n.

The derivation is independent of t, j, so we call $X(W_1, W_2) = \mathrm{Re}[\langle\phi_{t,j}|W_2^\dagger X_n W_1|\psi_t\rangle\langle\psi_t|W_1^\dagger Z_n W_2|\phi_{t,j}\rangle]$ for an arbitrary t, j. From Jensen's inequality $g(\mathbb{E}[Y]) \leqslant \mathbb{E}[g(Y)]$ applied to $g : x \mapsto x^2$ and $Y = |X|$, we get

$$\mathbb{E}_{W_1,W_2}[|X(W_1, W_2)|] \leqslant \sqrt{\mathbb{E}_{W_1,W_2}[X(W_1, W_2)^2]} = \sqrt{\mathbb{E}_{W_1,W_2}[\mathrm{Re}[\langle\phi_{t,j}|W_2^\dagger X_n W_1|\psi_t\rangle\langle\psi_t|W_1^\dagger Z_n W_2|\phi_{t,j}\rangle]^2]}. \tag{C33}$$

From here, we make use of $\mathrm{Re}[z]^2 = (\frac{z+\bar{z}}{2})^2 = \frac{1}{4}(z^2 + \overline{z^2} + 2z\bar{z})$ for $z = \langle\phi_{t,j}|W_2^\dagger X_n W_1|\psi_t\rangle\langle\psi_t|W_1^\dagger Z_n W_2|\phi_{t,j}\rangle$ to decompose this expected value into

$$\tfrac{1}{4}(\mathbb{E}_{W_1,W_2}[z^2] + \mathbb{E}_{W_1,W_2}[\overline{z^2}] + 2\mathbb{E}_{W_1,W_2}[z\bar{z}]). \tag{C34}$$

The tensor product structure of $|\psi_t\rangle, |\phi_{t,j}\rangle, W_1, W_2$ and the following relation for a single-qubit unitary W sampled from a 2-design and an arbitrary single-qubit state $|x\rangle$ (Eq. (2.26) in Ref. [59]):

$$\mathbb{E}_W[(W|x\rangle\langle x|W^\dagger)^{\otimes 2}] = \frac{I + \mathrm{SWAP}}{6} \tag{C35}$$

are then useful to compute

$$\begin{aligned}
\mathbb{E}_{W_1,W_2}[z^2] &= \mathbb{E}_{W_1,W_2}[\mathrm{Tr}[W_2|\phi_{t,j}\rangle\langle\phi_{t,j}|W_2^\dagger X_n W_1|\psi_t\rangle\langle\psi_t|W_1^\dagger Z_n]^2] \\
&= \prod_{i=1}^n \mathbb{E}_{W_1^{(i)},W_2^{(i)}}[\mathrm{Tr}[W_2^{(i)}|\phi_{t,j}^{(i)}\rangle\langle\phi_{t,j}^{(i)}|W_2^{(i)\dagger} X W_1^{(i)}|\psi_t^{(i)}\rangle\langle\psi_t^{(i)}|W_1^{(i)\dagger} Z]^2] \\
&= \prod_{i=1}^n \mathbb{E}_{W_1^{(i)},W_2^{(i)}}[\mathrm{Tr}[\left(W_2^{(i)}|\phi_{t,j}^{(i)}\rangle\langle\phi_{t,j}^{(i)}|W_2^{(i)\dagger} X W_1^{(i)}|\psi_t^{(i)}\rangle\langle\psi_t^{(i)}|W_1^{(i)\dagger} Z\right)^{\otimes 2}]] \\
&= \prod_{i=1}^n \mathrm{Tr}[\frac{I + \mathrm{SWAP}}{6} X^{\otimes 2} \frac{I + \mathrm{SWAP}}{6} Z^{\otimes 2}]
\end{aligned}$$

[2]For conciseness, we do not index each ψ_i, F_i, ϕ_i by the previous choices and measurement outcomes, although these in general depend on $\psi_1, F_1, \phi_1, \ldots, \psi_{i-1}, F_{i-1}, \phi_{i-1}$.

