

Classical shadows for quantum process tomography on near-term quantum computers

Ryan Levy^{1,2,*}, Di Luo,^{1,3,4,5,*} and Bryan K. Clark¹

¹*Institute for Condensed Matter Theory and IQIUST and NCSA Center for Artificial Intelligence Innovation and Department of Physics, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*

²*Center for Computational Quantum Physics, Flatiron Institute, New York, New York 10010, USA*

³*The NSF AI Institute for Artificial Intelligence and Fundamental Interactions*

⁴*Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

⁵*Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA*



(Received 16 November 2021; revised 6 July 2023; accepted 20 September 2023; published 10 January 2024)

Quantum process tomography is a powerful tool for understanding quantum channels and characterizing the properties of quantum devices. Inspired by recent advances using classical shadows in quantum state tomography [H.-Y. Huang, R. Kueng, and J. Preskill, *Nat. Phys.* **16**, 1050 (2020).], we have developed ShadowQPT, a classical shadow method for quantum process tomography. We introduce two related formulations with and without ancilla qubits. ShadowQPT stochastically reconstructs the Choi matrix of the device allowing for an *a posteriori* classical evaluation of the device on arbitrary inputs with respect to arbitrary outputs. Using shadows, we then show how to compute overlaps, generate all k -weight reduced processes, and perform reconstruction via Hamiltonian learning. These latter two tasks are efficient for large systems as the number of quantum measurements needed scales only logarithmically with the number of qubits. A number of additional approximations and improvements are developed, including the use of a pair-factorized Clifford shadow and a series of postprocessing techniques that significantly enhance the accuracy for recovering the quantum channel. We have implemented ShadowQPT using both Pauli and Clifford measurements on the IonQ trapped ion quantum computer for quantum processes up to $n = 4$ qubits, and we achieved good performance.

DOI: [10.1103/PhysRevResearch.6.013029](https://doi.org/10.1103/PhysRevResearch.6.013029)

I. INTRODUCTION

Quantum technologies have been developing rapidly in recent years. Quantum devices are used for various applications, such as quantum metrology, quantum teleportation, and quantum simulation [1–5]. One important task in the development of near-term quantum devices is their characterization.

For noisy intermediate-scale quantum (NISQ) technology, a first step to characterization is to use quantum state tomography (QST), which attempts to characterize the output state of the quantum circuit. A next crucial step, quantum process tomography (QPT), characterizes not just one output state but the entire quantum dynamics of the device. Earlier attempts at QPT used the linear inversion method [6,7]. Later various statistical methods were developed, including maximum-likelihood methods [8–12], Bayesian methods [13–15], compressed sensing methods [16], tensor network methods [17], and other optimization techniques [18–24]. Theoretically, quantum process tomography can be related to quantum state tomography through the Jamiołkowski

process-state isomorphism [25,26]. One recent important advancement in QST comes from classical shadow tomography [27], which allows for the prediction of multiple observables with few quantum measurements both theoretically and experimentally [28–41].

In this work, we propose a classical shadows algorithm for quantum process tomography on near-term quantum computers giving both a theoretical analysis and explicit implementations on an IonQ quantum computer. Using Jamiołkowski process-state isomorphism and techniques in process tomography, we develop the first classical shadow quantum process tomography (ShadowQPT).

In Sec. II, we prove two theorems related to the effectiveness of ShadowQPT. Theorem II.2 bounds the amount of ShadowQPT data needed to compute the overlap between any target density matrix σ with the output of the measured quantum process on any input ρ . Theorem II.1 discusses the complexity of process matrix tomography. In Sec. III, we show Corollary [cor1], which describes the sample complexity scaling of applying ShadowQPT to Hamiltonian learning, which is logarithmic in system size, following from Theorem II.1. In Sec. IV, we develop shadow algorithms with both Pauli and Clifford measurements. Both an ancilla-based scheme as well as a two-sided scheme which applies unitaries both before and after the channel are introduced. We proceed to show practical improvements to these algorithms for near-term devices by using two qubit Clifford unitaries instead of a global Clifford and using multiple repetitions per circuit.

*These authors contributed equally to this paper.

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.

Furthermore, we develop a series of postprocessing techniques, such as projecting into the space of physical channels and purifying outputs, which significantly improve our results. In Sec. V, we implement our algorithms on IonQ quantum computers for both unitary and nonunitary process tomography on $n = 2, 3, 4$ qubits systems. We benchmark the process matrix construction, and we test predicting the overlap for pairs of input and output pure states and compare to direct measurement outcomes. In Sec. VI, we numerically simulate Hamiltonian learning and reconstruct a random one dimensional (1D) Ising model using ShadowQPT, showing efficient scalability with system size. We conclude in Sec. VII with a discussion on opportunities for and future explorations of our ShadowQPT algorithms.

II. THEORETICAL ANALYSIS OF SHADOWQPT

In this work, we have developed shadow tomography algorithms for quantum process tomography. The key idea is based on the recent development of shadow tomography on a quantum state [27] and the Jamiołkowski process-state isomorphism [25]. An important result in classical shadow tomography states that [27] classical shadows of size N suffice to predict M linear functions $\{\text{tr}(O_i \rho)\}_{i=1, \dots, M}$ up to additive error ϵ given that $N \geq (\text{order}) \log(M) \max_i \|O_i\|_{\text{shadow}}^2 / \epsilon^2$.

To connect quantum process tomography with quantum state tomography, we utilize the Jamiołkowski process-state isomorphism [25]. It provides a positive-semidefinite operator representation $\Lambda_{\mathcal{E}}$ for an n -qubit process $\mathcal{E}$, which is the Choi matrix representation [26]

$$\Lambda_{\mathcal{E}} = (\mathbb{I} \otimes \mathcal{E})(|\phi^+\rangle\langle\phi^+|^{\otimes n}), \quad (1)$$

where $|\phi^+\rangle = |00\rangle + |11\rangle$. For an input density matrix ρ to the channel $\mathcal{E}$, we have $\mathcal{E}(\rho) = \text{tr}[(\rho^T \otimes \mathbb{I})\Lambda_{\mathcal{E}}]$. In the rest of the paper, we implicitly denote $\Lambda_{\mathcal{E}}$ as Λ . Notice that $\text{tr}(\Lambda_{\mathcal{E}}) = 2^n$, and we further denote the normalized process matrix Λ with trace 1 as $\rho_{\Lambda} = \frac{1}{2^n} \Lambda$. We notate a reduced k -qubit ($k \leq n$) process Choi matrix as the Choi matrix representation for the quantum channel on a particular k -qubit subsystem where the other $n - k$ qubits are traced out.

Based on the classical shadow tomography property and the Choi matrix representation, we state the following two informal versions of the theorems for classical shadow tomography for a quantum process.

Theorem II.1. For an n -qubit process Choi matrix ρ_{Λ} and $\epsilon, \delta \in (0, 1)$, the number of random global Clifford measurements N that suffice to simultaneously predict any reduced k -qubit process Choi matrix $\rho_{\Lambda^{(k)}}$ with Frobenius norm error up to ϵ with probability $1 - \delta$ is of order $\frac{4^{n+k}}{\epsilon^2} \log(2(8n)^{2k}/\delta)$.

Meanwhile, the number of Pauli-6 positive operator-valued measurements (POVMs) N that suffice to simultaneously predict any reduced k -qubit process Choi matrix $\rho_{\Lambda^{(k)}}$ with Frobenius norm error up to ϵ with probability $1 - \delta$ is of order $\frac{36^k}{\epsilon^2} \log(2(8n)^{2k}/\delta)$, and trace norm error up to ϵ with probability $1 - \delta$ is of order $\frac{144^k}{\epsilon^2} \log((24n)^{2k}/\delta)$.

The Pauli-6 POVM measurement includes $\{\frac{1}{3}|0\rangle\langle 0|, \frac{1}{3}|1\rangle\langle 1|, \frac{1}{3}|+\rangle\langle +|, \frac{1}{3}|-\rangle\langle -|, \frac{1}{3}|r\rangle\langle r|, \frac{1}{3}|l\rangle\langle l|\}$, where $(|0\rangle, |1\rangle), (|+\rangle, |-\rangle), (|r\rangle, |l\rangle)$ are eigenvectors of σ_z, σ_x , and σ_y . We note that the Pauli-6 POVM and the random

single-qubit Clifford measurement usually share similar results since the two measurements have the same shadow norm for factorized Pauli observables [27, 36]. In this work, the term “random Pauli measurement” or “Pauli” in the figures refers to a Pauli-6 POVM, though one can also apply a random single-qubit Clifford in those contexts.

Notice that for any fixed k , the shadow methods allow simultaneous prediction of all $\binom{n}{k}$ reduced k -qubit processes in time logarithmic in n . This is useful for learning quantum dynamics of local observables, and it provides a foundation for us to develop an efficient k -local Hamiltonian learning scheme (see Sec. III). For full process tomography ($k = n$), Theorem II.1 states that the quantum measurement complexity of ShadowQPT scales exponentially, which is consistent with the known lower bound. We note that each measurement of ShadowQPT gives a full unbiased (potentially noisy) stochastic representation of the entire Choi matrix; this stands in contrast with standard maximum-likelihood estimation (MLE), which requires exponential classical postprocessing to generate the Choi matrix even from small amounts of data.

Theorem II.1 focuses on predicting the complete set of reduced k -qubit process Choi matrices. Instead, one can ask about computing the result of the quantum channel run on a series of (input ρ , output σ) pairs, i.e., the quantum device run on the density matrix ρ and traced against the observable σ .

Theorem II.2. For an n -qubit quantum process ρ_{Λ} and $\epsilon, \delta \in (0, 1)$, given a set of density matrix pairs $\{(\rho_i^{\text{in}}, \sigma_i), \dots, (\rho_M^{\text{in}}, \sigma_M)\}$, the number of measurements N that suffice to predict $\text{tr}((\rho_i^{\text{in}T} \otimes \sigma_i)\rho_{\Lambda})$ for any i up to error ϵ with probability $1 - \delta$ is of order

$$\frac{\log(2M/\delta)}{\epsilon^2} \max_i \left\| O_i - \frac{\text{tr}(O_i)}{2} \mathbb{I} \right\|_{\text{shadow}}^2, \quad (2)$$

where $O_i = \rho_i^{\text{in}} \otimes \sigma_i$, and $\|\cdot\|_{\text{shadow}}$ is the shadow norm (see Theorem A.1 in the Appendix for the definition).

