





Tensor network noise characterization for near-term quantum computers

Stefano Mangini ^{1,2,*} Marco Cattaneo ^{1,2} Daniel Cavalcanti,¹ Sergei Filippov ¹
Matteo A. C. Rossi ¹ and Guillermo García-Pérez¹

¹*Algorithmiq Ltd, Kanavakatu 3C 00160 Helsinki, Finland*

²*QTF Centre of Excellence, Department of Physics, University of Helsinki, P.O. Box 43, FI-00014 Helsinki, Finland*



(Received 15 February 2024; accepted 29 July 2024; published 26 August 2024)

Characterization of noise in current near-term quantum devices is of paramount importance to fully use their computational power. However, direct quantum process tomography becomes unfeasible for systems composed of tens of qubits. A promising alternative method based on tensor networks was recently proposed [Nat. Commun. 14, 2858 (2023)]. In this paper, we adapt it for the characterization of noise channels on near-term quantum computers and investigate its performance thoroughly. In particular, we show how experimentally feasible tomographic samples are sufficient to accurately characterize realistic correlated noise models affecting individual layers of quantum circuits, and study its performance on systems composed of up to 20 qubits. Furthermore, we combine this noise characterization method with a recently proposed noise-aware tensor network error mitigation protocol for correcting outcomes in noisy circuits, resulting accurate estimations even on deep circuit instances. This positions the tensor-network-based noise characterization protocol as a valuable tool for practical error characterization and mitigation in the near-term quantum computing era.

DOI: [10.1103/PhysRevResearch.6.033217](https://doi.org/10.1103/PhysRevResearch.6.033217)

I. INTRODUCTION

Near-term quantum computers are currently entering what has been termed as the “utility era” [1]. This key advancement has been possible thanks to quantum error mitigation [2], that is, a collection of techniques for mitigating and eventually eliminating the effects of noise in quantum computers without relying on quantum error correction. Despite encouraging recent progresses [3], quantum error correction is currently out of reach for useful quantum computing.

Many error mitigation techniques require (ideally perfect) knowledge of the noise channels on the device, i.e., of the actual physical processes implemented on the real machine instead of the ideal unitary gates or ideal sharp measurements. Characterizing quantum processes for tens to hundreds of qubits is, however, not a trivial task. Standard state and process tomography [4] requires an exponential amount of resources as a function of the number of qubits [5]. Different techniques have been proposed to overcome this key issue and other inherent difficulties in tomographic methods, such as the enforcement of physical constraints. Examples include, but are not limited to, twirling methods that tailor the investigated processes to specific simpler forms [6,7], classical shadow methods to reconstruct quantum processes [8,9], tensor network methods to characterize non-Markovian

evolution processes [10,11], compressed sensing techniques for low-rank quantum processes [12,13], and methods that ensure meaningful tomography by appropriately restricting the reconstructed process to physical subspaces, for instance by means of constrained gradient descent [14] or by projection onto the set of processes [15].

In this paper, we develop the recently introduced tensor-network-based quantum process tomography method proposed by Torlai *et al.* [16], and apply it to the problem of noise characterization in near-term devices. This method consists in finding an efficient tensor network representation [17–20] of the quantum process under scrutiny. Torlai *et al.* benchmark their method with ideal circuits (i.e., unitary transformations) up to 10 qubits, and a noisy operation on five qubits subject to single-qubit errors. However, some applications, such as probabilistic error cancellation (PEC) [21,22] or tensor network error mitigation (TEM) [23], require the knowledge of the performance of individual gates or layers of gates. Therefore, here we focus on the characterization of each individual layer of a given circuit. Moreover, we exploit the fact that the process can be split into an ideal and a noisy part, and characterize only the latter. These two modifications have the advantage of alleviating the numerical requirements of the method. Finally, we consider realistic types of correlated noise, including the noisy model observed on IBM devices, and investigate systems of up to 20 qubits in size.

We study the tensor-network-based noise learning procedure by running several numerical experiments for the characterization of various correlated noise channels with brickwork-like structure and realistic noise parameters, which are of great relevance in near-term quantum computing. While our analysis is restricted to tensor networks with 1D connectivity (such as MPOs), the proposed method could be

*Contact author: stefano.mangini@algorithmiq.fi

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. (a) Description of the tensor-network-based noise characterization pipeline. The goal is to characterize the noise map \mathcal{N} unavoidably accompanying an ideal unitary operation \mathcal{U} (in the plot, a layer of CNOT gates) when this is executed on real quantum hardware. A number of experimental tomographic samples obtained with random preparations and measurements are collected on the noisy quantum computer to learn the noise map. For a single experimental shot, the noise channel \mathcal{N} we aim to reconstruct (red rectangles) acts on the state that we denote by tomographic state ρ_α , prepared by the single-qubit gates (green squares) followed by the unitary channel \mathcal{U} . The state is then measured through a collection of POVMs with effects Π_β (blue squares). (b) Representation of the tomographic experiment as tensor network, where the noise channel under investigation is written as a locally purified density operator (LPDO) Λ_θ parameterized by the quantities θ . The noise channel is learned by training the LPDO according to a suitable cost function, so that it best explains the tomographic measurement statistics observed on the quantum device. (c) Tensor network error mitigation (TEM) applied to the full noisy process \mathcal{E} using the results of the noise characterization experiment.

generalized to more complex topologies of qubit connectivity, albeit at the cost of more demanding classical computations.

In this paper, we discuss the necessary amount of experimental settings and measurement shots to obtain an accurate reconstruction, and find that collecting statistics on just a limited number of random experiments with informationally complete states and measurements provides sufficient data to accomplish the task. In particular, we observe that linearly many experimental samples in the number of qubits suffice to ensure very good reconstructions. Clearly, one can improve the reconstruction by increasing the number of experimental settings (input states and measurement bases) or the number of allocated shots per setting. For example, only 10^3 different experiments with 10^3 measurement shots each are sufficient to characterize a correlated brickwork layer of depolarizing noise channel on $n = 20$ qubits with an error of $\approx 10^{-4}$, as measured in terms of Frobenius distance between the ideal and reconstructed channels. We also confirm the good reconstruction accuracy by comparing values in the Pauli transfer matrix of the true and the reconstructed processes, and again find a good agreement. Additionally, we address the effect of state preparation and measurement (SPAM) errors on reconstruction accuracy, demonstrating not only its robustness against small errors but also that SPAM error-free performance can be achieved by calibrating the quantum device using existing quantum detector tomography methods [24].

We also investigate the performance of the method in conjunction with the tensor network error mitigation (TEM) protocol [23] in noisy Clifford circuits of up to 10 qubits and 30 layers. The combined characterization and mitigation approach is capable of mitigating noise and predicting the expected value of heavy Pauli observables with high accuracy (relative error of the order of 10^{-2}). This suggests that the characterization protocol is a valuable tool for practical error

mitigation in the near-term era. In Fig. 1 we summarize the main idea of the presented analysis.

The paper is structured as follows. In Sec. II, we review the tensor network representation of processes, along with Torlai *et al.*'s process tomography method (with some technical modifications). In Sec. III, we introduce the numerical experiments used to test the method, and describe the noise models that we have taken into consideration. The results of these numerical experiments are then presented in Sec. IV, while Sec. V is devoted to analyzing the combination of the process characterization method with the recently proposed tensor-network-based error mitigation protocol (TEM). Finally, we offer some concluding remarks in Sec. VI.

II. TENSOR NETWORK PROCEDURE FOR NOISE CHARACTERIZATION

In this section we introduce all the necessary tools for describing the tensor network noise characterization protocol, graphically summarized in Figs. 1(a) and 1(b).

A. Tensor network representation of noise

Let us consider a system of n qubits, the total Hilbert space \mathcal{H} of which has dimension 2^n . We aim at estimating a generic noise channel, which is formally described by the completely positive and trace-preserving (CPTP) quantum map \mathcal{N} [25] belonging to the space of bounded operators acting on the set of density matrices of the qubits.

Different representations of \mathcal{N} are available [5,25,26], such as the Choi matrix or the Liouville superoperator representation [27]. For our purposes, we choose to represent \mathcal{N} through

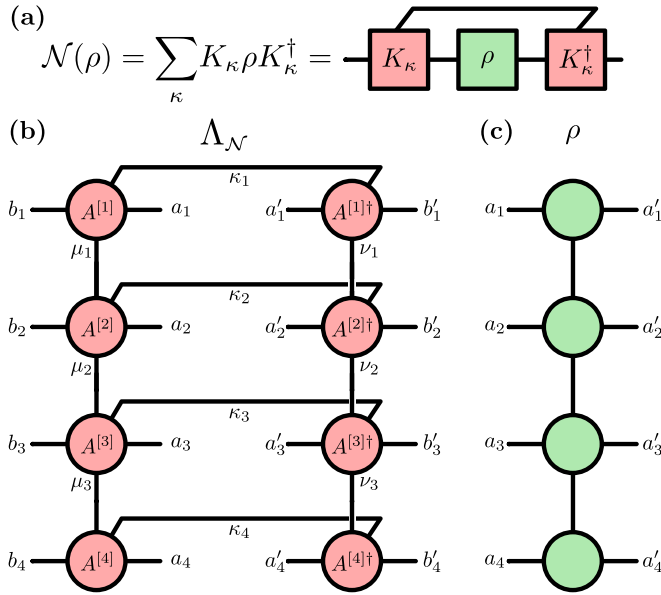


FIG. 2. (a) Kraus representation of the noise channel \mathcal{N} applied on a state ρ . (b) Tensor network representation of the channel \mathcal{N} as the locally purified density operator (LPDO) $\Lambda_{\mathcal{N}}$. (c) Tensor network representation of the state ρ as matrix product operator (MPO). The action of \mathcal{N} on ρ in tensor network notation is obtained by connecting the two tensor networks according to the indices as in the figure.

its *Kraus decomposition*, defined implicitly by

$$\mathcal{N}[\rho] = \sum_{\kappa} K_{\kappa} \rho K_{\kappa}^{\dagger}, \quad (1)$$

where ρ is a quantum state, and K_{κ} are *Kraus operators* acting on the Hilbert space of the qubits, satisfying the trace preserving condition

$$\sum_{\kappa} K_{\kappa}^{\dagger} K_{\kappa} = \mathbb{I}. \quad (2)$$

For any n -qubit map \mathcal{N} , the Kraus operators can be chosen in such a way that their number is at most 4^n .