$$= \prod_{i=1}^{n} \frac{1}{36}(\mathrm{Tr}[X^{\otimes 2}Z^{\otimes 2}] + \mathrm{Tr}[\mathrm{SWAP}X^{\otimes 2}Z^{\otimes 2}] + \mathrm{Tr}[X^{\otimes 2}\mathrm{SWAP}Z^{\otimes 2}] + \mathrm{Tr}[\mathrm{SWAP}X^{\otimes 2}\mathrm{SWAP}Z^{\otimes 2}])$$

$$= \prod_{i=1}^{n} \frac{1}{36}(2\mathrm{Tr}[X^{\otimes 2}Z^{\otimes 2}] + 2\mathrm{Tr}[\mathrm{SWAP}X^{\otimes 2}Z^{\otimes 2}])$$

$$= \prod_{i=1}^{n} \frac{1}{36}(0-4) = \left(-\frac{1}{9}\right)^{n}, \tag{C36}$$

which is also the value of $\mathbb{E}_{W_1,W_2}[\overline{z^2}]$ since it is real. Here, we also made use of $(A \otimes B)\mathrm{SWAP} = \mathrm{SWAP}(B \otimes A)$ [since $\mathrm{SWAP}(A \otimes B)\mathrm{SWAP} = (B \otimes A)$] and $\mathrm{Tr}[\mathrm{SWAP}A \otimes B] = \mathrm{Tr}[AB]$.

Similarly, by making use of $\overline{\mathrm{Tr}[A]} = \mathrm{Tr}[\,\overline{A}\,] = \mathrm{Tr}[A^{\dagger}]$, we have

$$\mathbb{E}_{W_1,W_2}[z\bar{z}] = \mathbb{E}_{W_1,W_2}[\mathrm{Tr}[W_2|\phi_{t,j}\rangle\langle\phi_{t,j}|W_2^{\dagger}X_nW_1|\psi_t\rangle\langle\psi_t|W_1^{\dagger}Z_n]\mathrm{Tr}[W_2|\phi_{t,j}\rangle\langle\phi_{t,j}|W_2^{\dagger}Z_nW_1|\psi_t\rangle\langle\psi_t|W_1^{\dagger}X_n]]$$

$$= \prod_{i=1}^{n} \mathbb{E}_{W_1^{(i)},W_2^{(i)}}\left[\mathrm{Tr}\left[\left(W_2^{(i)}|\phi_{t,j}^{(i)}\rangle\langle\phi_{t,j}^{(i)}|W_2^{(i)\dagger}\right)^{\otimes 2}(X\otimes Z)\left(W_1^{(i)}|\psi_t^{(i)}\rangle\langle\psi_t^{(i)}|W_1^{(i)\dagger}\right)^{\otimes 2}(Z\otimes X)\right]\right]$$

$$= \prod_{i=1}^{n} \mathrm{Tr}\left[\frac{I+\mathrm{SWAP}}{6}(X\otimes Z)\frac{I+\mathrm{SWAP}}{6}(Z\otimes X)\right]$$

$$= \prod_{i=1}^{n} \frac{1}{36}(\mathrm{Tr}[XZ\otimes ZX] + \mathrm{Tr}[\mathrm{SWAP}(XZ\otimes ZX)] + \mathrm{Tr}[(X\otimes Z)\mathrm{SWAP}(Z\otimes X)] + \mathrm{Tr}[\mathrm{SWAP}(X\otimes Z)\mathrm{SWAP}(Z\otimes X)])$$

$$= \prod_{i=1}^{n} \frac{1}{36}(\mathrm{Tr}[XZ\otimes ZX] + \mathrm{Tr}[\mathrm{SWAP}(XZ\otimes ZX)] + \mathrm{Tr}[\mathrm{SWAP}] + \mathrm{Tr}[I\otimes I])$$

$$= \prod_{i=1}^{n} \frac{1}{36}(0+2+2+4) = \left(\frac{2}{9}\right)^{n}. \tag{C37}$$