For random global Clifford measurement, it requires order $\frac{\log(2M)}{\epsilon^2} \max_i \text{tr}(O_i^2)$. In particular, if ρ_i^{in} and σ_i are all pure states, then $\text{tr}(O_i^2) = 1$. For random single-qubit Clifford measurement, it requires order $\frac{\log(2M)}{\epsilon^2} \max_i 4^k \|O_i\|_{\infty}^2$, where O_i acts nontrivially on k -qubits, and $\|\cdot\|_{\infty}$ is the spectral norm. If ρ_i^{in} and σ_i are all pure states, then $\|O_i\|_{\infty}^2 = 1$.

If σ_i and ρ_i^{in} are k -qubit reduced density matrices with the same support, then the overlap between σ_i and ρ_i^{in} through the channel corresponds to $\text{tr}(\mathcal{E}(\rho_i^{\text{in}})\sigma_i) = 2^k \text{tr}(\rho_{\Lambda^{(k)}}(\rho_i^{\text{in}T} \otimes \sigma_i))$. Note that the Clifford ShadowQPT allows us to compute the overlap of pairs of pure states $(\rho^{\text{in}}, \sigma)$ using a number of measurements which scale logarithmically with the number of pairs being considered. If ρ^{in} and σ are stabilizer states, the classical resources required to compute are also efficient as one can use the stabilizer algorithm [42].

Notice that to compute M pairs of $\text{tr}(\mathcal{E}(\rho_i^{\text{in}})\sigma_i)$, ShadowQPT requires time independent of n and scales as $\log(M)$, when ρ_i^{in} and σ_i have support over a k -qubit system. The detailed proofs for Theorems II.1 and II.2 and are given in Appendix 1.

III. HAMILTONIAN LEARNING VIA SHADOWQPT

A corollary of Theorem II.1 implies that certain applications of ShadowQPT can also be efficient. One such application is Hamiltonian learning [43,44], which seeks to determine an unknown Hamiltonian given access to a time evolution operator $U(t)$. In this context, we consider applications of ShadowQPT to Hamiltonian learning. For a k -local Hamiltonian $H = \sum_i c_i h_i$, h_i is a k -body operator as a Pauli string and c_i is the coefficient to learn.

One scheme that we develop is to learn c_i through ShadowQPT on e^{-iHt} . Notice that access to e^{-iHt} also gives us an ability to time evolve to all t' which are integer multiples of t . Consider t sufficiently small such that $U = e^{-iHt} \approx I - iHt = I - it \sum_i c_i h_i$, then U can be approximated by a sum of k -local operators. To characterize U , we first perform ShadowQPT on U and attain a representation of $\rho_{\Lambda^{(k)}}$. To learn c_i , we set $\Lambda(\rho) = U \rho U^\dagger$ and we choose a k' -local operator p_i where $k' < k$ such that $q_i \equiv [h_i, p_i] \neq 0$, and then classically evaluate

$$\text{tr}(\Lambda(p_i)q_i) = \text{tr}((p_i - it[H, p_i])q_i) + O(t^2) \quad (3)$$

$$= \text{tr}\left(\left(p_i - it \sum_j c_j q_j\right)q_i\right) + O(t^2); (4)$$

$$\text{tr}((p_i^T \otimes q_i)\rho_\Lambda) = -it \left(c_i + \sum_{j \neq i} g_j c_j\right) + O(t^2), \quad (5)$$

where $g_j = \text{tr}(q_j q_i) / \text{tr}(q_i q_i) = \text{tr}(q_j q_i) / 2^n$. Hence we can attain c_i from evaluating data from ShadowQPT and solving a system of equations. Note that if $\text{tr}(q_j q_i) = \delta_{ij} \forall i, j$, then $\text{tr}((p_i^T \otimes q_i)\rho_\Lambda) = -it c_i + O(t^2)$. We focus on this case, which is true for, e.g., the transverse field Ising model. According to Theorem II.1, the sampling complexity S of estimating c_i is given by the following.

Corollary III.1. Consider an unknown k local Hamiltonian $H = \sum_i c_i h_i$, where h_i is a k -body operator and c_i are real coefficients. Given access to $U(t) = e^{iHt}$, then ShadowQPT can estimate c_i up to error ϵ with probability $1 - \delta$ with sample complexity S ,

$$S \sim \frac{36^k}{t^2 \epsilon^2} 2k \log(8n) \log(1/\delta), \quad (6)$$

where t is sufficiently small that the linear approximation error is much smaller than ϵ .

Thus for fixed k , ShadowQPT can efficiently predict c_i in $O(\log(n))$ samples. Note that one can systematically reduce the (currently linear) approximation error by a higher-order Taylor expansion at the cost of additional classical processing. We also note that our approach is related to the latest version of Ref. [44] on real-time dynamics of Hamiltonian learning, which is proved in an alternative way by extending the results of finite-temperature Hamiltonian learning.

IV. SHADOWQPT POSTPROCESSING ALGORITHMS

In this work, we perform classical shadow process tomography algorithms on a unitary process as well as a nonunitary




FIG. 1. (a) Unitary channel $\mathcal{E}$, which is the GHZ process that can create a GHZ state for n qubits from $|0\rangle^{\otimes n}$. (b) Circuit for two-sided quantum process tomography to produce a Choi matrix Λ of a quantum channel $\mathcal{E}$. The application of two unitary circuits, denoted $U_L^\dagger$ and U_R (e.g., random Pauli/Clifford), can be applied to reconstruct the channel. (c) Circuit for ancilla-based quantum process tomography to produce a Choi matrix Λ of a quantum channel $\mathcal{E}$. After preparation of a Bell state input (dashed box) to $\mathcal{E}$, a unitary rotation U (e.g., random Pauli/Clifford) is applied before measurement.

reduced process. We exemplify our approach on the unitary channel Λ , which is the Greenberger-Horne-Zeilinger (GHZ) process and its reduced process. The GHZ process Λ is shown in Fig. 1(a) and it generates an n -qubit GHZ state [45] $|\text{GHZ}_n\rangle = (|0\rangle^{\otimes n} + |1\rangle^{\otimes n})/\sqrt{2}$ from $|0\rangle^{\otimes n}$. The GHZ processes are constructed such that the circuits have a depth of 2,3,3 for $n = 2, 3, 4$ qubits, respectively [46].

A. Unitary full process classical shadows

In this section, we will introduce two schemes for ShadowQPT, namely the two-sided scheme and the ancilla-based scheme. The two-sided scheme is shown in Fig. 1(b), where random unitary circuits $U_L^\dagger$ and U_R are applied to the left- and right-hand side of the channel. For the ancilla-based scheme, we use the circuit shown in Fig. 1(c). This circuit produces a $2n$ qubit output density matrix which represents the n qubit normalized Choi matrix ρ_Λ , for the n qubit process $\mathcal{E}$. This approach is often called ancilla-assisted quantum process tomography [47] (AAPT). In our work, for the ancilla-based approach we perform shadow tomography on this $2n$ qubit state. It turns out that the two-sided scheme can be viewed as a special case of the ancilla-based scheme, where the random unitary is factorized into a product of two unitaries U_R and U_L , respectively, which act on the top and bottom n wires each. The details for the proof of equivalence can be found in Appendix 2.

Both schemes can realize a random Pauli measurement scheme. For the ancilla-based scheme, the Pauli-6 POVM measurement can be realized by randomly applying a series of basis rotation gates chosen from $G_R = \{I, H, SH\}$ as the unitary U after the channel and measuring in the computational basis, where H is the Hadamard gate, S is the phase gate, and I is the identity. For the two-sided scheme, one can realize the Pauli-6 POVM measurements by randomly selecting U_L, U_R as a tensor product from $G = \{I, H, SH, X, HX, SHX\}$, where X is the Pauli X gate. In practice, this selection from G only needs to be done for U_L ; for U_R we can still randomly draw from G_R , just as in the ancilla-based scheme, since the application of the last X is irrelevant if a measurement is immediately performed afterwards. It is also possible to replace the Pauli-6 POVM measurements with random single-qubit Clifford gates ($k = 1$). This may be advantageous on certain quantum devices [33]. We numerically simulate their equivalence in Fig. 14.

Additionally, both schemes can utilize random Clifford measurements, which can have better performance or scaling compared to Pauli measurements. Since most NISQ devices have limited gate depth, a global n or $2n$ qubit random Clifford unitary needs gate depth $O(n^2/\log n)$ in general, which could be challenging to implement accurately; instead, we use a tensor product of two-qubit Clifford unitaries ($k = 2$) to realize the improved scaling of the Clifford group while minimizing the additional gate depth to the circuit. Random Clifford circuits are generated via the Qiskit library [48], which uses an optimal number of control-NOT gates for two qubits [49].

After we attain the random Pauli/Clifford measurements, we can construct a Choi matrix from the classical shadows (denoted *Pauli* or *Clifford* in figures). To be precise, after applying a unitary U and measuring in the computational basis a bitstring $|b_i\rangle \in \{0, 1\}^{\otimes n}$, we form the i th classical shadow $\hat{\Lambda}_i$ as

$$\hat{\Lambda}_i = 2^n \mathcal{M}^{-1}(U_i^\dagger |b_i\rangle \langle b_i| U_i), \quad (7)$$

where the inverse channel is $\mathcal{M}_n^{-1}(X) = (2^n + 1)X - \mathbb{I}$ for n -qubit Clifford circuits or $\mathcal{M}_n^{-1}(X) = \otimes_i \mathcal{M}_k^{-1}(X)$ for Pauli ($k = 1$) and two-qubit Clifford ($k = 2$) circuits. These objects can be held in memory efficiently, as Clifford circuits and

bitstrings have a polynomial representation as well as polynomial evaluation cost [42]. We average over all classical shadows $\hat{\Lambda}_i$ to obtain the final shadow reconstructed Choi matrix Λ^O . Note that it is suggested in Ref. [27] to utilize a median-of-means rather than direct mean. Our data have little dependence on this difference after projection (see Fig. 8), which agrees with the observations of Ref. [32].