Each representation of \mathcal{N} can be described in different ways using the tensor network formalism [27]. In our case, we focus on a tensor network representation of the Kraus operators known in the literature as *locally purified density operator* (LPDO) [16,28,29], which is depicted in Fig. 2. We denote the LPDO representation of a channel \mathcal{N} as $\Lambda_{\mathcal{N}}$, and we refer to Appendix A 1 for the explicit expression of the tensor components of such tensor network.

Let $\rho \in \mathbb{C}^{2^n \times 2^n}$ be the density matrix of a system of n qubits. Such matrix can be written in matrix product operator (MPO) form as [17,19,20,27,30]

$$\rho = \sum_{\lambda_1, \dots, \lambda_{n-1}} \rho_{\lambda_1}^{[1]} \otimes \rho_{\lambda_1, \lambda_2}^{[2]} \otimes \dots \otimes \rho_{\lambda_{n-1}}^{[n]}. \quad (3)$$

Then, using the LPDO representation for the quantum channel $\Lambda_{\mathcal{N}}$, and the MPO form for a state ρ , the action of the quantum channel $\mathcal{N}[\rho]$ can be written in tensor network

notation as

$$\begin{aligned} \Lambda_{\mathcal{N}}[\rho] &= \sum_{\kappa=1}^{\chi_{\kappa}} \sum_{\mu=1}^{\chi_{\mu}} (A_{\mu_1, \kappa_1}^{[1]} \otimes A_{\mu_1, \mu_2, \kappa_2}^{[2]} \otimes \dots \otimes A_{\mu_{n-1}, \kappa_n}^{[n]}) \\ &\times \sum_{\lambda} \rho_{\lambda_1}^{[1]} \otimes \rho_{\lambda_1, \lambda_2}^{[2]} \otimes \dots \otimes \rho_{\lambda_{n-1}}^{[n]} \\ &\times \sum_{\nu=1}^{\chi_{\nu}} (A_{\nu_1, \kappa_1}^{[1]\dagger} \otimes A_{\nu_1, \nu_2, \kappa_2}^{[2]\dagger} \otimes \dots \otimes A_{\nu_{n-1}, \kappa_n}^{[n]\dagger}), \quad (4) \end{aligned}$$

where $A_{i,j,k}^{[m]} \in \mathbb{C}^{2 \times 2}$ are local operators acting on single-qubit sites, $A_{i,j,k}^{[m]\dagger} := (A_{i,j,k}^{[m]})^{\dagger}$ are their transposed conjugates, and they interact with the corresponding local matrices of ρ as shown in Fig. 2. The indices $\mu = (\mu_1, \dots, \mu_{n-1})$ and $\nu = (\nu_1, \dots, \nu_{n-1})$ with $\mu_j, \nu_j = 1, \dots, \chi_b^{(j)}$ are the so-called *virtual bond indices* of the LPDO, while $\kappa = (\kappa_1, \dots, \kappa_n)$ with $\kappa_j = 1, \dots, \chi_{\kappa}^{(j)}$ are called *Kraus indices*. The maximal *bond dimension* of the LPDO is defined as the size of the largest of the virtual bond indices, $\chi_b = \max_j \chi_b^{(j)}$. Similarly, the maximal *Kraus dimension* is instead given by $\chi_{\kappa} = \max_j \chi_{\kappa}^{(j)}$. The LPDO structure has also been used to represent the positive Choi matrix of the channel \mathcal{N} in a very similar way [16,29]. Finally, the Kraus decomposition expressed by the LPDO can be easily transformed into an MPO, which is the superoperator representation of the quantum channel in Liouville space [27], as described in Appendix A 2.

It is important to stress that we consider the task of characterizing shallow processes with a clear local structure, which admit an efficient classical tensor network representation with low bond dimensions. By *shallow*, we hereby refer to quantum operations acting on n qubits that create only short-range qubit-qubit correlations, with this range being independent of n . For instance, a layer of CNOTs is shallow because it creates entanglement between pairs of qubits only, so the range of the correlations is 2, independently of the total number of qubits. Our method may still be used to characterize noise in more complex circuits but, in order to avoid an exponential scaling of the bond dimension of their tensor network representations, each deep quantum circuit should be divided into elementary shallow layers, with our protocol applied independently to each of them.

Applying the LPDO on the MPO ρ in Eq. (4) corresponds to the action of the Kraus decomposition of the channel \mathcal{N} on a quantum state. Indeed, we can group the Kraus indices in a single multi-index $\kappa = \{\kappa_1, \dots, \kappa_n\}$, the upper summation limit of which is equal to the product of all the dimensions of the Kraus indices. Then, we obtain a single set of global Kraus operators K_{κ} , the MPO structure of which is given by

$$K_{\kappa} = \sum_{\mu_1, \dots, \mu_{n-1}} A_{\mu_1, \kappa}^{[1]} \otimes A_{\mu_1, \mu_2, \kappa}^{[2]} \otimes \dots \otimes A_{\mu_{n-1}, \kappa}^{[n]}, \quad (5)$$

and the action of the channel \mathcal{N} on the state ρ in Eq. (4) is equivalent to Eq. (1).

Note, however, that while the LPDO structure is completely positive by design, it does not automatically satisfy the trace preserving condition of Eq. (2), which translates to $\text{Tr}_{b,b'}[\Lambda_{\theta}] = \mathbb{I}$ in LPDO notation [27], where Λ_{θ} is a generic LPDO parameterized by quantities θ , and the indices b and

b' in the partial trace are as in Fig. 2. As proposed in [16], one can enforce such property by adding a trace preservation penalty term in the cost function at training time, but more explicit constraints have also been proposed [31]. In addition, as discussed in Sec. II C, one can initialize the tensor elements θ in such a way that the resulting LPDO is at least correctly normalized $\text{Tr}[\Lambda_\theta] \approx 2^n$ in expectation value.

B. Data sampling

In what follows, we describe the tomographic sampling strategy based on the recent paper by Torlai *et al.* [16]. It consists in generating N_{set} experimental tomographic settings in a randomized way in order to obtain sufficient information about the quantum channel we aim to characterize. For experiments on near-term quantum computers, we define an experimental tomographic setting as a collection of n separable, single-qubit input states, and a collection of n choices of measurement basis, one for each qubit of the device.

In our paper, for each qubit, we only use three possible measurement bases, corresponding to the Pauli measurements, i.e., we measure along the X , Y , or Z axes. The POVM effects of the measurement along the X axis, for instance, are $|x\rangle\langle x|$ and $|-x\rangle\langle -x|$, where $|x\rangle$ ($|-x\rangle$) is the eigenstate of X with eigenvalue $+1$ (-1), and equivalently for other axes. Since the typical native measurement on a quantum computer is in the computational basis, we can perform a measurements in the X or Y directions by simply applying suitable single-qubit rotations before the measurement.

Each single-qubit input state, which for the j th qubit is denoted by $\rho_{\alpha_j}^{\text{in}}$, is drawn from an informationally complete (IC) pool of states $\mathcal{P} = \{\rho_{\alpha_j}^{\text{in}}\}_j$, forming a basis in the space of single-qubit density matrices. The minimum number of states in the pool is 4, but this set may also be larger. For instance, one may choose the overcomplete set comprising the six Pauli eigenstates $\mathcal{P} = \{|\pm x\rangle\langle \pm x|, |\pm y\rangle\langle \pm y|, |\pm z\rangle\langle \pm z|\}$, which can be prepared by applying a single-qubit unitary gates to state $|+z\rangle = |0\rangle$. Alternatively, for the sake of employing less experimental settings, one may also use a symmetric and informationally complete (SIC) set of input states containing only four elements, for example given by [32]

$$\begin{aligned} \rho_j &= \frac{1}{4}(1 + \mathbf{a}_j \cdot \boldsymbol{\sigma}), \quad \text{with} \\ \boldsymbol{\sigma} &= (X, Y, Z), \\ \mathbf{a}_0 &= \frac{(+1, +1, +1)}{\sqrt{3}}, \quad \mathbf{a}_1 = \frac{(+1, -1, -1)}{\sqrt{3}}, \\ \mathbf{a}_2 &= \frac{(-1, +1, -1)}{\sqrt{3}}, \quad \mathbf{a}_3 = \frac{(-1, -1, +1)}{\sqrt{3}}. \end{aligned} \quad (6)$$

In practical scenarios, one usually seeks to characterize the noise channel \mathcal{N} accompanying an ideal unitary layer \mathcal{U} . Thus, we define the *tomographic state* as $\rho_\alpha = \mathcal{U}[\rho_\alpha^{\text{in}}]$, obtained by evolving the initial randomly chosen input states through the ideal unitary whose noise we want to characterize, see Fig. 1(a).

An experimental setting is then defined in the following way. We first draw one state $\rho_{\alpha_j}^{\text{in}}$ for each qubit $j = 1, \dots, n$ from a uniform distribution over the pool \mathcal{P} , which, for

simplicity, we assume to be equal for all qubits. Let us use the collective index $\alpha = \{\alpha_1, \dots, \alpha_n\}$ to denote the choice of initial states for all the qubits. Next, we draw one measurement basis $\beta_j \in \{X, Y, Z\}$ for each qubit j , again from a uniform distribution, and use the collective index $\beta = \{\beta_1, \dots, \beta_n\}$ to indicate which measurement basis has been chosen on each qubit.

A single tomographic experiment will then consist in preparing the state $\rho_\alpha^{\text{in}} = \bigotimes_{j=1}^n \rho_{\alpha_j}^{\text{in}}$ on all qubits, evolving it through the *full noisy process* $\mathcal{E} = \mathcal{N} \circ \mathcal{U}$, and finally measuring each qubit in the proper basis β_j . Note that, since we know the logical operation \mathcal{U} the noise channel of which we are characterizing, we isolate the noise channel \mathcal{N} from the full noisy process, and regard this experiment as the application of an unknown channel \mathcal{N} onto the known tomographic state $\rho_\alpha = \mathcal{U}[\rho_\alpha^{\text{in}}]$. Note also that the tomographic state is in general entangled, but its spatial correlations (which impacts the bond dimension of its MPO representation) are short range if the unitary circuit \mathcal{U} is shallow. This is the case we are interested in, as we are considering noise affecting single-layer instructions. This means that we can represent ρ_α as a tensor network efficiently.