Combining Eqs. (C33), (C34), (C36), and (C37), we get

$$\mathbb{E}_{W_1,W_2}[|X(W_1,W_2)|] \leqslant \frac{1}{2}\sqrt{2\left(\frac{2}{9}\right)^{n} - 2\left(\frac{1}{9}\right)^{n}} \leqslant \left(\sqrt{\frac{2}{9}}\right)^{n}, \tag{C38}$$

since we assumed n to be odd. A similar derivation can be done for n even, using $\mathrm{Im}[z]^2 = (\frac{z-\bar{z}}{2i})^2 = -\frac{1}{4}(z^2 + \overline{z^2} - 2z\bar{z})$, which leads to the same upper bound.

Now, combining this bound with Eqs. (C31) and (C32), we get

$$\mathbb{E}_{W_1,W_2}[\mathrm{TV}(p_+^t, p_-^t)] \leqslant \sum_j w_{t,j} 2^n \left(\sqrt{\frac{2}{9}}\right)^{n} \leqslant \left(\sqrt{\frac{8}{9}}\right)^{n}, \tag{C39}$$

where we made use of $\sum_j w_{t,j} = 1$.

Given that $\mathrm{TV}(p_+^t, p_-^t)$ is positive and exponentially small with respect to n on average over all assignments of $(W_1, W_2) \in \mathcal{W}$, then, if we call

$$G(\psi_t, \phi_t) = \left\{(W_1, W_2) \in \mathcal{W} : \mathrm{TV}(p_+^t, p_-^t) \leqslant \left(\frac{8}{9}\right)^{n/4}\right\}, \tag{C40}$$

we have necessarily that

$$\frac{|G(\psi_t, \phi_t)|}{|\mathcal{W}|} \geqslant 1 - \left(\frac{8}{9}\right)^{n/4}. \tag{C41}$$

We prove this by contradiction. Assuming that Eq. (C41) does not hold, we have at least $(\frac{8}{9})^{n/4}|\mathcal{W}|$ assignments of $(W_1, W_2) \in \mathcal{W}$ for which $\mathrm{TV}(p_+^t, p_-^t) > (\frac{8}{9})^{n/4}$, and therefore

$$\mathbb{E}_{W_1,W_2}[\mathrm{TV}(p_+^t, p_-^t)] = \frac{1}{|\mathcal{W}|}\sum_{W_1,W_2\in\mathcal{W}} \mathrm{TV}(p_+^t, p_-^t) > \left(\frac{8}{9}\right)^{n/4} \times \left(\frac{8}{9}\right)^{n/4} = \left(\frac{8}{9}\right)^{n/2}, \tag{C42}$$

which contradicts with Eq. (C39).

From here, if we consider an entire sequence of measurements rather than a single measurement, we find that the resulting set of assignments for which $\mathrm{TV}(p_+^t, p_-^t)$ is exponentially small along the sequence:

$$G(\psi_{1:T}, \phi_{1:T}) = \left\{ (W_1, W_2) \in \mathcal{W} : \mathrm{TV}(p_+^t, p_-^t) \leqslant \left(\frac{8}{9}\right)^{n/4} \right. \\ \left. \forall t \in 1, \ldots, T \right\} \tag{C43}$$

still has an overwhelming large size, for subexponential T:

$$\frac{|G(\psi_{1:T}, \phi_{1:T})|}{|\mathcal{W}|} \geqslant 1 - T\left(\frac{8}{9}\right)^{n/4}. \tag{C44}$$

This follows immediately from the lower bound (C41) on the size of $G(\psi_t, \phi_t)$ for each (ψ_t, ϕ_t) along the sequence.