We add an additional extension in which a given unitary is repeatedly measured under some number of repetitions. In practice, it is often the case that loading a different unitary on a quantum computer is significantly slower than taking additional measurements over the same unitary. The above theorems do not directly apply in this case, but we show via simulations (see Fig. 12) that including additional repetitions improves the quality of the result, up to a saturation threshold.

B. Nonunitary reduced process classical shadows

An interesting additional extension is the ability to characterize a nonunitary quantum process. By separating the n qubit process into subsystems of size 1 and up to size $n - 1$, we can produce a reduced process on a subsystem of up to $n - 1$ qubits. All subsets of the GHZ process are then nonunitary. Note that a general process $\Lambda_\mathcal{E}$ need not be unitary, nor must all subsystems be nonunitary. In this work, we focus on a unitary GHZ process Λ and simultaneously measuring all corresponding nonunitary reduced processes on each subsystems.

In a Pauli measurement based tomography scheme, including Pauli shadow tomography, we can directly use the operators that act over the relevant subsystem. Here, the partial trace over the Choi matrix is equivalent to this direct scheme, although for many expensive postprocessing schemes it is far easier to optimize in the subsystem space directly.

For a Clifford shadows scheme, however, a general n -qubit Clifford (or $2n$) will not be separable in the same manner as a Pauli-based scheme. Despite this, if a particular subsystem or set of subsystems is known beforehand, the Clifford unitaries can be chosen to respect separability. However, the tradeoff between the benefits of using nonseparable large n -qubit Clifford gates or using separable ones has not been well-studied.

C. Postprocessing of classical shadows

A physical quantum process Λ is a completely positive trace-preserving ($CPTP$) map between density matrices. For any positive-semidefinite density matrix ρ^{in} with trace 1, $\mathcal{E}(\rho^{\text{in}})$ should also be positive-semidefinite with unit trace. Even though the shadow reconstructed Choi matrix in Eq. (1) has the correct trace, it is not a valid $CPTP$ map in general. We further explore various postprocessing projection techniques for improving the classical shadow reconstruction in this section.

We denote the projection of Λ^O as Λ' . To project Λ^O into the space of completely positive preserving (CP) matrices, we use the technique of Ref. [12], which removes negative eigenvalues and rescales the remaining non-negative eigenvalues. We find this operation has the biggest reduction of trace distance from the target density matrix in our data. This

does not guarantee that Λ' is $\mathcal{CPTP}$, but it ensures that the resulting Λ' is normalized and positive-semidefinite.

In cases in which we know that the true underlying process should be unitary, we can remove extraneous statistical noise by purifying the shadow reconstructions to create a unitary shadow Choi matrix. Purifying in this manner can also be seen as an extreme $\mathcal{CP}$ -projection method, in which all but the dominant eigenvectors are thrown out. We find that this approach significantly improves our result in the case in which the target process is indeed unitary.

D. Predicting overlaps

Given a set of density matrix pairs $(\rho_i^{\text{in}}, \sigma_i)$ as in Theorem II.2, we can utilize shadow reconstructed Choi matrices to obtain predictions of the overlap $\text{tr}[\mathcal{E}(\rho_i^{\text{in}})\sigma_i]$; we work with the pure state $\sigma_i(\theta) = U(\theta)^\dagger|0\rangle$, where θ is a set of parameters which specify the quantum circuit that generates U , and we use $|0\rangle$ as a placeholder for $|0\rangle^{\otimes n}$.

When evolving an initial density matrix ρ^{in} through the process generated by Λ^O to produce ρ^{out} , we additionally project a $\mathcal{CP}$ -projected Choi matrix into the trace-preserving ($\mathcal{TP}$) space using the method outlined in Ref. [18], followed by projecting ρ^{out} into $\mathcal{CP}$ space. To be precise, we have $\Lambda' = \mathcal{TP}(\mathcal{CP}[\Lambda^O])$ and $\rho^{\text{out}} = \mathcal{CP}[\text{tr}(\rho^{\text{in}T} \otimes \mathbb{I})\Lambda']$.

We then take the pure state density matrix σ_i to calculate the overlap $\text{tr}[\rho_i^{\text{out}}\sigma_i] = \langle 0|U(\theta)\rho_i^{\text{out}}U(\theta)^\dagger|0\rangle$. For this to be classically efficient (even without postprocessing), σ_i has to be representable as a short tensor network for Pauli or be within the Clifford group for Clifford unitaries, respectively.

When scaling to large numbers of qubits, the postprocessing required to generate Λ' becomes infeasible, so we introduce a more efficient routine. To compute the overlap, we consider $\text{tr}[(\rho_i^{\text{in}T} \otimes \sigma_i)\Lambda'^O]$, where Λ'^O here represents a purified Λ^O without any additional projection.

V. EXPERIMENTAL RESULTS OF SHADOWQPT ON QUANTUM HARDWARE

We exemplify our ShadowQPT methods via quantum circuit measurements on the IonQ trapped ion quantum device. The IonQ system consists of 11 trapped ion qubits [50], available via Amazon Braket on Amazon Web Services (AWS). At the time of simulation, the machine was reported¹ to have a one-qubit gate average fidelity of 0.997 17, a two-qubit gate average fidelity of 0.9696, and a state preparation and measurement (SPAM) mean of 0.9961, which agrees with the data presented in Ref. [50]. We use 51 200 shadows each for $n = 2, 3, 4$ in the experiments with at most 1024 unitary circuits; see Appendix 3.

Included against our ShadowQPT results is a maximum-likelihood estimation (MLE) for reconstructing the Choi matrix using our random Pauli measurements. We use the iterative method outlined in Ref. [10]. Due to the size of the Clifford POVM, Clifford-MLE reconstructions are done within the space of observed measurements, as is $n = 4$ Pauli data. Notice that the MLE method is classically exponen-

tially costly in general. To perform MLE, we begin with a random complex matrix, taking at least 100 iteration steps such that the iteration has converged, i.e., $\|\rho_k - \rho_{k-1}\| \leq 10^{-3}$, or that the L_2 -norm of the difference between steps has reached a small value. We include purification of the resulting Choi matrix alongside our shadow reconstructions as well.

We explore the effect of using “fixed” and “nonfixed” schemes by using either a number of random Clifford unitaries or a set of random Clifford unitaries that are separable between the first qubit and the rest of the system. For a “nonfixed” unitary, a two-qubit Clifford unitary may be applied to any two wires of the circuit. For the “fixed” unitary, one Clifford is applied to qubits 1 and $N + 1$ (its corresponding ancilla), and the remaining Clifford unitaries are randomly applied in the “nonfixed” scheme. Our data have 0%, 50%, and 43.75% of the Clifford unitaries manually fixed for $n = 2, 3, 4$, respectively. These various schemes alternatively target, respectively, a situation in which one might ask about any reduced density matrix *a posteriori* (i.e., the nonfixed scheme), or target density matrices on a particular fixed set of qubits (i.e., the fixed scheme).

A. Full process

Using classical shadows, we directly reconstruct the Choi matrix and compare it using the normalized trace distance in Fig. 3,

$$T(\Lambda_1, \Lambda_2) = \frac{1}{2^{n+1}} \text{tr}[\sqrt{(\Lambda_1 - \Lambda_2)^\dagger(\Lambda_1 - \Lambda_2)}]. \quad (8)$$

First, to illustrate the scaling of ShadowQPT, a comparison of the Pauli and Clifford shadow measurements using a uniform number of shadows is shown by solid lines in Fig. 3. While each method is inherently exponential in scaling across process size, $k = 2$ Clifford measurements perform better than Pauli measurements, where a full $k = 2n$ Clifford measurement scheme should have performed the best.

In dashed lines, we then show the results when using all of our available measurements (i.e., multiple repetitions of the same unitary). A visualization of each of these Choi matrices is shown in Fig. 2 (see Fig. 9 for a 3D version of the plot). We see that even in the regime where all possible Pauli strings are measured ($n = 2, 3$), the Clifford circuits have a smaller trace distance.

If we consider that our underlying process should be unitary, we can additionally purify our noisy Choi matrix (dash-dotted lines). Purification appears to have an enormous effect, improving the trace distance by up to a factor of $34\times$. For all n , purified Clifford shadows have a smaller trace distance compared to the Pauli shadows, but this is followed closely by the purified Clifford MLE reconstruction, which has a narrowly smaller distance for $n = 3, 4$ (see a further comparison in Fig. 10).

Here we have selected only the purification postprocessing step, but there are additional projections, namely $\mathcal{CP}$ -projection of Ref. [12], that one may do. Shown in the right of Fig. 3 is an example of the effects on the trace distance each postprocessing method has for $n = 3$. We see a dramatic improvement in the trace distance after purification or

¹This information can be found via the AWS status page.




FIG. 2. Visualization of the GHZ process Choi matrix for various numbers of qubits (rows) constructed exactly or via classical shadow quantum process tomography. The first column corresponds to the exact unitary process Λ , which has no imaginary component. The next columns correspond to real and imaginary parts of the Pauli/Clifford shadow reconstructed Choi matrix Λ^O followed by their corresponding purification. The $N = 4$ visualization is zoomed into the first 60×60 elements.

projection, with Pauli shadows marginally performing better than the Clifford shadows for this system size.

As an additional comparison, we compare ShadowQPT with both an ancilla-based scheme and a two-sided scheme, shown in Fig. 4. On the left, we show simulations of both Pauli and $k = n$ Clifford measurements for a single and multiple repetitions. The simulations use 512 unitaries and 1 (50) repetitions for single (multiple) repetitions. In the case of single repetitions, we find nearly identical exponential scaling but with Clifford measurements with a smaller trace distance. With the addition of multiple repetitions, however, the ancilla schemes obtain a smaller trace distance than their two-sided counterparts. As each repetition effectively changes both the left and right unitary for the ancilla scheme but only changes the right unitary for the two-sided version, this result is un-

surprising. We provide more details on the equivalence of the schemes in Appendix 2.