A single-qubit Pauli measurement can be described by a POVM with only two effects corresponding to outcome $\zeta = +1$ or $\zeta = -1$, so the outcome of a single n -qubit tomographic experiment can then be represented as a vector $\boldsymbol{\zeta} = (\zeta_1, \dots, \zeta_n)$. The probability of obtaining the outcome $\boldsymbol{\zeta}$ for a fixed experimental setting, defined by the choice of input states α and measurement bases β , is given by the Born rule

$$p(\boldsymbol{\zeta}|\alpha, \beta) = \text{Tr} \left[\mathcal{N}[\rho_\alpha] \bigotimes_{j=1}^n \Pi_{\zeta_j}(\beta_j) \right], \quad (7)$$

where $\Pi_{\zeta_j}(\beta_j)$ is the effect corresponding to the outcome ζ_j for a measurement in the β_j basis performed on the j th qubit.

We point out that this sampling strategy assumes that we know the input states and the measurements perfectly well. It is, however, well known that this is not the case on near-term quantum computers, where state preparation and measurement (SPAM) errors are currently unavoidable [33]. This consideration has led to different self-consistent tomographic methods to determine the input states, the computational gates, and the measurement outcomes consistently and simultaneously [33,34]. Unfortunately, these procedures are usually too resource-expensive for useful near-term applications (i.e., involving tens of qubits), even when considering optimized strategies [35] (a more promising procedure is considered in Ref. [36], which performs SPAM-robust shadow estimation of some properties of a gate set from random gate sequences). For this reason, in the case of tensor network noise characterization, one may adopt a more practical solution on a real quantum computer: before running the noise tomography experiment, one can perform a calibration of the machine yielding a self-consistent description of the input states and measurements, such as the one based on semidefinite programming recently proposed by some of the authors [24]. Then, the output of the protocol, that is, a set of input states and POVM effects that are self-consistent and capture what is physically prepared and measured on the device, would be

used in the noise characterization procedure, as in Eq. (7). We point out that all the results of this paper, apart from the ones discussed in Sec. IV C, do not take SPAM errors into account and do not employ self-consistent tomography. We leave a detailed study of the efficacy of self-consistent strategies to ameliorate SPAM errors in the tensor network noise characterization protocol for future works.

1. Efficiency of the sampling strategy

Let $R = |\mathcal{P}|$ be the number of possible input states in the pool \mathcal{P} , for which we know that $R \geq 4$. For each qubit we then have a total of $3R$ possible different experimental settings, given by all the possible combinations of input states (R) and measurement bases (the three Pauli measurements). For a system of n qubits and assuming independent preparations and measurements, the total number of different settings therefore is equal to $(3R)^n$, which is a formidable number even for $n \sim 10$. Such exponential scaling, typical of quantum tomography, makes it unfeasible to implement all the possible different settings in a tomographic experiment when a large number of qubits is used.

We circumvent the issue by randomly generating only N_{set} different tomographic settings using the procedure described above, and allocating a number of N_{shots} measurement shots to each of these settings. Therefore, a complete tomographic experiment will use a total of $N = N_{\text{shots}} \times N_{\text{set}}$ measurements. Despite not having the exponentially large tomographic data required to reconstruct arbitrary channels, such limited information can still be sufficient in our scenario, where we are interested in learning processes with a local structure, which can be effectively described using a only limited number of values. Additionally, as already stated above, the local structure of the noise is also a key assumption for its efficient description through a LPDO tensor network with a small bond dimension.

We point out that here our analysis deviates from that performed in Ref. [16], where the effect of number of shots per setting is not taken into account. Instead, in this paper we consider a scenario that is more realistic for near-term quantum computers, and in particular for superconducting quantum devices available on the cloud [37], where the total measurement budget N is allocated by executing N_{shots} shots on each of a limited number of experimental settings N_{set} . In fact, given access constraints and the relatively long wall time needed to compile instructions on current near-term hardware [37,38], it is of paramount importance to be able to extract relevant information out of only a limited number of distinct experimental setups. The number of settings N_{set} is often the practical bottleneck for current experiments, while the number of shots per settings N_{set} comes at a much lower cost, both in terms of accessibility and execution time.

We point out that such random sampling strategy, based on generating only a reduced number of settings that in principle is not sufficient for full process tomography of the quantum channel, is equivalent to sampling for *shadow tomography* [39–41] of quantum processes, which has been explored in some recent papers [8,9,36,42]. The difference between our method and shadow process tomography lies in the post-processing of the sampled data. Instead of applying linear

inversion [8] or using more refined fitting methods [9,36,42] starting from the raw data, we train the LPDO structure to obtain the most accurate tensor network description of the channel, by finding the parameters in the tensor network that best explain the experimental data see Sec. II C. We leave a direct comparison of our approach with those based on shadow tomography as a topic for future studies.

2. Alternative local sampling strategies

The sampling strategy described so far generates different global random settings for the tomographic experiment. However, one may also adopt a different strategy, which assumes that the process \mathcal{N} under investigation only generates local correlations, and may therefore be well-characterized by using only local information about the subsystems. Broadly speaking, such methods select the tomographic settings in a way that one is able to collect data on all reduced subsystems of given locality, and then reconstruct the whole channel based on such local information. Local tomographic strategies building on such ideas have been successfully used in the literature to characterize quantum states with local correlations [43–45].

However, such strategies cannot be straightforwardly applied to the case of process tomography, where one has to probe the channel under investigation not only with informationally complete measurements, but also input states (we do not consider the reduction of channel to state tomography via the Choi–Jamiołkowski isomorphism as this requires the use of ancillary qubits [46]). Taking into account the burden of state generation, one can check that the resources needed to collect local tomographic data quickly become experimentally unfeasible, even for low locality. We refer the interested reader to Appendix B, where we discuss in detail possible tomographic strategies for accessing local data, also based on lightcone arguments stemming from the brickwork structure of the channel under investigation. In addition, in Appendix B, we also provide preliminary numerical evidence that the global random strategy performs better than a simple local strategy, both in terms of total measurements needed and reconstruction accuracy. All the numerical results presented in the following are thus obtained following the random generation of tomographic settings described in Sec. II B.

C. Tensor network optimization

The optimization of the LPDO Λ_θ over a set of parameters θ is based on the approach proposed in Ref. [16], in which the tensor network is trained so that the predicted distribution of outcomes best matches the observed measurement statistics.

Formally, for N total experimental shots, let $S = \{(\rho_{\alpha_m}, \Pi_{\zeta_m}(\beta_m))\}_{m=1}^N$ be the tomographic dataset collected on a real quantum device consisting of N pairs of tomographic states ρ_{α_m} and corresponding measured effects $\Pi_{\zeta_m}(\beta_m)$, where subscript m labels single experimental shots. Then, the LPDO can be fitted to the experimental data by minimizing the objective function [16]

$$D_{\text{kl}}(\theta; S) = -\frac{1}{N} \sum_{m=1}^N \log p(\zeta_m | \alpha_m, \beta_m; \theta) \\ = -\frac{1}{N} \sum_{m=1}^N \log \text{Tr}[\Lambda_\theta(\rho_{\alpha_m}) \Pi_{\zeta_m}(\beta_m)], \quad (8)$$

which is a Monte Carlo approximation to the Kullback-Leibler divergence between the true probability distribution of the quantum process under investigation (7), and the one generated by the parameterized tensor network.

In addition, the authors in Ref. [16] propose to using an additional penalty term in the loss function that favors physically valid LPDO satisfying the trace preservation (TP) condition (2). Such penalty term is given by the normalized Frobenius distance between the identity and the MPO obtained by contracting the outer legs in the LPDO (\mathbf{b} and \mathbf{b}' in Fig. 2), namely

$$\delta_{\text{TP}}(\boldsymbol{\theta}) = \frac{\|\text{Tr}_{\text{out}}[\Lambda_{\boldsymbol{\theta}}] - \mathbb{I}\|_F}{2^{n/2}}, \quad (9)$$

where $\|A\|_F^2 := \text{Tr}[A^\dagger A]$ is the operator Frobenius norm.

Finally, the complete loss function used to drive learning process is then given by

$$L(\boldsymbol{\theta}; S) = D_{\text{KL}}(\boldsymbol{\theta}; S) + \eta \delta_{\text{TP}}(\boldsymbol{\theta}), \quad (10)$$

where $\eta \in \mathbb{R}$ is a hyperparameter tuning the importance of the TP condition in the training process. In all our numerical simulations we set $\eta = 1.2$, which was heuristically found to consistently ensure a good convergence to a properly normalized LPDO, $\text{Tr}[\Lambda_{\boldsymbol{\theta}}] \approx 2^n$, at the end of training. While different choices do not impact the end results sensibly, these may result either in slower convergence times towards physically meaningful solutions, or to solutions having incorrect—but still close, if $\eta \approx 1$ —trace. Also, we note that such hyperparameter could be itself adapted during training, but we leave this investigation as a topic for future studies.

Despite its effectiveness in training the LPDO, the loss function (10) does not satisfy the symmetry requirement for a true distance and has a limited physical interpretation. For this reason, we measure the reconstruction error in the characterization procedure through the quantity

$$\Delta(\Lambda, \Lambda_{\boldsymbol{\theta}}) = \frac{\|\Lambda - \Lambda_{\boldsymbol{\theta}}\|_F^2}{2^{2n}}, \quad (11)$$

consisting of a properly normalized Frobenius distance between the true channel Λ and the trainable one $\Lambda_{\boldsymbol{\theta}}$, similar to the fidelity-like error measure used in [16]. Needless to say, this measure is only available in classical numerical simulations, where the true channel is known.

1. Normalized initialization of the LPDO

At the start of the training procedure, the parameterized LPDO is initialized with random values. However, this typically leads to unphysical quantum maps not respecting either the TP constraint (9) or the normalization condition $\text{Tr}[\Lambda] = 2^n$. In order to alleviate this issue, we first employ a parameter initialization method that yields, in expectation value, a correctly normalized LPDO, and then variationally pre-optimize the tensor network in order to appropriately satisfy the trace preserving condition. Both these strategies were heuristically found to improve convergence to good solutions and to stabilize the training process by avoiding numerical instabilities related to unphysical initializations of the tensor network.