All we have left to do now is relate the total variation distance $\mathrm{TV}(p_+, p_-)$ to that of single measurements in a sequence, i.e., $\mathrm{TV}(p_+^t, p_-^t)$, for $t = 1, \ldots, T$. For that, we show inductively that

$$1 - \mathrm{TV}(p_+, p_-) \geqslant \min_{\psi_{1:T-1}, \phi_{1:T-1}} \prod_{t=1}^{T} (1 - \mathrm{TV}(p_+^t, p_-^t)). \tag{C45}$$

We start from the first measurement:

$$1 - \mathrm{TV}(p_+, p_-) \\ = \sum_{\psi_{1:T}, \phi_{1:T}} \min\left(p_+(\psi_{1:T}, \phi_{1:T}), p_-(\psi_{1:T}, \phi_{1:T})\right) \tag{C46}$$

$$= \sum_{\psi_1, \phi_1} \sum_{\psi_{2:T}, \phi_{2:T}} \min\left(p_+(\psi_1, \phi_1) p_+(\psi_{2:T}, \phi_{2:T}), \right. \\ \left. p_-(\psi_1, \phi_1) p_-(\psi_{2:T}, \phi_{2:T})\right) \tag{C47}$$

$$\geqslant \sum_{\psi_1, \phi_1} \min\left(p_+(\psi_1, \phi_1), p_-(\psi_1, \phi_1)\right) \\ \times \sum_{\psi_{2:T}, \phi_{2:T}} \min(p_+(\psi_{2:T}, \phi_{2:T}), p_-(\psi_{2:T}, \phi_{2:T})) \tag{C48}$$

$$\geqslant \sum_{\psi_1, \phi_1} \min\left(p_+(\psi_1, \phi_1), p_-(\psi_1, \phi_1)\right) \\ \times \min_{\psi_1, \phi_1} (1 - \mathrm{TV}(p_+^{2:T}, p_-^{2:T})) \tag{C49}$$

$$\geqslant (1 - \mathrm{TV}(p_+^1, p_-^1)) \min_{\psi_1, \phi_1} (1 - \mathrm{TV}(p_+^{2:T}, p_-^{2:T})), \tag{C50}$$

where we used $\mathrm{TV}(p, q) = \sum_{i, p_i \geqslant q_i} p_i - q_i = \sum_i p_i - \min(p_i, q_i) = 1 - \sum_i \min(p_i, q_i)$ for the first equality and the last two inequalities, and $\min(ab, cd) \geqslant \min(a, c)\min(b, d)$ for positive a, b, c, d in the first inequality.

We apply the same inductive reasoning for $t = 2, \ldots, T$ to get Eq. (C45). Since Eq. (C44) applies to any sequence, we then have from Eq. (C45) that there exists a $(1 - T(\frac{8}{9})^{n/4})$ fraction of $\mathcal{W}$ for which

$$1 - \mathrm{TV}(p_+, p_-) \geqslant \left(1 - \left(\frac{8}{9}\right)^{n/4}\right)^T \tag{C51}$$

$$\geqslant 1 - T\left(\frac{8}{9}\right)^{n/4}, \tag{C52}$$

which is equivalent to

$$\mathrm{TV}(p_+, p_-) \leqslant T\left(\frac{8}{9}\right)^{n/4} \leqslant T/2^{0.04n}, \tag{C53}$$

and therefore T must be scaling as $2^{\Omega(n)}$ for the algorithm to have high probability of success on all possible choices of $(W_1, W_2) \in \mathcal{W}$. ■

[1] A. Bisio, G. Chiribella, G. M. D'Ariano, S. Facchini, and P. Perinotti, Optimal quantum learning of a unitary transformation, Phys. Rev. A **81**, 032324 (2010).

[2] K. Poland, K. Beer, and T. J. Osborne, No free lunch for quantum machine learning, arXiv:2003.14103.

[3] K. Sharma, M. Cerezo, Z. Holmes, L. Cincio, A. Sornborger, and P. J. Coles, Reformulation of the no-free-lunch theorem for entangled datasets, Phys. Rev. Lett. **128**, 070501 (2022).