Then on the right of Fig. 4, we compare two examples of Pauli and Clifford measurements and the effect of post-processing methods. For the two-sided scheme, we use 512 Clifford unitaries at 50 repetitions for each run on the IonQ device, and for the ancilla scheme we randomly resample our IonQ data to match. Note that our ancilla measurements are with randomized $k = 2$ Clifford unitaries. Pauli measurements are done with 25 600 random measured Pauli strings. We find that for both Pauli and Clifford measurements, the ancilla-based scheme has a lower normalized trace distance without postprocessing and under purification, and a larger trace distance for Clifford $\mathcal{CP}$ -projection.

B. Nonunitary reduced process

We further study two examples of a nonunitary reduced process on a subset of the qubits. In Figs. 5(a) and 5(b) we show the average over all one- and two-qubit processes. Analogous to the previous section, we show a reduced data set by solid lines, and the fully postprocessed data, here $\mathcal{CP}$ -projection, by dot-dashed lines. In general, we expect the full $2n$ -Clifford measurements, simulation shown in green, to grow exponentially with system size. As we have fixed the size of the process, we see that Pauli measurements do not scale with respect to system size, and there is a weak scaling for the two-qubit Clifford measurement scheme.

For all system sizes n , a classical shadow routine produces comparable or better trace distance compared with MLE, with an application of $\mathcal{CP}$ -projection after evolution. In addition, despite having limited separability compared to the Pauli matrices, the performance of the Clifford shadows is competitive with the Pauli shadows.

C. Predicting overlaps

We can also compute overlaps using the shadow reconstructions formalism of Theorem II.2. For the choice of σ_i , we use circuits parametrized by $R_x(\theta)$ or $R_y(\theta)$ as input to a GHZ circuit; see the Appendix for more details on the choices




FIG. 3. Left: Normalized trace distance $T(\Lambda, \Lambda^O)$ between the unitary Choi matrix Λ and a Pauli/Clifford shadow reconstructed Choi matrix Λ^O . Solid lines represent single repetition shadows with 512 unitaries on IonQ for Pauli/Clifford measurements, and in green, simulated $2n$ -Clifford circuits; dotted lines includes all data/repetitions collected; dot-dashed lines are purified. The corresponding purified MLE results are shown in blue/orange stars for Pauli/Clifford measurements, respectively. Right: Trace distance for various postprocessing for $n = 3$. The projection/postprocessing method is shown in parentheses above the bar.




FIG. 4. Left: Simulated normalized trace distance $T(\Lambda, \Lambda^O)$ scaling for 512 unitaries with a single repetition (top) and 50 repetitions (bottom) between the unitary Choi matrix Λ and a reconstruction via Pauli/Clifford ancilla based ShadowQPT or two-sided ShadowQPT. Ancilla results are nearly identical to the two-sided results with a single repetition, but with multiple repetitions ancilla simulations (squares) have lower trace distance than their two-sided counterparts (circles). Middle: Normalized trace distance between (IonQ measured) Pauli two-sided ShadowQPT and ancilla ShadowQPT with 512 unitaries and 50 repetitions for a three-qubit GHZ process. Right: Normalized trace distance between (IonQ measured) Clifford two-sided ShadowQPT and ancilla ShadowQPT with 512 unitaries and 50 repetitions for a two-qubit GHZ process. The ancilla Cliffords are decomposed into $k = 2$ as in Fig. 3.

of states. We compare to both the underlying unitary process and overlaps computed on the IonQ by constructing the trace circuit shown in Fig. 6(a); each circuit is measured with 1000 repetitions. Note that this latter approach requires additional quantum measurements for each target overlap, whereas the shadow tomography uses the same ShadowQPT data for the overlap of any output state.

In Fig. 6(b) we show the overlap prediction $w_i^O = \text{tr}[(\rho_i^{\text{in}T} \otimes \sigma_i) \Lambda^O]$ to the shadow reconstruction data Λ^O without postprocessing. For $n = 2, 3$, we are able to reasonably reconstruct the correct overlap, but for $n = 4$ there is not enough data to obtain the correct pure state behavior. Note that in each case the correct qualitative behavior is evident, despite predicting much smaller overlap values.

Then in Fig. 6(c) we use the the most efficient postprocessing approach to purify the Choi matrix. Purified Pauli and Clifford overlap predictions are now nearly identical with the expected noise-free values, even for the case of $n = 4$.

VI. NUMERICAL SIMULATION OF SHADOWQPT HAMILTONIAN LEARNING

Here we present simulation results with Hamiltonian learning. We study a 1D transverse field Ising model $H = \sum_i J_i X_i X_{i+1} + h_i Z_i$ with uniform random couplings $J_i, h_i \in [-1, 1)$ over 10 disorder realizations. When applying ShadowQPT, we apply the unitary e^{-itH} and sample N random Pauli measurements with no additional postprocessing. For the J_i terms we use $p_i = Z_i$, $q_i = [h_i, p_i] = -Y_i X_{i+1}$, and for the h_i terms $p_i = X_i$, $q_i = [h_i, p_i] = Y_i$.

Using the methods described in Sec. III, we plot the behavior of the average error in Fig. 7. In Fig. 7(a), we show the average absolute error $\langle |\tilde{c}_i - c_i| \rangle$ between the ShadowQPT learned couplings $\tilde{c}_i$ and the original couplings c_i ($c_i \equiv \{J, h\}$) with a simulation of $N = 100\,000$. Given a fixed t , the average error scales nearly independently of system size n [and should be bounded by $O(\log(n))$ scaling at large n]. For small t , while the first-order expansion of the time evolution operator is a good approximation, one needs many measurements to determine the couplings beyond statistical noise as the channel is very close to the identity. For large t , the higher-order terms generated by the exponential become important, and the learned couplings $\tilde{c}_i \approx c_i^{\text{renorm}}(t) = i \text{tr}(\rho_{\Lambda_i}(p_i) q_i) / t \not\approx c_i$. At intermediate t , we find a favorable regime where there are enough measurements to learn the random couplings to an average error of 10^{-1} .

In Fig. 7(b) (see the Appendix for further details) we more carefully analyze the errors in this process, particularly as we change the number of measurements and time. There are two sources of error in Hamiltonian learning using ShadowQPT. A systematic error occurs due to the linear approximation of the time evolution operator and is given by $\epsilon_s(t) \equiv \langle |c_i - c_i^{\text{renorm}}(t)| \rangle$ and scales as $O(t^2)$. Notice that this error does not depend on the number of measurements. There is also a statistical error coming from the stochastic nature of the ShadowQPT process, which scales as $\epsilon \propto 1/\sqrt{N}$ [see $t = 0.1$ in Fig. 7(b)] and empirically $\epsilon \propto 1/t$ (see Fig. 17). For a fixed t , we can therefore make the error better by increasing the number of measurements until we reach an error of $\epsilon_s(t)$. The optimal error, found when the systematic error is approxi-




FIG. 5. Normalized trace distance $T(\Lambda_B, \Lambda_B^O)$ of all subsystems of $k = 1$ qubit (a) and $k = 2$ qubits (b) between the unitary Choi matrix Λ_B and a Pauli/Clifford shadow reconstructed Choi matrix Λ_B^O . Solid lines represent single repetition shadows with 512 unitaries on IonQ for Pauli/Clifford measurements, and in green, simulated $2n$ -Clifford circuits; dot-dashed lines are $\mathcal{CP}$ -projected after the partial trace. The corresponding MLE results are shown in blue/orange stars for Pauli/Clifford measurements, respectively. Pauli MLE is done directly in the reduced problem space of k qubit reconstruction.

mately the same as the statistical error, is shown in purple diamonds in Fig. 7(b) and scales as $O(1/\epsilon^3)$. Despite the increase in scaling, the number of measurements required is overall fewer than from the $O(1/\epsilon^2)$ results of using a fixed time. As we also noted earlier in Sec. III, the systematic error can be further reduced by higher Taylor expansion and classical processing to achieve better overall scaling. Our study provides a prescription for choosing the number of measurements needed for a given value of ϵ with a weak dependence of n .

VII. CONCLUSION

In this work, we have developed classical shadow algorithms for quantum process tomography. We discuss the




FIG. 6. (a) Trace circuit measurement of the overlap $\text{tr}[\Lambda^O(\rho_i^{\text{in}T} \otimes \sigma_i)]$, where $\sigma = U(\theta)^\dagger |0\rangle\langle 0| U(\theta)$. We measure this circuit on IonQ to compare with a prediction from ShadowQPT. (b) Comparison of the overlap $w_i^O = \text{tr}[(\rho_i^{\text{in}T} \otimes \sigma_i) \Lambda^O]$ for Pauli and Clifford data on systems of size $n = 2, 3, 4$ for pairs of density matrices $[\rho^{\text{in}}, \sigma]$ formed from the state $[O_\rho^{\text{in}}|0\rangle, O_\sigma|0\rangle]$ for operators $[I, R_y(\theta)], [I, R_x(\theta)], [\otimes_j R_x(\phi_j), R_y(\theta)], [\otimes_j R_x(\phi_j), R_y(\theta)]$ (each column, respectively) over 50 different angles. For comparison, we include trace circuit measurements performed on the IonQ using 1000 repetitions per point; error bars are generated by using 10 batches of data. (c) Overlap predictions as in (b) but purifying Λ^O .




FIG. 7. (a) Hamiltonian learning simulation results for a 1D transverse field Ising model with n sites and random couplings between $[-1, 1]$. We average over 10 disorder realizations and use $N = 100\,000$ random Pauli measurements with no additional post-processing. The average error is given by average absolute error $\langle |\tilde{c}_i - c_i| \rangle$ to the original Hamiltonian coupling c_i . (b) Dependence on the number of Pauli measurements N vs average error of Hamiltonian learning both for fixed t and the optimal t (shown in purple). Notice at fixed t , as the error decreases it approaches the minimum error at that t , shown by the vertical dot-dashed line, coming from the systematic error caused by the linear approximation to the time evolution operator.

power of classical ShadowQPT in Theorems II.1 and II.2. Interestingly, ShadowQPT can be applied to low-weight inputs and outputs at a cost in quantum measurements which scales only logarithmically with N . To realize the ShadowQPT on near-term quantum devices, we study the effects of both random Pauli measurements and Clifford measurements. In addition, we explore different postprocessing techniques. We further benchmark our methods and achieve good performance for the unitary process, the nonunitary reduced process, and overlap estimation on IonQ quantum hardware up to $n = 4$ qubits. We find that our postprocessing techniques, particularly purification, result in the ability to simultaneously compute very accurate overlaps for any state. Our method not only provides a theoretical foundation, but also effectively

applies to NISQ quantum devices. Additionally, we show the equivalence between the ancilla-based ShadowQPT and a two-sided scheme, and we find both via simulations and measurements that the ancilla scheme under multiple repetitions achieves better performance.