The tensor elements $\theta_k \in \mathbb{C}$ of the LPDO $\Lambda_{\boldsymbol{\theta}}$ are randomly initialized from a complex Gaussian distribution

$$\begin{aligned} \theta_k &= \text{Re}(\theta_k) + i \text{Im}(\theta_k) \\ &\text{with } \text{Re}(\theta_k), \text{Im}(\theta_k) \sim \mathcal{G}(0, \sigma^2), \end{aligned} \quad (12)$$

where $\mathcal{G}(0, \sigma^2)$ denotes a Gaussian distribution with zero mean and variance σ^2 . As proven in Appendix E 1, under such circumstances, one can explicitly compute the expectation value of the trace of the LPDO $\text{Tr}[\Lambda_{\boldsymbol{\theta}}]$ upon initialization, which amounts to

$$\mathbb{E}_{\boldsymbol{\theta}}[\text{Tr}[\Lambda_{\boldsymbol{\theta}}]] = (8\sigma^2\chi_{\kappa})^n \chi_b^{n-1}, \quad (13)$$

where n is the number of qubits, and χ_{κ} and χ_b are the Kraus dimension and the virtual bond dimension of the LPDO, respectively. Thus, by sampling the initial parameters according to a Gaussian with variance

$$\sigma^2 = 2 / (8\chi_{\kappa}\chi_b^{1-1/n}) \quad (14)$$

the LPDO is properly normalized to the correct value $\mathbb{E}_{\boldsymbol{\theta}}[\text{Tr}[\Lambda_{\boldsymbol{\theta}}]] = 2^n$ on average upon random initialization.

Additionally, we further pre-optimize the initial LPDO to satisfy the TP constraint by variationally minimizing the penalty term $\delta_{\text{TP}}(\boldsymbol{\theta})$ (9) by means of an optimizer before the actual training of the LPDO starts.

2. Details on numerical simulations and optimization

All numerical experiments are run using the python tensor network library `quimb` [47], in combination with libraries for automatic differentiation and optimization `jax` [48] and `optax` [49].

The trace preserving pre-optimization of the LPDO is run using optimizer L-BFGS-B [50] provided in `quimb`. The training of the LPDO by minimization of the loss function $L(\boldsymbol{\theta}; S)$ (10) is done using the Adam optimizer [51], together with an additional custom exponential decay schedule of the learning rate, which was found to improve convergence. We refer to Appendix E 2 for further details on the optimization process, including details on the training batch size and dimension of the test set.

III. NOISE TOMOGRAPHY EXPERIMENTS

For the sake of benchmarking our noise characterization method, we consider the task of determining the noise \mathcal{N} accompanying the simple yet very common logical instruction \mathcal{U} consisting of an n -qubit even layer of CNOTs, as depicted in Fig. 1(a). We simulate different noise models applied to this circuit layer, which also take into account crosstalk errors between nearby qubits. We run some classical simulations of the tomographic experiment to characterize such a noisy circuit, and we compare the results of the tensor network reconstruction with the true noisy channel.

In the current section, we describe the different noise models we have employed in the classical simulation. The ensuing Sec. IV is devoted to the study of the numerical results and performance analysis. Finally, in Sec. V, we also validate the accuracy of the channel characterization scheme by employing it to mitigate the noise on a noisy circuit through the

recently proposed tensor network error mitigation protocol (TEM) [23].

In our simulations, we choose three different realistic noise models that are of particular importance for near-term quantum computers: the sparse Pauli-Lindblad noise model [7], the incoherent depolarizing noise model, and the coherent depolarizing noise model. Notably, all these multiqubit correlated noise models can be graphically represented by the *brickwork* circuit structure depicted in Fig. 1.

A. Sparse Pauli-Lindblad noise model

The sparse Pauli-Lindblad noise model is a locally correlated noise model that was recently introduced as an effective method to describe errors in superconducting quantum hardware [7]. Such noise model is described by the map

$$\mathcal{N}_{\text{SPL}}[\rho] = \prod_{k \in \mathcal{K}} (\omega_k \cdot + (1 - \omega_k) P_k \cdot P_k) \rho, \quad (15)$$

where \mathcal{K} is a poly(n)-size subset of the 4^n n -qubit Pauli operators, and \cdot is a placeholder for the argument of the quantum map, e.g., $(P_k \cdot P_k)\rho = P_k \rho P_k$. The coefficients ω_k are defined as

$$\omega_k = (1 + e^{-2\lambda_k})/2 \quad (16)$$

with $\lambda_k \geq 0$ being non-negative parameters defining the strength of each Pauli interaction term in the Lindblad master equation description of the noise $\mathcal{N}_{\text{SPL}}[\rho] = \exp[\mathbf{L}](\rho)$, generated by $\mathbf{L}(\rho) = \sum_{k \in \mathcal{K}} \lambda_k (P_k \rho P_k - \rho)$ [7]. These parameters can be estimated, for instance, by *cycle benchmarking* [6,7], which removes any preparation and measurement (SPAM) errors [33]. However, this technique is known to provide only an ambiguous reconstruction of the parameters [7,52]. The sparse Pauli-Lindblad noise model is typically a faithful description of the noise channels on the device if one performs randomized compiling [53] to approximately transform the possibly coherent true noise to an incoherent Pauli channel.

The expression of the noise in Eq. (15) may take into account crosstalk errors between very far away qubits, depending on how we choose the set \mathcal{K} . To opt for a more realistic description of spatially correlated errors on the CNOT layer, we choose \mathcal{K} such that it only accounts for first-neighbors crosstalk errors. That is, the coefficients k in \mathcal{K} can only refer to single- or two-qubit Pauli interaction terms acting on adjacent qubits, an assumption that has been experimentally validated several times [1,7]. Importantly, in our simulations we use realistic noise coefficients for SPL noise found in current superconducting quantum hardware (see Appendix C for further details and explicit coefficients).

We point out that the Pauli-Lindblad channel gives rise to a *Clifford* noise model [5], that is, the application of \mathcal{N}_{SPL} onto a Pauli operator returns the same Pauli operator scaled by a factor.

B. Incoherent depolarizing noise model

A two-qubit depolarizing noise channel is defined as

$$\mathcal{D}^{(p)}[\rho] = (1 - p)\rho + \frac{p}{4} \text{Tr}[\rho] \mathbb{I}, \quad (17)$$

where $p \in [0, 1]$ is the error rate. In the incoherent depolarizing noise model, for each CNOT gate in the unitary layer U , we apply one two-qubit depolarizing channel with error rate p on the target and control qubits. Moreover, in order to simulate first-neighbors crosstalk errors, we consider another layer of two-qubit depolarizing channels with error rate $p/2$ on the nearby qubits that are not connected by a CNOT gate. This then creates the *brickwork* structure of Fig. 1(a), as the noise model consists of one even layer of two-qubit depolarizing channels followed by an odd layer with a lower error rate. Formally, the total noise channel on n qubits can be expressed as follows:

$$\mathcal{N}_{\text{inc}}^{(p)} = \prod_{k=1}^{\lfloor (n-1)/2 \rfloor} \mathcal{D}_{2k, 2k+1}^{(p/2)} \circ \prod_{k=1}^{\lfloor n/2 \rfloor} \mathcal{D}_{2k-1, 2k}^{(p)}, \quad (18)$$

where the depolarizing channel $\mathcal{D}_{k, k+1}^{(p)}$ is acting on the qubits k and $k+1$.

For the simulations in the main text we set the depolarization strength to $p = 10^{-3}$, which is only slightly lower than two-qubit gate errors reported for state-of-art machines based on, e.g., superconducting circuits [37,54], neutral atoms [3], and ion-traps [55], and foreseeably achievable in the near future. However, for completeness and as discussed in the ensuing sections, we also report results for a stronger depolarizing rate in Appendix D.

We note that, as for the sparse Pauli-Lindblad noise model (15), also the incoherent depolarizing noise model is a Clifford map.

C. Coherent depolarizing noise model

We extend our analysis to coherent error sources by considering the more complex case where, in addition to the brickwork depolarizing channel described above, the qubits are also affected by undesired single-qubit unitaries. Specifically, we assume that the overall noise process consists of a first layer of single-qubit random rotations used to simulate coherent noise, and the aforementioned correlated incoherent depolarizing error channel.

The complete noise channel can then be written as

$$\mathcal{N}_{\text{coh}}^{(p)} = \mathcal{N}_{\text{inc}}^{(p)} \circ \bigotimes_{j=1}^n U_j, \quad (19)$$

where U_j are single-qubit random rotations. Given three angle parameters ψ , φ , and ϕ , these random rotations can be parameterized as

$$U(\phi, \varphi, \psi) = \begin{pmatrix} e^{i\varphi} \cos \phi & e^{i\psi} \sin \phi \\ -e^{-i\psi} \sin \phi & e^{-i\varphi} \cos \phi \end{pmatrix}. \quad (20)$$

A single-qubit Haar-random unitary rotation can then be obtained by sampling uniformly ψ and φ from $[0, 2\pi]$, and $\phi = \arcsin \sqrt{\zeta}$ where ζ is sampled uniformly from $[0, 1]$ [56]. Since we assume single-qubit errors to be small, we consider restricted rotation angles given by $\zeta \leftarrow \epsilon \zeta$ and $\varphi \leftarrow \epsilon \varphi$ with $\epsilon = 10^{-3}$.

We sample one different random rotation for each qubit. If we apply the noisy circuit layer more than once, the random rotations on each qubit are the same for all layers, that is, we always associate the same noise channel to the same




FIG. 3. Example training run of the parameterized LPDO Λ_θ for learning the brickwork depolarizing noise $\mathcal{N}_p^{\text{dep}}$ in Eq. (18) with $p = 10^{-3}$ on a system of $n = 10$ qubits. The LPDO was trained with a dataset of 10^6 samples (10^3 experimental settings, 10^3 shots per setting), and with bond dimensions $\chi_\kappa = 2$, $\chi_b = 16$. Both the reconstruction error $\Delta(\Lambda, \Lambda_\theta)$ and the test loss $L(\theta)$ are minimized and converge in a modest number of training epochs. The TP penalty term in the cost function (9) helps the LPDO to converge to the correct normalization.

logical instruction (the even CNOT layer in our cases). Contrary to the sparse Pauli-Lindblad and incoherent depolarizing noise models, the coherent depolarizing noise model can be non-Clifford.

IV. RESULTS

In this section we investigate the effectiveness of the proposed tensor network noise characterization technique, and analyze how its performance scales with the tomographic dataset size, the number of qubits and the accuracy in estimating noise coefficients. Importantly, as discussed in Sec. II B, we stress again that in our analysis we consider realistic scenarios with a limited number of experimental settings and multiple shots per setting, and we find that this is sufficient to provide a reliable approximation of the noisy process. In all the numerical results reported below, the tomographic data for the channel reconstruction is obtained by sampling, for each qubit, input states from the SIC set of four states defined in Eq. (6), and measurements from the Pauli basis as described in Sec. II B.