[4] S. Khatri, R. LaRose, A. Poremba, L. Cincio, A. T. Sornborger, and P. J. Coles, Quantum-assisted quantum compiling, Quantum **3**, 140 (2019).

[5] T. Jones and S. C. Benjamin, Robust quantum compilation and circuit optimisation via energy minimisation, Quantum **6**, 628 (2022).

[6] K. Heya, Y. Suzuki, Y. Nakamura, and K. Fujii, Variational quantum gate optimization, arXiv:1810.12745.

[7] C. Crstoiu, Z. Holmes, J. Iosue, L. Cincio, P. J. Coles, and A. Sornborger, Variational fast forwarding for quantum simulation beyond the coherence time, npj Quantum Inf. **6**, 82 (2020).

[8] J. Gibbs, K. Gili, Z. Holmes, B. Commeau, A. Arrasmith, L. Cincio, P. J. Coles, and A. Sornborger, Long-time simulations for fixed input states on quantum hardware, npj Quantum Inf. **8**, 135 (2022).

[9] J. Gibbs, Z. Holmes, M. C. Caro, N. Ezzell, H.-Y. Huang, L. Cincio, A. T. Sornborger, and P. J. Coles, Dynamical simulation via quantum machine learning with provable generalization, Phys. Rev. Res. **6**, 013241 (2024).

[10] H.-Y. Huang, R. Kueng, and J. Preskill, Information-theoretic bounds on quantum advantage in machine learning, Phys. Rev. Lett. **126**, 190505 (2021).

[11] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, J. Preskill, and J. R. McClean, Quantum advantage in learning from experiments, Science **376**, 1182 (2022).

[12] M. C. Caro, H.-Y. Huang, M. Cerezo, K. Sharma, A. Sornborger, L. Cincio, and P. J. Coles, Generalization in quantum machine learning from few training data, Nat. Commun. **13**, 4919 (2022).

[13] M. C. Caro, H.-Y. Huang, N. Ezzell, J. Gibbs, A. T. Sornborger, L. Cincio, P. J. Coles, and Z. Holmes, Out-of-distribution generalization for learning quantum dynamics, Nat. Commun. **14**, 3751 (2023).

[14] H.-Y. Huang, S. Chen, and J. Preskill, Learning to predict arbitrary quantum processes, PRX Quantum **4**, 040337 (2023).

[15] M. C. Caro, Learning quantum processes and Hamiltonians via the Pauli transfer matrix, ACM Trans. Quantum Comput. **5**, 1 (2024).

[16] H. Cai, Q. Ye, and D.-L. Deng, Sample complexity of learning parametric quantum circuits, Quantum Sci. Technol. **7**, 025014 (2022).

[17] Y. Mao, M. Gong, K. Nemoto, W. J. Munro, and J. Majer, Perspective on witnessing entanglement in hybrid quantum systems, Appl. Phys. Lett. **119**, 110501 (2021).

[18] E. Waks and C. Monroe, Protocol for hybrid entanglement between a trapped atom and a quantum dot, Phys. Rev. A **80**, 062330 (2009).

[19] C. Eichler, C. Lang, J. M. Fink, J. Govenius, S. Filipp, and A. Wallraff, Observation of entanglement between itinerant microwave photons and a superconducting qubit, Phys. Rev. Lett. **109**, 240501 (2012).

[20] A. Stute, B. Casabone, P. Schindler, T. Monz, P. O. Schmidt, B. Brandstätter, T. E. Northup, and R. Blatt, Tunable ion–photon entanglement in an optical cavity, Nature (London) **485**, 482 (2012).

[21] A. N. Craddock, J. Hannegan, D. P. Ornelas-Huerta, J. D. Siverns, A. J. Hachtel, E. A. Goldschmidt, J. V. Porto, Q. Quraishi, and S. L. Rolston, Quantum interference between photons from an atomic ensemble and a remote atomic ion, Phys. Rev. Lett. **123**, 213601 (2019).