We additionally develop Hamiltonian learning using ShadowQPT, and we discuss its scaling in Corollary [cor1]. Then using numerical simulations we show that there is a logarithmic dependence on system size, as expected, and using intermediate time propagators can be advantageous for minimizing the total number of samples needed for a given error threshold. For future exploration, one can consider integrating our methods with recently developed Hamiltonian driven shadow methods [35,40] or full Clifford circuit decomposition [39], which could allow for more flexible choices for unitaries as well as applications to different experimental platforms. ShadowQPT can further study the dynamics of observables and correlation functions, which is explored in the concurrent work of Ref. [51], which also studied the use of classical shadows in QPT. We anticipate that ShadowQPT will be an important tool for future exploration and validation of quantum devices.

ACKNOWLEDGMENTS

We acknowledge support from the NSF Quantum Leap Challenge Institute for Hybrid Quantum Architectures and Networks (NSF Award No. 2016136). This work is supported by the National Science Foundation under Cooperative Agreement PHY-2019786 (the NSF AI Institute for Artificial Intelligence and Fundamental Interactions [52]). This material is based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Co-design Center for Quantum Advantage (C2QA) under Contract No. DE-SC0012704. A portion of this work was completed during the 2021 QHack hackathon. We gratefully acknowledge funding for quantum computing machine-time from AWS and Xanadu. R.L. and D.L. thank our two hackathon team members E. Chertkov and A. Khan, and acknowledge helpful discussions from Hongye Hu, Nicholas Rubin, Juan Carrasquilla, and Yizhuang You, as well as insightful comments on the shadow schemes from Daniel K. Mark, Minh Tran, Tianci Zhou, and Soonwon Choi. The Flatiron Institute is a division of the Simons Foundation.

APPENDIX: CLASSICAL SHADOWS OF QUANTUM PROCESS TOMOGRAPHY ON NEAR-TERM QUANTUM COMPUTERS

1. Proofs for shadow process tomography

We first review the formal version of a classical shadow quantum state tomography theorem in Ref. [27].

Theorem A.1. Fix a measurement primitive U , a collection $\{O_i\}_{i=1,\dots,M}$ of $2^n \times 2^n$ Hermitian matrices and accuracy parameters $\epsilon, \delta \in (0, 1)$. Set $K = 2 \log(2M/\delta)$ and $N = \frac{34}{\epsilon^2} \max_i \|O_i - \frac{\text{tr}(O_i)}{2} \mathbb{I}\|_{\text{shadow}}^2$, where $\|\cdot\|_{\text{shadow}}$ is the shadow norm defined with the inverse channel $\mathcal{M}^{-1}$ induced by the

measurement primitive as follows:

$$\|O\|_{\text{shadow}} = \max_{\sigma} \left(\mathbb{E}_{U \sim \mathcal{U}} \sum_{b \in \{0,1\}^n} \langle b|U\sigma U^\dagger|b \rangle \langle b|U\mathcal{M}^{-1} \right. \\ \left. \times (O)U^\dagger|b \rangle^2 \right)^{1/2}. \quad (\text{A1})$$

Then, a collection of NK independent classical shadows allows for accurately predicting all features via median of means prediction $\hat{o}_i$:

$$|\hat{o}_i(N, K) - \text{tr}(O_i \rho)| \leq \epsilon \quad (\text{A2})$$

for all $1 \leq i \leq M$ with probability at least $1 - \delta$.

In particular, for random global Clifford measurement primitive, $\|O - \frac{\text{tr}(O)}{2}\mathbb{I}\|_{\text{shadow}}^2 \leq 3 \text{tr}(O^2)$. For random single-qubit Clifford measurement primitive, $\|O - \frac{\text{tr}(O)}{2}\mathbb{I}\|_{\text{shadow}}^2 \leq 4^k \|O\|_{\infty}^2$, where k is the locality of the operator O and $\|\cdot\|_{\infty}$ is the spectral norm. Furthermore, if O is a single k -local Pauli observable, $\|O - \frac{\text{tr}(O)}{2}\mathbb{I}\|_{\text{shadow}}^2 \leq 3^k$.

We now combine the above theorem and the Jamiołkowski process-state isomorphism [25] to reach the following two theorems for shadow quantum process tomography.

Theorem A.2. For an n -qubit quantum process ρ_{Λ} and $\epsilon, \delta \in (0, 1)$, given a set of density matrix pairs $\{(\rho_1^{\text{in}}, \sigma_1), \dots, (\rho_M^{\text{in}}, \sigma_M)\}$, the number of measurements N that suffice to predict the overlaps $\text{tr}(\rho_{\Lambda}(\rho_i^{\text{in}T} \otimes \sigma_i))$ for any i up to error ϵ with probability $1 - \delta$ is

$$N = \frac{68}{\epsilon^2} \log(2M/\delta) \max_i \left\| O_i - \frac{\text{tr}(O_i)}{2} \mathbb{I} \right\|_{\text{shadow}}^2, \quad (\text{A3})$$

where $O_i = \rho_i^{\text{in}T} \otimes \sigma_i$. For random global Clifford measurement, $N = \frac{68}{\epsilon^2} \log(2M/\delta) \max_i \text{tr}(O_i^2)$. If ρ_i^{in} and σ_i are all pure states, then $\text{tr}(O_i^2) = 1$. For random single-qubit Clifford measurement, $N = \frac{68}{\epsilon^2} \log(2M/\delta) \max_i 4^k \|O_i\|_{\infty}^2$, where O_i acts nontrivially on k_i -qubits. If ρ_i^{in} and σ_i are all pure states, then $\|O_i\|_{\infty}^2 = 1$.

Proof. Under the Jamiołkowski process-state isomorphism [25],

$$\text{tr}(\rho_{\Lambda}(\rho_i^{\text{in}T} \otimes \sigma_i)) = \text{tr}(\rho_{\Lambda}(O_i)), \quad (\text{A4})$$

The result follows directly from applying Theorem A.1 with M measurements. For random global Clifford measurement and pure states ρ_i^{in} and σ_i , $O_i = \rho_i^{\text{in}T} \otimes \sigma_i$ is also a pure state of rank 1 so that $\text{tr}(O_i^2) = 1$. This implies that $\|O_i - \frac{\text{tr}(O_i)}{2}\mathbb{I}\|_{\text{shadow}}^2 \leq 3 \text{tr}(O_i^2) = 3$. For random single-qubit Clifford measurement and pure states ρ_i^{in} and σ_i , $O_i = \rho_i^{\text{in}T} \otimes \sigma_i$ is also a pure state so that $\|O_i\|_{\infty}^2 = 1$. ■

Theorem A.3. For an n -qubit process Choi matrix ρ and $\epsilon, \delta \in (0, 1)$, the following statements hold:

(i) The number of measurements N that suffice to simultaneously predict any reduced k -qubit process Choi matrix ρ_k up to error ϵ in the Frobenius norm and probability $1 - \delta$ is

$$N = \frac{68}{\epsilon^2} 4^k \log(2(8n)^{2k}/\delta) \max_i \left\| O_i^{(2k)} - \frac{\text{tr}(O_i^{(2k)})}{2} \mathbb{I} \right\|_{\text{shadow}}^2, \quad (\text{A5})$$

where $O_i^{(2k)}$ is any $2k$ -qubit Pauli observable.




FIG. 8. Normalized trace distance $T(\Lambda, \Lambda^O)$ between the unitary Choi matrix Λ and a Pauli/Clifford shadow reconstructed Choi matrix Λ^O for the $n = 3$ GHZ process. The construction is either via mean, median of means via randomizing and batching individual shadows (with many derived from the same unitary), or median of means for each unique unitary (which has $50/\approx 70$ repetitions each). We include the effects of $\mathcal{CP}$ projection on the resulting Choi matrices for comparison.

In particular, for random global Clifford measurement, $N = \frac{204}{\epsilon^2} 4^{n+k} \log(2(8n)^{2k}/\delta)$ and for Pauli-6 POVM measurement, $N = \frac{68}{\epsilon^2} 36^k \log(2(8n)^{2k}/\delta) \max_i \|O_i^{(2k)}\|_{\infty}^2$.

(ii) The number of Pauli-6 POVM measurements N that suffice to simultaneously predict any reduced k -qubit process Choi matrix ρ_k up to error ϵ in the trace norm and probability $1 - \delta$ is

$$N = \frac{8}{3} \frac{144^k}{\epsilon^2} \log((24n)^{2k}/\delta). \quad (\text{A6})$$

Proof. (i) According to the Jamiołkowski process-state isomorphism [25], a quantum process on an n -qubit system can be represented by a density matrix ρ on a $2n$ -qubit system. Any reduced k -qubit process matrix ρ_k could be expressed as

$$\rho_k = \frac{1}{4^k} \sum_{\alpha} \alpha_i O_i^{(2k)}, \quad (\text{A7})$$

where $\alpha_i = \text{tr}(\rho_k O_i^{(2k)}) = \text{tr}(\rho O_i^{(2k)} \otimes \mathbb{I}_{2n-2k})$.

Denote the estimated k -qubit reduced density matrix from the shadow tomography as $\hat{\rho}_k$ and $\hat{\rho}_k = \frac{1}{4^k} \sum_{\alpha} \hat{\alpha}_i O_i^{(2k)}$, where $\hat{\alpha}_i$ is the shadow estimation of observable $O_i^{(2k)}$. It follows that

$$\langle \hat{\rho}_k - \rho_k, \hat{\rho}_k - \rho_k \rangle_F = \left\langle \frac{1}{4^k} \sum_{\alpha} (\hat{\alpha}_j - \alpha) O_j^{(2k)}, \right. \\ \left. \frac{1}{4^k} \sum_{\alpha} (\hat{\alpha}_i - \alpha) O_i^{(2k)} \right\rangle_F \\ = \frac{1}{4^k} \sum_{\alpha} (\hat{\alpha}_i - \alpha_i)^2, \quad (\text{A8})$$

where $\langle \cdot, \cdot \rangle_F$ is the Frobenius inner product.