From the classical computational viewpoint, the number of trainable parameters in the LPDO scales as $\mathcal{O}n\chi_\kappa\chi_b^2$, so the learning process remains efficient as long as the bond dimensions are small, which is the case for our task of characterizing shallow noisy operations (the largest Kraus and virtual bond dimensions used in our simulations are, respectively, $\chi_\kappa = 16$ and $\chi_b = 4$). Indeed, our largest training experiment with an LPDO of $n = 20$ qubits, with bond dimensions $\chi_b = \chi_\kappa = 4$, on a tomographic dataset consisting of $N = 10^6$ samples, can be run in about one hour on a laptop.

In Fig. 3, we report an example of the characterization of a brickwork depolarizing noise channel for $n = 10$ qubits with $N = 10^6$ shots, using an LPDO with $\chi_\kappa = 2$, $\chi_b = 16$. The reconstruction error (11) and the loss function evaluated on a




FIG. 4. Frobenius distance between true and reconstructed noise channel as a function of the number of shots, for the different noise models introduced in Sec. III, and with $n = 10$ qubits. We also plot the expected “shot-noise” scaling, decreasing as $1/N$. Each point in the plot shows the best value obtained in three different training runs of the LPDO initialized with different parameters, but on the same training dataset. We note that all training runs eventually converge to similar performances.

test set of samples both decrease along the training process and converge to a minimum value in a few training epochs. Also, after starting from the correct value (see Sec. II C 1), the TP penalty (9) in the loss function enforces the LPDO to converge to the correct normalization $\text{Tr}[\Lambda_\theta] \approx 2^n$ by the end of training.

In all the analyses presented below, we show the results obtained with the trained LPDOs $\Lambda_{\theta_{\text{opt}}}$ attaining the lowest test error (10) during the training process, a measure, which is accessible in real experiments and does not require knowledge of the process under characterization. In all experiments with brickwork depolarizing channels, $\chi_b = 2$ and $\chi_\kappa = 16$ were used, while $\chi_b = 4$ and $\chi_\kappa = 4$ were used for experiments involving sparse Pauli-Lindblad noise.

A. Accuracy vs number of shots

In order to evaluate the viability of tensor network noise learning with random tomographic settings in realistic experimental scenarios, we start by analyzing how the reconstruction accuracy behaves with the size N of the tomographic dataset. In particular, we consider a fixed budget of $N_{\text{set}} = 10^3$ experimental settings, and vary the number of shots allocated to each measurement setting, $N_{\text{shots}} \in \{1, 10, 10^2, 10^3, 10^4\}$. In Fig. 4, we report results for the characterization procedure of the three realistic noise models discussed in Sec. III, for a system of $n = 10$ qubits.

Interestingly, the reconstruction accuracy follows a shot-noise behavior—we consider the square of the usual shot-noise scaling \sqrt{N} to account for the square in the definition of the reconstruction error Δ (11), which signals that the learning procedure is able to take full advantage of additional tomographic samples. However, when the number of shots per setting is large enough $N_{\text{shots}} = 10^4$ ($N = 10^7$), the reconstruction accuracy starts deviating from the shot-noise scaling, at which point it would be beneficial to increase the number of settings rather than the shots per setting.




FIG. 5. Frobenius distance between true and reconstructed noise channels as a function of the number of qubits, for the different noise models introduced in Sec. III, and with $N_{\text{set}} = N_{\text{shots}} = 10^3$ (then, $N = 10^6$). The dashed lines are linear fits with parameters reported in the figure. Each point in the plot shows the best value obtained in three different training runs of the LPDO initialized with different parameters, but on the same training dataset. We note that all training runs eventually converge to similar performances.

This is especially evident for the sparse Pauli-Lindblad noise model, for which not only the training yields in general a slightly lower reconstruction accuracy, but the Frobenius distance also displays a significant deviation at large number of shots. We believe such behavior to be a consequence both of the intrinsically more complex structure of the sparse Pauli-Lindblad noise, and also of this channel being more noisy overall (see the noise coefficients in Figs. 11 and 7 below). In Appendix D we report results for the characterization of brickwork depolarizing noise with stronger intensity $p = 0.1$ (much larger than current two-qubit error rates [1,54]). The analysis is in agreement with similar but simpler results in [16], where the reconstruction accuracy was found to decrease in the presence of stronger noise sources. This could be understood as a consequence of the decrease of visibility of the useful signal with increasing noise, indicating that either more resources or a more fine-tuned training routine are needed to distinguish the signal from a background white noise.

B. Accuracy vs number of qubits

We now turn our attention to the investigation of the behavior of the reconstruction accuracy as a function of size of the system.

In Fig. 5, we report the accuracy obtained with $N = 10^6$ shots on systems of varying size, up to $n = 20$ qubits. We observe a favorable linear scaling of the reconstruction error with the number of qubits n for all noise models considered, which indicates the feasibility of the proposed approach for characterization purposes on near-term devices with a limited amount of qubits. Overall, the results in Figs. 4 and 5 suggest the use of linearly larger tomographic datasets to compensate for the linear decrease in reconstruction accuracy for larger system sizes.

C. Accuracy in the presence of SPAM errors

In this section we show how the proposed noise characterization method can be used also in the presence of




FIG. 6. Effect of SPAM errors on the noise characterization procedure on $n = 10$ qubits subject to sparse Pauli-Lindblad noise, with $N_{\text{set}} = N_{\text{shots}} = 10^3$. Both state preparation and measurement errors are parameterized as single-qubit depolarizing channels with intensities p_{prep} and p_{meas} , respectively. Measurements error are mitigated by using quantum detector tomography (QDT) [24] with 10^4 shots to reconstruct the noisy POVM effects.

SPAM errors, by combining it with techniques aimed at characterizing such state preparation and measurement noise. In particular, we employ the quantum detector tomography (QDT) procedure described in [24] to first reconstruct the noisy POVM effects that are actually implemented on the device, and then use such reconstructed effects in the noise characterization procedure. In fact, if state preparation errors are small compared to the other sources of error—as it is usually the case in current quantum hardware, we observe that the use of measurement tomography alone is already sufficient to recover the reconstruction accuracy obtained in the SPAM-free regime.

Quantum detector tomography is implemented by executing a set of circuits implementing only state preparation and measurement instructions. Assuming state preparation errors to be negligible compared to measurement errors, by probing the chosen POVM with a set of informationally complete states, one can realize a tomography of the quantum detector and hence reconstruct the real noisy effects $\Pi_{\zeta}(\beta) \rightarrow \tilde{\Pi}_{\zeta}(\beta)$ composing the POVM. These effects are then used in the loss function (10) to drive the noise characterization process. In the experiments below, QDT is run using a set of four informationally complete state to reconstruct the 6-outcome POVM obtained by performing Pauli measurements. This requires a total of $4 \times 3 = 12$ circuits, each of which is executed with 10^4 shots. Note that the reconstruction of the effects is itself only approximate, with better performance obtained with larger measurement budgets [24]. Additionally, the QDT procedure is run assuming an ideal preparation of the input states, while these are in fact also subject to errors. These two effects combined, namely the limited measurement budget and the occurrence of unknown preparation errors, then result in an imperfect reconstruction accuracy of the POVM effects.

In Fig. 6 we report the simulation results obtained by characterizing the sparse Pauli-Lindblad noise on a system of $n = 10$ qubits in the presence of realistic SPAM noise, with and without employing QDT to mitigate measurement errors. For the sake of simplicity we consider only incoherent errors: both




FIG. 7. Some diagonal coefficients of the MPOs in Pauli transfer matrix (21) of the true (green circles) and reconstructed (orange crosses) noise channels, for the different noise models introduced in Sec. III and $n = 10$ qubits. The numbers above the plot indicate the Pauli weight (i.e., the number of nonidentities) of the operators. The mismatch between the green circles and orange crosses arises from the error in channel tomography, obtained with a training set of $N = 10^7$ ($N_{\text{set}} = 10^3$, $N_{\text{shots}} = 10^4$) samples.

state preparation and measurement errors are parameterized as single-qubit depolarizing channels, with state preparation having depolarizing strength $p_{\text{prep}} = 10^{-4}$ (which is a reasonable value for single-qubit gates on near-term computers), and measurement error having $p_{\text{meas}} \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$.

When measurement errors are large, the noise learning procedure is unable to provide an accurate description of the noisy evolution, but this can be readily solved by using QDT to calibrate the device and train the LPDO using the reconstructed noisy effects. When SPAM errors are small enough instead, noise characterization obtains good reconstruction accuracy irrespective of the use of QDT. This can be understood by noticing that the true and noisy effects are now very close to each other, and QDT is unable of precisely distinguishing them using a limited number of shots. Additionally, in the regime where state preparation and measurement errors are of the same order of magnitude, QDT yields incorrect noisy effects since it was run assuming ideal state preparations, which can then impact the noise learning procedure. This issue may be solved by using self-consistent characterization protocols [24,34], but we leave this as a subject of future studies.

Overall, our results not only indicate the proposed tensor network noise learning procedure is stable against small SPAM errors, but also that it can be straightforwardly combined with existing detector tomography methods to calibrate

the measurement apparatus and cancel the effects of large measurement errors.

D. Reconstructed noise coefficients

While the Frobenius distance captures the overall difference between the two LPDOs, we investigate more physical figures of merit as well, such as specific coefficients within the tensor network representing the noise channel.

In particular, let us transform the LPDO representing the channel into an MPO, i.e., let us switch to the superoperator representation of the channel (see Appendix A 2 for details). Moreover, we perform a suitable change of basis such that this MPO is written in the basis of Pauli matrices, and then consider the coefficients in the Pauli transfer matrix representation of the noise channel

$$c_{ij} = \frac{1}{2^n} \text{Tr}[P_i \mathcal{N}[P_j]], \quad (21)$$

where $P_{i,j}$ are n -qubit Pauli operators.

In Fig. 7, we report some values of these coefficients for both the true and reconstructed noise channels, for the different noise models introduced in Sec. III defined on $n = 10$ qubits. As it is clearly unfeasible to investigate all the 4^{2n} Pauli coefficients c_{ij} , we restrict our analysis to diagonal terms ($i = j$), and report data for some randomly sampled Pauli strings

having different Pauli weight (number of nonidentities). Note that while for the incoherent brickwork depolarizing (18) and sparse Pauli-Lindblad noise models (15) the Pauli Transfer matrix is indeed diagonal, this is not the case for coherent brickwork depolarizing channel which also has nondiagonal elements.