[22] P. Scarlino, D. J. Van Woerkom, U. C. Mendes, J. V. Koski, A. J. Landig, C. K. Andersen, S. Gasparinetti, C. Reichl, W. Wegscheider, K. Ensslin, *et al.*, Coherent microwave-photon-mediated coupling between a semiconductor and a superconducting qubit, Nat. Commun. **10**, 3011 (2019).

[23] G. Vidal, Efficient classical simulation of slightly entangled quantum computations, Phys. Rev. Lett. **91**, 147902 (2003).

[24] R. Orús, A practical introduction to tensor networks: Matrix product states and projected entangled pair states, Ann. Phys. **349**, 117 (2014).

[25] A. Anshu, Z. Landau, and Y. Liu, Distributed quantum inner product estimation, in *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing* (Association for Computing Machinery, New York, 2022), pp. 44–51.

[26] J. Wang, S. Paesani, R. Santagati, S. Knauer, A. A. Gentile, N. Wiebe, M. Petruzzella, J. L. O'Brien, J. G. Rarity, A. Laing, and M. G. Thompson, Experimental quantum Hamiltonian learning, Nat. Phys. **13**, 551 (2017).

[27] N. Wiebe, C. Granade, C. Ferrie, and D. Cory, Quantum Hamiltonian learning using imperfect quantum resources, Phys. Rev. A **89**, 042314 (2014).

[28] A. A. Gentile, B. Flynn, S. Knauer, N. Wiebe, S. Paesani, C. E. Granade, J. G. Rarity, R. Santagati, and A. Laing, Learning models of quantum systems from experiments, Nat. Phys. **17**, 837 (2021).

[29] F. Wilde, A. Kshetrimayum, I. Roth, D. Hangleiter, R. Sweke, and J. Eisert, Scalably learning quantum many-body Hamiltonians from dynamical data, arXiv:2209.14328.

[30] V. Gebhart, R. Santagati, A. A. Gentile, E. Gauger, D. Craig, N. Ares, L. Banchi, F. Marquardt, L. Pezze, and C. Bonato, Learning quantum systems, Nat. Rev. Phys. **5**, 141 (2023).

[31] D. Stilck França, L. A. Markovich, V. Dobrovitski, A. H. Werner, and J. Borregaard, Efficient and robust estimation of many-qubit Hamiltonians, Nat. Commun. **15**, 311 (2024).

[32] J. Haah, R. Kothari, and E. Tang, Optimal learning of quantum Hamiltonians from high-temperature Gibbs states, Nat. Phys. **20**, 1027 (2024).

[33] A. Anshu, S. Arunachalam, T. Kuwahara, and M. Soleimanifar, Sample-efficient learning of interacting quantum systems, Nat. Phys. **17**, 931 (2021).

[34] L. Pastori, T. Olsacher, C. Kokail, and P. Zoller, Characterization and verification of Trotterized digital quantum simulation via Hamiltonian and Liouvillian learning, PRX Quantum **3**, 030324 (2022).

[35] D. Hangleiter, I. Roth, J. Fuksa, J. Eisert, and P. Roushan, Robustly learning the Hamiltonian dynamics of a superconducting quantum processor, Nat. Commun. **15**, 9595 (2024).

[36] D. K. Mark, J. Choi, A. L. Shaw, M. Endres, and S. Choi, Benchmarking quantum simulators using ergodic quantum dynamics, Phys. Rev. Lett. **131**, 110601 (2023).

[37] M. Fanizza, Y. Quek, and M. Rosati, Learning quantum processes without input control, PRX Quantum **5**, 020367 (2024).

[38] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, Nat. Phys. **16**, 1050 (2020).

[39] A. Elben, S. T. Flammia, H.-Y. Huang, R. Kueng, J. Preskill, B. Vermersch, and P. Zoller, The randomized measurement toolbox, Nat. Rev. Phys. **5**, 9 (2023).