FIG. 9. (a) Choi matrix visualization using Pauli (left two) and Clifford (right two) shadows, separated into real and imaginary components. (b) Purification of matrices in (a).

According to Theorem A.1, with N given in Eq. (5), we have

$$|\hat{\alpha}_i - \alpha_i| = |\hat{\alpha}_i - \text{tr}(\rho O_i^{(2k)} \otimes \mathbb{I}_{2n-2k})| \leq \frac{\epsilon}{2^k} \quad \forall \quad i. \quad (\text{A9})$$

This implies that $\langle \hat{\rho}_k - \rho_k, \hat{\rho}_k - \rho_k \rangle_F \leq \frac{16^k \epsilon^2}{4^k 4^k} = \epsilon^2$, so that $\|\hat{\rho}_k - \rho_k\|_F \leq \epsilon$, where $\|\cdot\|_F$ is the Frobenius norm. For random global Clifford measurement, the result follows from $\|O - \frac{\text{tr}(O)}{2} \mathbb{I}\|_{\text{shadow}}^2 \leq 3 \text{tr}(O^2) = 3 \times 4^n$ for $O = O_i^{(2k)}$. For Pauli-6 POVM measurement, the result follows from $\|O - \frac{\text{tr}(O)}{2} \mathbb{I}\|_{\text{shadow}}^2 \leq 9^k$ for $O = O_i^{(2k)}$ [36].

(ii) For a given $\epsilon, \delta \in (0, 1)$, the Pauli-6 POVM measurement of size $N = \frac{8}{3\epsilon^2} 12^r (\log(p^r 12^r / \delta))$ is sufficient to predict all reduced density matrices on subsystem size $r \leq p$ on a p -qubit system with trace norm error ϵ and probability at least $1 - \delta$ according to Lemma 1 in Ref. [41]. With the Jamiołkowski process-state isomorphism [25], a k -qubit pro-

cess Choi matrix is equivalent to a $2k$ -qubit density matrix on a $2n$ -qubit system. The result follows from choosing $r = 2k$ and $p = 2n$, which implies that with probability $1 - \delta$, $N = \frac{8}{3} \frac{144^k}{\epsilon^2} \log((24n)^{2k} / \delta)$ number of Pauli-6 POVM measurements is sufficient to simultaneously predict any reduced k -qubit process Choi matrix ρ_k up to error ϵ in the trace norm, i.e., any estimator $\hat{\rho}_k$ from the shadow measurement is closer to the exact ρ_k by $\|\hat{\rho}_k - \rho_k\|_1 < \epsilon$. ■

2. Equivalence of two-sided scheme with the top-bottom factorized ancilla-based scheme

We would like to show that the two-sided scheme is equivalent to the top-bottom factorized ancilla-based scheme. To be precise, consider a top-bottom factorized ancilla-based scheme with unitaries $\{U_j\}$ acting on the top half and unitaries $\{U_i\}$ acting on the bottom half, where U_j, U_i are sampled

uniformly from a set G . Up to permutation, the reconstruction formula for the Choi matrix is the following:

$$\rho_{\Lambda}^a = \mathbb{E}_{U_j \sim G} \mathbb{E}_{U_i \sim G} \mathbb{E}_{b_j, b_i \sim P(i, j)} \mathcal{M}_n^{-1}(U_i^{\dagger} |b_i\rangle \langle b_i| U_i) \otimes \mathcal{M}_n^{-1}(U_j^{\dagger} |b_j\rangle \langle b_j| U_j), \quad (\text{A10})$$

where b_j, b_i are measurement outcome with respect to U_j, U_i in the computational basis, $P(i, j) = \text{tr}(\rho_{\Lambda}(U_i^{\dagger} |b_i\rangle \langle b_i| U_i \otimes U_j^{\dagger} |b_j\rangle \langle b_j| U_j))$. $\mathcal{M}_n^{-1}$ is the inverse channel defined by $\mathcal{M}_n^{-1}(X) = (2^n + 1)X - \mathbb{I}$ for n -qubit Clifford unitaries or $\mathcal{M}_n^{-1}(X) = \otimes_i M_k^{-1}(X)$ for Pauli ($k = 1$) for smaller k qubit Cliffords.

Meanwhile, consider a two-sided scheme where the input state is $|0\rangle$, followed by $\{U_i^{\dagger}\}$, process channel $\mathcal{E}$, $\{U_j\}$, and finally measurement in the computational basis shown in Fig. 1(b). U_j, U_i are also sampled uniformly from group G . The reconstruction formula for the Choi matrix is given as follows:

$$\rho_{\Lambda}^t = \mathbb{E}_{U_j \sim G} \mathbb{E}_{U_i \sim G} \mathbb{E}_{b_j \sim P(j|U_i^{\dagger}|0\rangle)} \mathcal{M}_n^{-1}((U_i^{\dagger}|0\rangle \langle 0| U_i)^T) \otimes \mathcal{M}_n^{-1}(U_j^{\dagger} |b_j\rangle \langle b_j| U_j), \quad (\text{A11})$$

where $P(j|U_i^{\dagger}|0\rangle) = \text{tr}(\mathcal{E}(U_i^{\dagger}|0\rangle \langle 0| U_i) U_j^{\dagger} |b_j\rangle \langle b_j| U_j) = 2^n \text{tr}(\rho_{\Lambda}((U_i^{\dagger}|0\rangle \langle 0| U_i)^T \otimes U_j^{\dagger} |b_j\rangle \langle b_j| U_j))$. The last line comes from the Choi matrix identification. For any set G that is right invariant with respect to any Pauli- X matrix such that $GX = G$, we have

$$\rho_{\Lambda}^t = \mathbb{E}_{U_j \sim G} \mathbb{E}_{U_i \sim G} \mathbb{E}_{b_i \sim \frac{1}{2^n}} \mathbb{E}_{b_j \sim P(j|U_i^{\dagger}|b_i\rangle)} \mathcal{M}_n^{-1}((U_i^{\dagger}|b_i\rangle \langle b_i| U_i)^T) \otimes \mathcal{M}_n^{-1}(U_j^{\dagger} |b_j\rangle \langle b_j| U_j). \quad (\text{A12})$$

This is true because $|b_i\rangle = \otimes_{\{\ell\}} X|0\rangle$, where $\{\ell\}$ are the locations where the measurement outcome is 1, and hence for any $U_i^{\dagger}|0\rangle$ there exists U_i' from G such that $U_i'|b_i\rangle = U_i|0\rangle$. Notice that this condition is true for random global Clifford measurement, random single-qubit Clifford measurement, as well as Pauli-6 measurement. It follows that

$$\rho_{\Lambda}^t = \mathbb{E}_{U_j \sim G} \mathbb{E}_{U_i \sim G} \mathbb{E}_{b_j, b_i \sim P^t(i, j)} \mathcal{M}_n^{-1}(U_i^{\dagger} |b_i\rangle \langle b_i| U_i)^T \otimes \mathcal{M}_n^{-1}(U_j^{\dagger} |b_j\rangle \langle b_j| U_j), \quad (\text{A13})$$

where $P^t(i, j) = \frac{1}{2^n} P(j|U_i^{\dagger}|b_i\rangle) = \text{tr}(\rho_{\Lambda}((U_i^{\dagger}|b_i\rangle \langle b_i| U_i)^T \otimes U_j^{\dagger} |b_j\rangle \langle b_j| U_j))$. For a Pauli-6 POVM, we have $(U_i^{\dagger}|b_i\rangle \langle b_i| U_i)^T = U_i^{\dagger}|b_i\rangle \langle b_i| U_i$, which implies $\rho_{\Lambda}^t = \rho_{\Lambda}^a$. For Clifford group G and any fixed $|b\rangle$, we have $\{U_i^{\dagger}|b\rangle \langle b| U_i | U_i \in G\} = \{U_i^{\dagger}|b\rangle \langle b| U_i^T | U_i \in G\}$, which also implies $\rho_{\Lambda}^t = \rho_{\Lambda}^a$.




FIG. 10. Normalized trace distance $T(\Lambda, \Lambda^O)$ between the unitary Choi matrix Λ and a Pauli/Clifford shadow reconstructed Choi matrix Λ^O . Included are two projections of the Λ^O into the space of positive-semidefinite (PSD) matrices, the iterative maximum-likelihood estimator (MLE) method of Ref. [10], and the eigenvalue rescaling method of Ref. [12], as well as the corresponding purification.

We remark that the above argument holds for a single repetition measurement. For multiple repetition measurements with respect to one U_i , the equivalence breaks because multiple $|b_i\rangle$ are generated with respect to one U_i in the ancilla-based scheme while the two-sided scheme always has input $|0\rangle$ fixed for one U_i .

3. Details on shadow measurements

We present the number of random unitary circuits used for both Clifford and Pauli measurements. For those with unitary entries with *, all possible unitaries have been used:

Qubits	Type	Unitaries	Repetitions/ Unitary (Total)
2	Pauli	81*	≈632 (51 200)
	Clifford	1024	50 (51 200)
3	Pauli	729*	≈70 (51 200)
	Clifford	1024	50 (51 200)
4	Pauli	1024	50 (51 200)
	Clifford	1024	50 (51 200)

4. Comparison between median-of-means and mean

In Ref. [27], the authors suggest a median-of-means procedure for the classical shadows, or more precisely taking the median of K averaged shadows. Then in Ref. [32] they show that most of their results are unaffected by using the direct mean over a median-of-means.