First of all, we note that all the reconstructed noise channels display the correct necessary behavior for trace preservation, as the coefficient belonging to the zero-weight Pauli string (all identities) is correctly normalized to one, as it holds $\text{Tr}[\mathbb{I} \mathcal{N}[\mathbb{I}]]/2^n = 1$. More importantly, we observe that the errors in the learned noise coefficients are relatively small, with a typical error of order 10^{-3} in all cases analyzed. Interestingly, we note that larger coefficients belonging to low-weight Pauli strings and larger noise levels are easier to learn, a fact, which we will investigate more deeply in future studies. We have also checked that the accuracy of the method still holds when characterizing stronger noise, as discussed in detail in Appendix D.

Overall, these results provide an additional and more direct evidence of the potential of the tensor network approach to characterize noisy processes.

V. APPLICATION TO ERROR MITIGATION

Despite their broad applicability and straightforward definition, distance measures like the norm in Eq. (11) may not be directly relevant for practical scenarios when one is interested in applying error mitigation techniques using characterized noise. For example, when it comes to calculating rigorous bounds for, e.g., estimation errors in experiments, the use of distances may lead to very loose bounds of little practical use [23].

In this section, we test the proposed noise learning procedure on the very timely task of error mitigation, showing how the proposed approach is able to provide accurate enough descriptions of the noise processes to achieve good noise-free estimates of expectation values when used in tandem with error mitigation techniques.

A. Tensor network error mitigation strategy

The error mitigation strategy we adopt in this paper is the tensor network error mitigation (TEM) algorithm recently introduced by some of the authors [23]. TEM relies on the (ideally perfect) characterization of the noise channels that affect the quantum circuit. This characterization is then employed to invert and cancel the effect of the noise channels, in the same spirit as in one of the most successful methods for quantum error mitigation, probabilistic error cancellation (PEC) [7,21,22]. At variance with PEC, TEM is applied completely in postprocessing and, moreover, it provides a quadratic advantage in the sampling overhead with respect to the former [23]. It also provides a sampling advantage with respect to zero-noise extrapolation with probabilistic error amplification (ZNE-PEA). In fact, for specific cases, it can be shown that its sampling overhead is optimal [57].

Suppose that the circuit we want to run on the quantum computer is composed of M layers represented by the ideal

unitaries

$$C_{\text{ideal}} = U_M \circ \dots \circ U_1. \tag{22}$$

However, as a result of inevitable noise in the quantum processor, the evolution we implement on hardware is instead given by

$$C_{\text{noisy}} = \mathcal{N}_M \circ U_M \circ \dots \circ \mathcal{N}_1 \circ U_1, \tag{23}$$

where \mathcal{N}_j is the noise channel associated with the ideal unitary operation U_j in the j th layer. After running the noisy circuit on hardware and obtaining the final outcome through a proper measurement procedure, our goal is to improve the accuracy of the outcome by mitigating the detrimental effect of the noise channels \mathcal{N}_j . The way we can achieve this through TEM is the following.

First, we characterize the noise channels \mathcal{N}_j in tensor network formalism. This characterization should be as accurate as possible and, crucially, we should be able to characterize the same layers we are using during the actual execution of the quantum circuit in Eq. (23). That is, the noise on the hardware should not change in the time between characterization and execution. Then, by computing the inverse of the noise channels \mathcal{N}_j^{-1} (see Appendix F for more details), we can finally postprocess the informationally complete measurement results obtained from the noisy state by applying the nonphysical map

$$C_{\text{TEM}} = U_M \circ \dots \circ U_1 \circ U_1^{-1} \circ \mathcal{N}_1^{-1} \circ \dots \circ U_M^{-1} \circ \mathcal{N}_M^{-1}, \tag{24}$$

for which it is easy to see that $C_{\text{TEM}} \circ C_{\text{noisy}} = C_{\text{ideal}}$, that is, we recover the ideal output.

The mitigation map C_{TEM} is represented as a tensor network and it is thus computed, i.e., contracted, on a classical computer. If C_{TEM} were as complex—from a tensor network perspective—as C_{ideal} , then TEM would not be of any use, since we would only be able to mitigate noise through classical tensor network methods if we were also able to directly compute the evolution driven by C_{ideal} through the same techniques. The core idea of TEM, however, is that only the inverse of the aggregated noise in the circuit must be classically simulated. If the noise in the channels \mathcal{N}_j is small enough, then the postprocessing map approaches the identity operator $C_{\text{TEM}} \approx \mathcal{I}$, and thus its contraction can be computed efficiently through tensor network methods, even if we are dealing with a large number of qubits. We refer the interested readers to the original paper [23] for more details and discussions about TEM.

B. Numerical results

To test the noise characterization method, we numerically simulate a noise mitigation experiment on $n = 10$ qubits in which we employ the noise channel returned by the characterization protocol together with TEM to mitigate the noisy circuit depicted in Fig. 8 (left).

The circuits we analyze consist in a repeated structure of operations sampled from Clifford gates, which allows for an easy computation of ideal noise-free expectation values from the circuit [5]. In order to study the accumulation of errors in deep circuits resulting from noise happening on several




FIG. 8. (Left) Schematics of the noisy Clifford circuits we consider into account. One ideal layer consists of a random single-qubit Clifford operations (green squares) followed by a (even or odd) CNOT layer. Each ideal layer is followed by the noise channel \mathcal{N} , which is a sparse Pauli-Lindblad noise channel introduced in Sec. III A, with coefficients sampled in order to resemble real experiments on IBM computers, see Appendix C for details. (Right) Results of the numerical experiment of error mitigation applied to the noisy circuit on the left. The expectation value of a different Pauli stabiliser at each circuit depth is shown, for either the unmitigated noisy circuit (gray diamonds), the mitigated circuit through TEM based on the true noise channel (blue crosses), and the mitigated circuit through TEM based on the reconstructed noise channel (orange circles) using a training set of 10^6 samples. (Inset) Mismatch between the mitigated results and the true result (equal to 1 for all depths), for either TEM based on the true noise channel (blue line) or the TEM based on the reconstructed noise channel (orange line). The blue line is different from zero because of the bond dimension truncation in the TEM method only. The orange line, in contrast, comprehends errors arising from the inaccuracy in the channel reconstruction (dominant contribution), the bond dimension truncation, and also errors in the inversion procedure of the MPO representing the noise necessary to run TEM. Nonetheless, the characterization procedure is able to provide a remarkably accurate description of the noise, so that TEM is able to provide almost ideal noise-mitigated values even at large depth.

computational layers, we run the tensor mitigation strategy on several circuits of different depths, obtained by iteratively appending additional layers one after another.

One step of the ideal (noise-free) circuit comprises a layer of random single-qubit Clifford gates followed by a layer of CNOTs, with the CNOT gates in each layer acting either on even or odd links between the qubits, depending on the step. Note that such alternating brickwork circuit structure is of practical interest as it can be used, for example, to study properties of many-body quantum systems via Trotterized evolution, see e.g., [1,58]. At the end of the circuit, we assume we are measuring the stabilizing Pauli operator O having expectation value $\langle O \rangle = +1$, which can be calculated by evolving the initial Pauli string $Z^{\otimes n}$, whose $+1$ eigenstate is the initial ground state $|0\rangle^{\otimes n}$ of the computation, with the Clifford operations in the circuit.

To take noise into account, we assume that each ideal circuit layer is followed by a noise channel \mathcal{N} , the effects of which we aim to mitigate through TEM in post-processing. For these experiments, we set the noise channel \mathcal{N} to be a sparse Pauli-Lindblad noise (15) with the coefficients as in Appendix C, sampled to resemble publicly available data by IBM on recent experiments leveraging SPL noise models [1,7]. Notably, as discussed in Sec. III A, since such noise is also a Clifford map, its effects on the output of the circuit can be computed efficiently.

We perform the TEM experiments with both the exact noise model used in the noisy simulations, and with the noise

model obtained with the characterization procedure using a total of 10^6 random measurement shots, as in Fig. 4. In order to use the characterized LPDO of the noise with TEM (24), we first transform it into an MPO and then compute its inverse by combining the explicit linear-algebra-based approach proposed in [29] together with an additional variational minimization. We refer to Appendixes A 2 and F for further details on the LPDO to MPO transformation and inversion of an MPO, respectively. In the simulations below the bond dimension of the MPO used to represent the tensor error mitigation map (24) is $\chi = 200$.

As discussed before, the mean values of the different Pauli operators O considered in the right panel of Fig. 8 are always equal to $+1$ for the ideal noise-free circuits, as for each step we are measuring the Pauli operator stabilized by that circuit. The same expectation values but for the noisy circuit are also shown (grey diamonds) up to 30 steps, with the signal almost vanishing at the last step.

In the right panel of Fig. 8, we show the TEM-mitigated results of the mean values of the Pauli operators using as an input for TEM either the true noise channel (blue crosses), which we can perfectly know only in a numerical experiment, or the reconstructed noise channel obtained through tensor-network-based noise characterization (orange circles). The difference between the blue crosses and the ideal value $+1$ arises from the truncation of the bond dimension in the TEM method and, as shown in the inset of Fig. 8 (right), it is small and noticeable only for higher circuit depths. In

other words, for all practical purposes, TEM reproduces the exact ideal result in the first circuit steps. In contrast, there is a visible difference between the ideal noise-free estimates and the mitigated ones obtained with the characterized noise model. However, in the inset, we observe that such mismatch is always of the order of 10^{-2} and, importantly, it does not increase with the circuit depth, so we are able to recover an almost perfect result even at step 30, where the noise almost wiped out the signal entirely. This is remarkable, given that three different sources of error are at play at the same time: (i) reconstruction error inherent to the noise learning procedure, (ii) errors in the inversion of the MPO of the characterized noise, and eventually (iii) truncation errors introduced by TEM to compute the tensor network mitigation map (24), with the first one dominating over the other two.

Our results thus show that the tensor-network-based noise characterization scheme studied in this paper can provide an accurate description of the noise even with a modest number of training experimental data, with direct applications in error mitigation techniques that rely on the knowledge of the noise.

VI. CONCLUSIONS

Accurate noise characterization is of utmost importance for attaining the best performance out of near-term quantum computers, and especially for state-of-the-art error mitigation methods, many of which rely on accurate knowledge of the noisy gates physically applied [7,23]. Standard process tomography [5], however, is unfeasible in the era of quantum utility, as circuit layers with tens or hundreds of qubits would require a huge amount of tomographic resources (e.g., experimental setups, state preparations, measurement shots, etc.), the scaling of which is exponential in the number of qubits.