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

[41] We note that a similar argument was made in passing in Ref. [13], but the argument was not formalized and its implications for incoherent learning were not appreciated.

[42] D. Maslov and M. Roetteler, Shorter stabilizer circuits via Bruhat decomposition and quantum circuit transformations, IEEE Trans. Inform. Theory **64**, 4729 (2018).

[43] One way around this would be to build prior knowledge about the target unitary into the guess ansatz. It has been shown, for example, that highly symmetrized ansätze can avoid globality-induced barren plateaus [60]. Nonetheless, such approaches are unlikely to work in all cases.

[44] M. P. da Silva, O. Landon-Cardinal, and D. Poulin, Practical characterization of quantum devices without tomography, Phys. Rev. Lett. **107**, 210404 (2011).

[45] A. Basheer, Y. Feng, C. Ferrie, and S. Li, Alternating layered variational quantum circuits can be classically optimized efficiently using classical shadows, in *Proceedings of the AAAI Conference on Artificial Intelligence* (AAAI Press, Washington, DC, 2023), Vol. 37, pp. 6770–6778.

[46] M. Cerezo, M. Larocca, D. Garcia-Martin, N. L. Diaz, P. Braccia, E. Fontana, M. S. Rudolph, P. Bermejo, A. Ijaz, S. Thanasilp, *et al.*, Does provable absence of barren plateaus imply classical simulability, Nat. Commun. **16**, 7907 (2025).

[47] H.-Y. Hu, M. Ma, W. Gong, Q. Ye, Y. Tong, S. T. Flammia, and S. F. Yelin, Ansatz-free Hamiltonian learning with Heisenberg-limited scaling, PRX Quantum **6**, 040315 (2025).

[48] Random product states enjoy the same out-of-distribution generalization guarantees as products of random single-qubit stabilizer states [13].

[49] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: Composable transformations of Python + NumPy programs, GitHub (2018), https://github.com/jax-ml/jax.

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

[51] K. Mizuta, Y. O. Nakagawa, K. Mitarai, and K. Fujii, Local variational quantum compilation of large-scale Hamiltonian dynamics, PRX Quantum **3**, 040302 (2022).

[52] E. Younis, K. Sen, K. Yelick, and C. Iancu, QFAST: Conflating search and numerical optimization for scalable quantum circuit synthesis, in *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, New York, 2021), p. 232.

[53] T. Patel, E. Younis, C. Iancu, W. de Jong, and D. Tiwari, QUEST: Systematically approximating quantum circuits for higher output fidelity, in *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'22)* (Association for Computing Machinery, New York, 2022), p. 514.

[54] F. Vatan and C. Williams, Optimal quantum circuits for general two-qubit gates, Phys. Rev. A **69**, 032315 (2004).

[55] M. M. Wilde, *Quantum Information Theory*, 2nd ed. (Cambridge University Press, Cambridge, 2017).

[56] H. Zhao, L. Lewis, I. Kannan, Y. Quek, H.-Y. Huang, and M. C. Caro, Learning quantum states and unitaries of bounded gate complexity, PRX Quantum **5**, 040306 (2024).

[57] B. Yu, Assouad, Fano, and Le Cam, in *Festschrift for Lucien Le Cam: Research Papers in Probability and Statistics*, edited by P. David, T. Erik, and L. Y. Grace (Springer, New York, 1997), p. 423.

[58] M. StackExchange, Total variation inequality for the product measure (2011), https://math.stackexchange.com/questions/72315/total-variation-inequality-for-the-product-measure.

[59] D. A. Roberts and B. Yoshida, Chaos and complexity by design, J. High Energy Phys. 04 (2017) 121.

[60] L. Schatzki, M. Larocca, Q. T. Nguyen, F. Sauvage, and M. Cerezo, Theoretical guarantees for permutation-equivariant quantum neural networks, npj Quantum Inf. **10**, 12 (2024).