To understand the role of the procedure in ShadowQPT, we first distinguish between two median-of-means procedures. The first is a median on the “shadow” level: take each classical shadow, many of which are from the same unitary, and perform the median-of-means procedure using this data set. An alternative procedure is to first average the repetitions




FIG. 11. Normalized trace distance $T(\Lambda_B, \Lambda_B^O)$ of the subsystem of the first qubit ($k = 1$) (a) and (b) the remaining $n - 1$ qubits ($k = n - 1$) between the unitary Choi matrix Λ_B and a Pauli/Clifford shadow reconstructed Choi matrix Λ_B^O .

corresponding to each unitary, and then obtain the median by using sets of K averaged unitaries.

In Fig. 8 is shown the results of using the mean and the two median-of-mean procedures for $n = 3$ data, with a batch size of $K = 23$. While using a median-of-means at the shadow level does show a slight improvement for both Pauli and Clifford measurements without projection, after projection we find that there is little difference between the mean and median-of-means results. Given the weak dependence on the differences after postprocessing and the potential introduction of a tunable variable (K), we chose to utilize the mean for our computing results of the classical shadows.

5. 3D choi matrix plots

In Fig. 9 we present a 3D representation of $n = 2, 3$ in Fig. 2. The axes are shown enumerated in binary for the Choi matrices 4^n index values.

6. Postprocessing dependence

To show the dependence on the postprocessing method, Fig. 10 shows the normalized trace distance between MLE reconstruction, $\mathcal{CP}$ -projection, and purification, while Fig. 11 shows the same for a subsystem of $k = 1$. Overall the results are roughly consistent across system size n , in that purified shadows are competitive with purification of (Clifford) MLE results. Just considering $\mathcal{CP}$ projection and




FIG. 12. (a) Normalized trace distance $T(\Lambda, \Lambda^O)$ between the unitary Choi matrix Λ and various simulated reconstruction methods. Simulated data include a prefix of “Sim.” The vertical line represents saturation of all Pauli strings. (b) Normalized trace distance $T(\Lambda, \Lambda^O)$ between the unitary Choi matrix Λ and a Clifford shadow reconstruction from resampled data (markers) on top of simulated data (dotted lines with shading). A star marks the full IonQ data set; error bars represent an average of five resampling trials. Resampled IonQ data consist of equal fixed and random Clifford orderings.




FIG. 13. (a) Left: Normalized Frobenius norm $\|\Lambda - \Lambda^O\|_F/2^n$ between the unitary Choi matrix Λ and a Pauli/Clifford shadow reconstructed Choi matrix Λ^O . Solid lines represent single repetition shadows with 512 unitaries on IonQ for Pauli/Clifford measurements, and in green, simulated $2n$ -Clifford circuits; dotted lines includes all data/repetitions collected; dot-dashed lines are purified. The corresponding purified MLE results are shown in blue/orange stars for Pauli/Clifford measurements, respectively. Right: Frobenius norm distance for various postprocessing for $n=3$. The projection/postprocessing method is shown in parentheses above the bar. (b) $k=1$ qubit and (c) $k=2$ qubits reduced process reconstruction between the unitary Choi matrix Λ_B and a Pauli/Clifford shadow reconstructed Choi matrix Λ_B^O using the Frobenius norm. Solid lines represent single repetition shadows with 512 unitaries (including simulated $2n$ -Clifford circuits), dotted lines includes all data/repetitions, and dot-dashed lines are $\mathcal{CP}$ -projected after the partial trace. The corresponding MLE results are shown in blue/orange stars for Pauli/Clifford measurements, respectively. Pauli MLE is done directly in the reduced problem space of k qubit reconstruction.

standard MLE, Pauli measurements have the minimal trace distance.

7. Repetition dependence on trace distance

Next we compare the dependence on the trace distance with the number of unitaries and classical shadows for $n=3$ in Fig. 12.

Second, we resample the Clifford measurements from the IonQ alongside simulated data, computing the normalized trace distance in Fig. 12(b). We find that the measured data track the simulation very closely for a number of chosen unitary Clifford measurements. Although we predict performance could be increased by increasing the number of

repetitions/circuits [i.e., we had not reached instability in trace distance seen in the low unitary counts, shown by the blue and orange lines of Fig. 12(b)], the number of individual circuits and repetitions/circuits chosen appears to be reasonable in the study of our method. Although there is an instability in the regime where the number of repetitions is greater than the number of unitary, the trace distance remains improved compared to a single measurement.

8. Distance using the frobenius norm

We provide additional plots of Figs. 3, 4, 5(a), and 5(b) using a Frobenius norm distance shown in Figs. 13 and 14. We find nearly the same qualitative behavior for both distance




FIG. 14. Left: Simulated Frobenius norm $\|\Lambda - \Lambda_O\|_F / 2^n$ scaling for 512 unitaries with a single repetition (top) and 50 repetitions (bottom) between the unitary Choi matrix Λ and a reconstruction via Pauli or Clifford ancilla based ShadowQPT or two-sided ShadowQPT. “Sim-Clifford/Pauli” corresponds to single-qubit Clifford measurements ($k = 1$) for U_L and Pauli measurements on the right, which is nearly identical to Pauli two-sided measurements. Ancilla results are nearly identical to the two-sided results with a single repetition, but with multiple repetitions ancilla simulations (squares) have a lower trace distance than their two-sided counterparts (circles). Middle: Normalized trace distance between (measured) Pauli two-sided ShadowQPT and ancilla ShadowQPT with 512 unitaries and 50 repetitions for a three-qubit GHZ process. Right: Normalized Frobenius norm between (measured) Clifford two-sided ShadowQPT and ancilla ShadowQPT with 512 unitaries and 50 repetitions for a two-qubit GHZ process. The ancilla Cliffords are decomposed into $k = 2$ as in Fig. 13.

measures. This is unsurprising as, at least for the Pauli case, there are similar scaling bounds of $\log(n)$ for fixed k (see Theorem II.1).

9. Details on overlap estimation

To illustrate Theorem II.2, we compute $\text{tr}[\Lambda'(\rho^{\text{in}T} \otimes \sigma(\theta))]$ for several sets of ρ and $\sigma(\theta)$. The input to the channel, ρ^{in} can be represented by an initial state of which there are

three categories: $|0\rangle^{\otimes n}$, $|+i\rangle^{\otimes n}$, and $\otimes_j R_x(\phi_j)|0\rangle^{\otimes n}$ states.² The first two states are already measured within the Choi matrix, and they effectively measure the reconstruction of the channel process. After running the circuit in Fig. 6(a), we measure $|0\rangle^{\otimes n}$ or $P(|0\rangle^{\otimes n})$ to obtain the value of the overlap.

²where $R_x(\phi_j)$ acts on the j th qubit and $\phi_j = \{0.1717, 0.1234, 0.9876, 0.888\}$




FIG. 15. Average overlap error $\langle \Delta \rangle = \langle |w_i^O - w_i| \rangle_i$ for Pauli reconstruction data (left), Clifford data (middle), and simulated Clifford data (right) for systems of size $n = 2, 3, 4$. We average over three different ρ^{in} and four different σ each with 51 different angles. For comparison, we include trace circuit measurements (Tr circuit) performed on the IonQ using the data shown in Fig. 6(b).




FIG. 16. Normalized purity measurement $\text{tr}[\Lambda^2]/4^n$ of the $\mathcal{CP}$ -projected Pauli/Clifford measured Choi matrices as well as the MLE reconstructions shown as a star. We use the BBA algorithm run directly on the IonQ with 1000 repetitions shown in diamonds to compute $\text{tr}[\rho_{\text{GHZ}_n}^2]$, where $\rho_{\text{GHZ}_n} = \mathcal{E}(|0\rangle\langle 0|^{\otimes n})$; error bars are done by batching the measurements into 10 sets.

We only display $|0\rangle^{\otimes n}$ and $\otimes_j R_x(\phi_j)|0\rangle^{\otimes n}$ states in the main work.

We compute the average overlap error in Fig. 15 compared to the trace circuit measurements performed on the IonQ device. We use 51 pairs for each ρ^{in} for the shadow reconstructed data iterating uniformly $\theta \in [0, 2\pi]$ and $\approx 5 - 8$ pairs for each ρ^{in} . Error bars represent the standard error over the overlap data set. We see that our reconstructed shadow data have similar trends to the trace circuit measurement, with a significant range of values. With purification, we can significantly decrease the error to the noiseless case even compared to the trace circuit, particularly in the $n = 4$ case.

Purity of the channel and GHZ state can also be calculated. The Bell-basis algorithm (BBA) swap test [53] provides an algorithm to determine the overlap between two general quantum states. We compute the overlap of two GHZ states $\text{tr}[\rho_{\text{GHZ}_n}^2]$, where $\rho_{\text{GHZ}_n} = \mathcal{E}(|0\rangle\langle 0|^{\otimes n})$, using the BBA algorithm, alongside computing the shadow reconstructed $\text{tr}[\Lambda^2]/4^n$ in Fig. 16. When measuring the purity, we find a much higher value than our shadow reconstructions. A noiseless quantum computer should produce unit purity for the full GHZ process. These large measurements further motivate our choice to use purification as a postprocessing technique. Note that purity has a strong dependence on postprocessing techniques.

10. Hamiltonian learning optimal error fits

In Fig. 17 we show for a fixed number of Pauli measurements (shadows) N the average error versus t of the unitary operator $\exp(-iHt)$. These points have been simulated and averaged over the same 10 trials as in Fig. 7. For a given value of t , because of the linear approximation to $\exp(-iHt)$, the error of which is shown in blue, the error grows as $O(t^2)$ and is the minimum achievable error for a given t . ShadowQPT reconstructed couplings $\tilde{c}_i$ are inherently measuring c_i^{renorm} , however, with an error as $O(1/t)$ shown in green. Thus as these two approach, the true absolute error between $\tilde{c}_i$ and c_i is shown in orange, which mostly follows either the linear approximation error or the renormalized error except near their crossing. The minimum error can then be found for a given t , shown in a dashed gray line. We additionally fit this line for $N = 1000, 1\,000\,000$, which are not shown.




FIG. 17. Hamiltonian learning simulation results for a 1D transverse field Ising model with n sites and random couplings between $[-1, 1]$. We average over 10 disorder realizations and use $N = 100, 10\,000$, and $100\,000$ random Pauli measurements, respectively, with no additional postprocessing. The average error is given by average absolute error $\langle |b_i - c_i| \rangle$ to the original Hamiltonian coupling c_i , where b_i is either c_i^{renorm} (the renormalized couplings), $\tilde{c}_i$ (the ShadowQPT reconstructed couplings), or $\tilde{c}_i + c_i - c_i^{\text{renorm}}$, or the error between ShadowQPT and the renormalized couplings. A dashed line shows the optimal error where the linear approximation error is approximately the same as the renormalized error.