In this paper, we propose a protocol for noise characterization based on the tensor network procedure introduced by Torlai *et al.* [16]. Our method not only avoids the exponential scaling of measurement resources by sampling the different possible tomographic settings in a randomized way, but it also enables an efficient, meaningful, and scalable description of the reconstructed noise channel by means of tensor network techniques (more specifically, a locally purified density operator structure, LPDO) with low bond dimension. The investigated method does not require any twirling of the noise maps [53], and is therefore suited to learn generic noisy processes. As the output of our protocol is a tensor network representation of the noise channel, it can be directly used as an input for the tensor network error mitigation (TEM) algorithm recently introduced in [23].

Whereas the original proposal in [16] mainly focused on learning unitary processes coming from arbitrarily deep quantum circuits, the originality of our approach lies in specializing the channel tomography technique to the case of learning shallow noise maps that accompany imperfect circuit layer instructions. This makes the procedure practical, as it requires low bond dimensions, and highly relevant for many noise-aware mitigation protocols. Additionally, we extensively tested the method in several scenarios that are experimentally relevant, including the effect of SPAM errors. We also compared our method with a similar proposal based on a local tomographic strategy, which demonstrated worse

performance. Finally, we tested the protocol for error mitigation, which is crucial for the success of near-term quantum computation and the primary reason for needing noise characterization

Our protocol was tested through several numerical experiments for realistic multiqubit correlated noise model learning. We specifically addressed three different noise channels that are of great relevance for current quantum computation, namely the sparse Pauli-Lindblad noise model [7], the incoherent depolarizing noise model with crosstalk, and the depolarizing noise model with crosstalk and coherent errors.

To assess the accuracy of the reconstruction, we used two figures of merit: the Frobenius distance between the true and reconstructed LPDOs, and the difference between the elements of the true and reconstructed superoperators expressed in the Pauli basis. We found that a limited and experimentally feasible number of shots (around 10^6 per characterization experiment) suffices for accurate noise channel characterization. We also explored how accuracy scales with the number of shots and qubits, observing favorable linear behavior in both cases.

Importantly, we have also tested the efficacy of the method in the presence of SPAM errors, demonstrating its resilience against small errors and how it can be combined with existing detector tomography techniques to mitigate undesired measurement errors and retain good channel reconstruction accuracies.

Moreover, we benchmarked the method with the timely and relevant task of quantum error mitigation. Specifically, we used the output of the noise characterization procedure as input for the TEM protocol to mitigate measurement outcomes from a noisy quantum circuit. We showed that TEM with characterized noise is able to provide mitigated expectation values with good accuracy (relative error of the order of 10^{-2}), even on deep circuit instances with tens of layers.

Summarizing, our analysis suggests that the tensor network noise characterization protocol may be a valuable tool for error mitigation for near-term quantum computers. The accuracy of our method is corroborated by the precision with which we can both reconstruct the noise channel as a tensor network and recover the ideal result of a noisy circuit when we employ this channel in conjunction with TEM.

ACKNOWLEDGMENTS

The authors would like to thank C. J. Wood for interesting discussions about noise characterization on near-term quantum computers. The code for the noise characterization method based on tensor networks presented in this paper is integrated into *Aurora*, a proprietary quantum chemistry platform developed by Algorithmiq Ltd.

G.G.P. proposed the idea and supervised the project and together with D.C., M.C., and S.M. conducted the research. S.M. wrote the main numerical code for this paper. M.C. and S.M. ran numerical simulations of this code on classical computers and wrote the first draft of the manuscript. M.R. and S.F. provided the code and support for TEM mitigation, S.F. provided support for TN parametrization. All authors contributed to scientific discussions and to the writing of the manuscript.

APPENDIX A: TENSOR NETWORK DETAILS

1. Explicit expression of the quantum channel in tensor network notation

The diagrammatic representation of the LPDO in Fig. 2 can be written in tensor network notation with explicit indices as

$$\begin{aligned} [\Lambda_{\mathcal{N}}]_{a,a'}^{b,b'} &= \sum_{\mu_1, \dots, \mu_{n-1}} \sum_{\nu_1, \dots, \nu_{n-1}} \sum_{\kappa_1, \dots, \kappa_n} \prod_{j=2}^{n-1} \\ &\times [A_1]_{b_1 a_1}^{\mu_1 \kappa_1} [A_1^\dagger]_{a'_1 b'_1}^{\nu_1 \kappa_1} \dots [A_j]_{b_j a_j}^{\mu_{j-1} \mu_j \kappa_j} [A_j^\dagger]_{a'_j b'_j}^{\nu_{j-1} \nu_j \kappa_j} \dots \\ &\times [A_n]_{b_n a_n}^{\mu_{n-1} \kappa_n} [A_n^\dagger]_{a'_n b'_n}^{\nu_{n-1} \kappa_n}. \end{aligned} \quad (\text{A1})$$

The coefficients $[A_j]_{b_j a_j}^{\mu_{j-1} \mu_j \kappa_j}$ are obtained by choosing a basis for the single-qubit operators, and then decomposing each global (i.e., acting on all the qubits) Kraus operator K_κ into a linear combination of tensor products of single-qubit operators, as given by Eq. (5). One general approach to split each Kraus operator into single-qubit operators (and the corresponding index κ into local indexes $\kappa_1, \dots, \kappa_n$) is based on a recursive application of the singular value decomposition, and the choice of $[A_j]_{b_j a_j}^{\mu_{j-1} \mu_j \kappa_j}$ depends on the specific decomposition one applies, see e.g., Ref. [17]. In practical cases, however, the local decomposition is evident from the structure of the channel under investigation. For example, a general two-qubit Pauli channel reads

$$\begin{aligned} \mathcal{P}[\rho] &= \sum_{\kappa=1}^{16} c_\kappa P_\kappa \rho P_\kappa \\ &= \sum_{\kappa_1, \kappa_2=1}^4 c_{\kappa_1 \kappa_2} (P_{\kappa_1} \otimes P_{\kappa_2}) \rho (P_{\kappa_1} \otimes P_{\kappa_2}), \end{aligned} \quad (\text{A2})$$

with $P_{\kappa_i} \in \{\mathbb{I}, X, Y, Z\}$ being single-qubit Pauli matrices. Such channel has Kraus operators $K_{\kappa_1 \kappa_2} = \sqrt{c_{\kappa_1 \kappa_2}} P_{\kappa_1} \otimes P_{\kappa_2}$, which have a clear local structure. Starting from such decomposition, one can realize that a LPDO representation of the channel as in Eq. (5) is achievable starting from the local tensors

$$\begin{aligned} [A^{[1]}]_{\mu \kappa_1} &= \begin{bmatrix} \mathbb{I} & 0 & 0 & 0 \\ 0 & X & 0 & 0 \\ 0 & 0 & Y & 0 \\ 0 & 0 & 0 & Z \end{bmatrix}, \\ [A^{[2]}]_{\mu \kappa_2} &= \begin{bmatrix} \sqrt{c_{11}} \mathbb{I} & \sqrt{c_{12}} X & \sqrt{c_{13}} Y & \sqrt{c_{14}} Z \\ \sqrt{c_{21}} \mathbb{I} & \sqrt{c_{22}} X & \sqrt{c_{23}} Y & \sqrt{c_{24}} Z \\ \sqrt{c_{31}} \mathbb{I} & \sqrt{c_{32}} X & \sqrt{c_{33}} Y & \sqrt{c_{34}} Z \\ \sqrt{c_{41}} \mathbb{I} & \sqrt{c_{42}} X & \sqrt{c_{43}} Y & \sqrt{c_{44}} Z \end{bmatrix}. \end{aligned} \quad (\text{A3})$$

Note, however, that such representation is nonunique, because of the inherent gauge freedom of tensor networks (for example, exchanging the two local tensors $A^{[1]} \leftrightarrow A^{[2]}$ give rise to the same channel). More complex channels arising from combinations of single- and two-qubit channels—as the ones discussed in the main text—can be obtained by combining and contracting together the LPDO representation of each of these channels.

2. From locally purified density operators (LPDO) to matrix product operators (MPO)

In Fig. 9(a) it is represented the LPDO representation of the quantum channel \mathcal{N} we want to characterize. In many applications (e.g., for running the TEM algorithm) we need the superoperator representation of \mathcal{N} in the Liouville space, which consists of transforming the channel into a matrix acting on the vectorized space of density matrices [27].

In the superoperator formalism, the action of some Kraus operators K_κ acting on the state ρ like $K_\kappa \rho K_\kappa^\dagger$ is represented as $K_\kappa \otimes (K_\kappa^\dagger)^T |\rho\rangle\rangle$, where $|\rho\rangle\rangle$ is a suitable vectorization of the density matrix. Then, the superoperator associated with \mathcal{N} in the tensor network formalism can be easily obtained from the LPDO structure as shown in Figs. 9(b) and 9(c): the indices of the Kraus operators acting on the left and on the right of the density matrix are suitably reshuffled and then merged to create a MPO representing \mathcal{N} [27]. Additionally, for the sake of running the tensor error mitigation (TEM) algorithm described in Sec. V A, we need a MPO superoperator representation of \mathcal{N} in the Pauli basis. To do this we apply on each site of the MPO a local change-of-basis unitary operator that transforms the computational basis into the desired Pauli basis, as depicted in Fig. 9(d). Finally, the MPO we will invert to run TEM is shown in Fig. 9(e).

APPENDIX B: LOCAL SAMPLING STRATEGIES

The data sampling employed in this paper is based on the random strategy described in Sec. II B. We have also explored a different strategy that assumes that the correlations between different sites of the tensor network representing the quantum channel \mathcal{N} are only ℓ local, thus focusing on the reconstructing of ℓ -reduced channels only. This strategy is motivated by the similar methods that has been successfully applied to the state tomography of matrix product states (MPS) [44,45] and mixed states [43].

There is, however, a fundamental difference between the local strategy for state tomography and for process tomography. Suppose that we are employing the Pauli measurements (i.e., three different measurement bases per qubit, as discussed in the Sec. II B) for ℓ -qubit state tomography; then we need 3^ℓ different experimental tomographic settings, corresponding to six different POVM outcomes. As discussed in Sec. II B, for ℓ -qubit process tomography, in contrast, even in the best possible scenario we need 12^ℓ settings in order to take into account also the preparation of informationally complete input states. This means that the number of experimental settings grows much faster than for state tomography as a function of the locality ℓ . For a real experiment on current near-term quantum computers with limited access and capabilities, it is already quite difficult to gather statistics on pretty low locality, for example $\ell = 4$ implying 12^4 different experimental settings (quantum circuits), and absolutely unfeasible to reach locality $\ell = 5$.