- [1] D. Bouwmeester, J.-W. Pan, K. Mattle, M. Eibl, H. Weinfurter, and A. Zeilinger, Experimental quantum teleportation, *Nature (London)* **390**, 575 (1997).
- [2] V. Giovannetti, S. Lloyd, and L. Maccone, Advances in quantum metrology, *Nat. Photon.* **5**, 222 (2011).
- [3] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature (London)* **574**, 505 (2019).
- [4] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You, Z. Wang *et al.*, Quantum computational advantage using photons, *Science* **370**, 1460 (2020).
- [5] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho *et al.*, Quantum phases of matter on a 256-atom programmable quantum simulator, *Nature (London)* **595**, 227 (2021).
- [6] I. L. Chuang and M. A. Nielsen, Prescription for experimental determination of the dynamics of a quantum black box, *J. Mod. Opt.* **44**, 2455 (1997).
- [7] J. F. Poyatos, J. I. Cirac, and P. Zoller, Complete characterization of a quantum process: The two-bit quantum gate, *Phys. Rev. Lett.* **78**, 390 (1997).
- [8] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White, Measurement of qubits, *Phys. Rev. A* **64**, 052312 (2001).
- [9] F. Bouchard, F. Hufnagel, D. Koutný, A. Abbas, A. Sit, K. Heshami, R. Fickler, and E. Karimi, Quantum process tomography of a high-dimensional quantum communication channel, *Quantum* **3**, 138 (2019).
- [10] A. I. Lvovsky, Iterative maximum-likelihood reconstruction in quantum homodyne tomography, *J. Opt. B* **6**, S556 (2004).
- [11] R. Blume-Kohout, Hedged maximum likelihood quantum state estimation, *Phys. Rev. Lett.* **105**, 200504 (2010).
- [12] J. A. Smolin, J. M. Gambetta, and G. Smith, Efficient method for computing the maximum-likelihood quantum state from measurements with additive Gaussian noise, *Phys. Rev. Lett.* **108**, 070502 (2012).
- [13] C. Granade, C. Ferrie, and S. T. Flammia, Practical adaptive quantum tomography, *New J. Phys.* **19**, 113017 (2017).
- [14] M. Christandl, R. König, and R. Renner, Postselection technique for quantum channels with applications to quantum cryptography, *Phys. Rev. Lett.* **102**, 020504 (2009).
- [15] R. Blume-Kohout, Optimal, reliable estimation of quantum states, *New J. Phys.* **12**, 043034 (2010).
- [16] A. V. Rodionov, A. Veitia, R. Barends, J. Kelly, D. Sank, J. Wenner, J. M. Martinis, R. L. Kosut, and A. N. Korotkov, Compressed sensing quantum process tomography for superconducting quantum gates, *Phys. Rev. B* **90**, 144504 (2014).
- [17] G. Torlai, C. J. Wood, A. Acharya, G. Carleo, J. Carrasquilla, and L. Aolita, Quantum process tomography with unsupervised learning and tensor networks, *Nat. Commun.* **14**, 2858 (2023).
- [18] G. C. Knee, E. Bolduc, J. Leach, and E. M. Gauger, Quantum process tomography via completely positive and trace-preserving projection, *Phys. Rev. A* **98**, 062336 (2018).
- [19] T. Surawy-Stepney, J. Kahn, R. Kueng, and M. Guta, Projected least-squares quantum process tomography, *Quantum* **6**, 844 (2022).
- [20] G. Zhang, M. Huo, and Y. Li, Quantum operation of fermionic systems and process tomography using majorana fermion gates (2021), [arXiv:2102.00620](https://arxiv.org/abs/2102.00620) [quant-ph]
- [21] G. A. L. White, F. A. Pollock, L. C. L. Hollenberg, K. Modi, and C. D. Hill, Non-markovian quantum process tomography, *PRX Quantum* **3**, 020344 (2022).
- [22] L. C. G. Govia, G. J. Ribeill, D. Ristè, M. Ware, and H. Krovi, Bootstrapping quantum process tomography via a perturbative ansatz, *Nat. Commun.* **11**, 1084 (2020).
- [23] S. Xue, Y. Liu, Y. Wang, P. Zhu, C. Guo, and J. Wu, Variational quantum process tomography of unitaries, *Phys. Rev. A* **105**, 032427 (2022).
- [24] E. Onorati, T. Kohler, and T. Cubitt, Fitting quantum noise models to tomography data, *Quantum* **7**, 1197 (2023).
- [25] A. Jamiołkowski, Linear transformations which preserve trace and positive semidefiniteness of operators, *Rep. Math. Phys.* **3**, 275 (1972).
- [26] M.-D. Choi, Completely positive linear maps on complex matrices, *Lin. Alg. Appl.* **10**, 285 (1975).
- [27] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [28] C. Hadfield, Adaptive pauli shadows for energy estimation (2021), [arXiv:2105.12207](https://arxiv.org/abs/2105.12207) [quant-ph].
- [29] S. Chen, W. Yu, P. Zeng, and S. T. Flammia, Robust shadow estimation, *PRX Quantum* **2**, 030348 (2021).
- [30] D. E. Koh and S. Grewal, Classical shadows with noise, *Quantum* **6**, 776 (2022).
- [31] S. Hillmich, C. Hadfield, R. Raymond, A. Mezzacapo, and R. Wille, Decision diagrams for quantum measurements with shallow circuits, in *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, Piscataway, NJ, 2021), pp. 24–34.
- [32] G. I. Struchalin, Y. A. Zagorovskii, E. V. Kovlakov, S. S. Straupe, and S. P. Kulik, Experimental estimation of quantum state properties from classical shadows, *PRX Quantum* **2**, 010307 (2021).
- [33] T. Zhang, J. Sun, X.-X. Fang, X.-M. Zhang, X. Yuan, and H. Lu, Experimental quantum state measurement with classical shadows, *Phys. Rev. Lett.* **127**, 200501 (2021).
- [34] A. Zhao, N. C. Rubin, and A. Miyake, Fermionic partial tomography via classical shadows, *Phys. Rev. Lett.* **127**, 110504 (2021).
- [35] H.-Y. Hu and Y.-Z. You, Hamiltonian-driven shadow tomography of quantum states, *Phys. Rev. Res.* **4**, 013054 (2022).
- [36] A. Acharya, S. Saha, and A. M. Sengupta, Shadow tomography based on informationally complete positive operator-valued measure, *Phys. Rev. A* **104**, 052418 (2021).
- [37] C. Rouzé and D. S. França, Learning quantum many-body systems from a few copies (2021), [arXiv:2107.03333](https://arxiv.org/abs/2107.03333) [quant-ph].
- [38] C. Hadfield, S. Bravyi, R. Raymond, and A. Mezzacapo, Measurements of quantum hamiltonians with locally-biased classical shadows, *Commun. Math. Phys.* **391**, 951 (2022).
- [39] W. J. Huggins, B. A. O’Gorman, N. C. Rubin, D. R. Reichman, R. Babbush, and J. Lee, Unbiasing fermionic quantum Monte

- Carlo with a quantum computer, *Nature (London)* **603**, 416 (2022).
- [40] H.-Y. Hu, S. Choi, and Y.-Z. You, Classical shadow tomography with locally scrambled quantum dynamics, *Phys. Rev. Res.* **5**, 023027 (2023).
- [41] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, Provably efficient machine learning for quantum many-body problems, *Science* **377**, eabk3333 (2022).
- [42] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, *Phys. Rev. A* **70**, 052328 (2004).
- [43] Z. Li, L. Zou, and T. H. Hsieh, Hamiltonian tomography via quantum quench, *Phys. Rev. Lett.* **124**, 160502 (2020).
- [44] J. Haah, R. Kothari, and E. Tang, Optimal learning of quantum Hamiltonians from high-temperature Gibbs states, *arXiv:2108.04842* (2021).
- [45] D. M. Greenberger, M. A. Horne, and A. Zeilinger, Going beyond bell's theorem, in *Bell's Theorem, Quantum Theory and Conceptions of the Universe* (Springer, Dordrecht, 1989), pp. 69–72.
- [46] D. Cruz, R. Fournier, F. Gremion, A. Jeannerot, K. Komagata, T. Tosić, J. Thiesbrummel, C. L. Chan, N. Macris, M. A. Dupertuis, and C. Javerzac-Galy, Efficient quantum algorithms for GHZ and W states, and implementation on the IBM quantum computer, *Adv. Quantum Technol.* **2**, 1900015 (2019).
- [47] J. B. Altepeter, D. Branning, E. Jeffrey, T. C. Wei, P. G. Kwiat, R. T. Thew, J. L. O'Brien, M. A. Nielsen, and A. G. White, Ancilla-assisted quantum process tomography, *Phys. Rev. Lett.* **90**, 193601 (2003).
- [48] M. S. ANIS *et al.*, Qiskit: An open-source framework for quantum computing, *zenodo* (2023), doi:10.5281/zenodo.8190968.
- [49] S. Bravyi and D. Maslov, Hadamard-free circuits expose the structure of the clifford group, *IEEE Trans. Inf. Theor.* **67**, 4546 (2021).
- [50] K. Wright, K. Beck, S. Debnath, J. Amini, Y. Nam, N. Grzesiak, J.-S. Chen, N. Pisić, M. Chmielewski, C. Collins *et al.*, Benchmarking an 11-qubit quantum computer, *Nat. Commun.* **10**, 1 (2019).
- [51] J. Kunjummen, M. C. Tran, D. Carney, and J. M. Taylor, Shadow process tomography of quantum channels, *Phys. Rev. A* **107**, 042403 (2023).
- [52] <http://iaifi.org/>.
- [53] L. Cincio, Y. Subaşı, A. T. Sornborger, and P. J. Coles, Learning the quantum algorithm for state overlap, *New J. Phys.* **20**, 113022 (2018).