Fixing the value of the locality ℓ , a basic local sampling strategy can be implemented by preparing all the possible tomographic settings on subsets of ℓ qubits. Specifically, as described in Sec. II B, if we use a set of R informationally complete input states and Pauli measurements, we will need

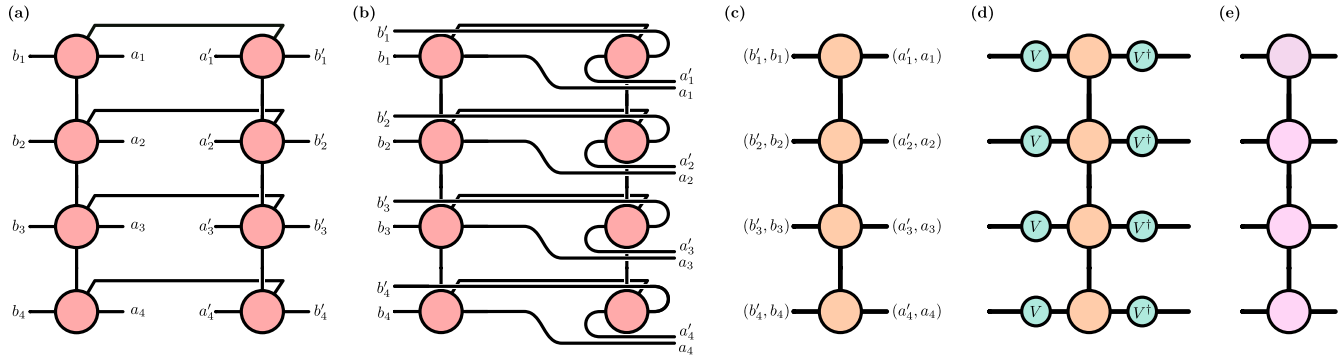


FIG. 9. (a) LPDO representation of a quantum channel \mathcal{N} , analogously to Fig. 2. (b) Indices are reshuffled according to the superoperator representation. (c) The reshuffled indices are merged to give rise to a MPO representing \mathcal{N} as a superoperator. (d) A change of basis transformation is applied locally on each qubit to write the MPO in the Pauli basis (starting from the computational one). (e) Final MPO representation of \mathcal{N} in the Pauli basis.

to execute at least $(3R)^\ell$ different experimental settings. In fact, for a linear chain of qubits, one can see that by using a scheme of correlated preparations and measurements comprising $(3R)^\ell$ settings is enough to provide the necessary ℓ -local tomographic data on all ℓ local reduced channels on neighboring qubits.

For instance, suppose we want to characterize five qubits and we choose locality $\ell = 3$. Consider experimental settings where the qubits are prepared in a correlated fashion, that is with input states of the form $\rho = \rho_A \otimes \rho_B \otimes \rho_C \otimes \rho_A \otimes \rho_B$, where $\rho_{A,B,C}$ are sampled from a set of IC states; and also measured on correlated bases, that is with measurement operators of the form $P = P_A \otimes P_B \otimes P_C \otimes P_A \otimes P_B$, where $P_{A,B,C}$ are Pauli operators. Using such correlated tomographic scheme, by considering all the $(3R)^3$ tomographic settings obtained by considering all combinations of states $\rho_{A,B,C}$ and measurements $P_{A,B,C}$, one covers the experimental settings needed to reconstruct all the 3-reduced channels acting on subsets of qubits $\{0, 1, 2\}$, $\{1, 2, 3\}$, and $\{2, 3, 4\}$. This is because if states $\rho_A \otimes \rho_B \otimes \rho_C$ on qubits $\{0, 1, 2\}$ spans all R^3 possibilities, then also $\rho_B \otimes \rho_C \otimes \rho_A$ on qubits $\{1,2,3\}$ will span all possibilities, and similarly for $\rho_C \otimes \rho_A \otimes \rho_B$ on qubits $\{2,3,4\}$. Same goes for the measurement operators. This optimal scheme holds for a linear chain of qubits, for more complex topologies the choice of settings may be different [59].

Alternatively, we may implement the local strategy by keeping a different locality for the input states and the measurement bases. This is motivated by *lightcone* arguments. That is, if we want to characterize all the ℓ -local outcomes (i.e., we characterize up to ℓ -local correlations in the measurements over the n qubits), then the outcomes over ℓ qubits will in general depend on the input states over more than ℓ qubits, depending on the entangling structure of the channel. For instance, consider a single layer of noisy CNOTs in which crosstalk errors affect only the first neighbors, which is the case treated in the main text and depicted in Fig. 1(a). Then, it is easy to see that the outcomes on a single qubit can be influenced by the initial states of at most four qubits. If we aim to characterize 2-local outcomes instead, these will depend on the input state of at most six qubits. This analysis

tells us immediately that the scaling of this lightcone-based local strategy is again quite unfavorable: in the best scenario, we need $4^6 \times 3^2 = 36\,864$ settings for exactly characterizing 2-local correlations in the measurement outcomes, which is hardly feasible on current near-term quantum computers.

Independent of the chosen strategy to collect local tomographic data, one could then still use the same machinery discussed in Sec. II C to train a tensor network for the total channel \mathcal{N} starting, however, from tomographic data on the ℓ -local channels. Of course, as a general n -qubit quantum channel cannot be written in terms of products of ℓ -local ones, the reconstruction accuracy of the whole channel will be impacted, with good accuracy reached only when the experiments locality used to collect the tomographic data approaches the actual locality of the channel [44,45].

For completeness, in Fig. 10 we report some numerical results obtained by training the LPDO on local data of different locality ℓ obtained using the basic sampling strategy described above, to learn the brickwork depolarizing channel (Sec. III B). As clear from the picture, the accuracy improves when considering a larger locality for the data sampling, but the global strategy still results in better reconstruction performances despite using a smaller number of experimental settings.

APPENDIX C: SPARSE PAULI-LINDBLAD NOISE MODEL

In Fig. 11 we report the coefficients used in the experiments involving the sparse Pauli-Lindblad noise model, as defined in Eq. (15). Such coefficients were sampled randomly to match publicly available data by IBM on noise characterization procedures run on superconducting quantum hardware [1,7]. Whenever we consider instances of such noise model on systems with less than 20 qubits ($n < 20$), we proceed by simply restricting the noise model to those Pauli-Lindblad operators, which act nontrivially on qubits $q \in \{0, \dots, n-1\}$.

APPENDIX D: STRONG DEPOLARIZING NOISE

In this Appendix, we report numerical results for the characterization of a stronger noise channel, namely the brickwork

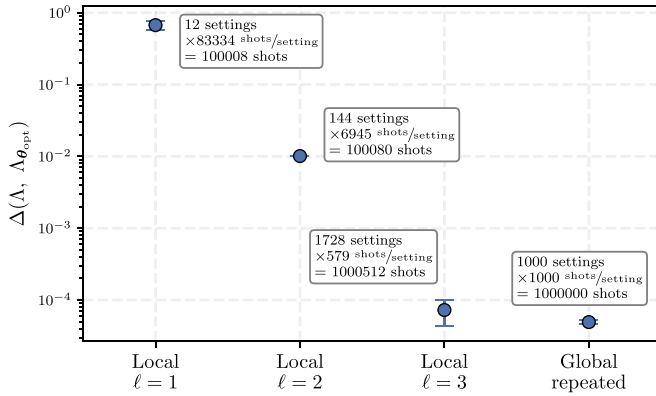


FIG. 10. Characterizing the brickwork depolarizing channel with local data. Frobenius distance between true and reconstructed noise channels through the local strategy, for different localities $\ell \in \{1, 2, 3\}$. The number of shots per setting in each scenario is tuned so that all characterization experiments use roughly the same total number of shots ($\approx 10^6$). The results are compared with the global repeated strategy used in the main text (see Sec. II B), consisting of 10^3 settings and 10^3 shots per setting, for a total of 10^6 total shots. The reconstruction error improves by considering larger localities, but the global strategy achieves better performances. Each point in the plot is obtained as the mean value of three different training runs of the LPDO initialized with different parameters, with the error bars being the standard deviation.

depolarizing channel of Eq. (18) but with noise parameter $p = 10^{-1}$, as opposed to $p = 10^{-3}$ used in the main text. Note that such error rate is widely larger than those found in already available state-of-the-art quantum computers.

In Fig. 12 we show the scaling of the Frobenius distance between the true and reconstructed channel as a function of the number of shots, and compare it with the other noise models we have explored in this paper. As argued in the main text, we witness a clear dependence of the reconstruction error on the noise intensity, which can be understood as a consequence of the whole learning procedure being unable to distinguish the signal from a background white noise.

In addition, in Fig. 13 we report some coefficients of the true and reconstructed noise channel in the MPO representation and in the Pauli basis. Despite the lesser performance in terms of channel distance, we observe that the accuracy of our reconstruction procedure is, however, once again remarkable even in the strong noise scenario.

APPENDIX E: INITIALIZATION AND OPTIMIZATION OF THE TENSOR NETWORK

In this Appendix we discuss the custom random initialization of the LPDO tensor network and provide details on the optimization routines to train it.

1. Initialization of the LPDO parameters

The tensor elements in the LPDO Λ_θ are initialized as random complex variables with Gaussian distributed real

and imaginary part, namely

$$\theta_k = \text{Re}(\theta_k) + i \text{Im}(\theta_k)$$

$$\text{with } \text{Re}(\theta_k), \text{Im}(\theta_k) \sim \mathcal{N}(0, \sigma^2). \quad (\text{E1})$$

Given such choice, it is possible to compute the expectation value of the trace of the LPDO upon initialization, which amounts to

$$\mathbb{E}_\theta[\Lambda_\theta] = (8\sigma^2 \chi_\kappa)^n \chi_b^{n-1}, \quad (\text{E2})$$

where n is the number of qubits, and χ_κ and χ_b are the Kraus bond dimension and the virtual bond dimension of the LPDO, respectively. Then, by setting the variance to be $\sigma^2 = 2/(8\chi_\kappa \chi_b^{1-1/n})$ one has that, in expectation value upon random initialization, the LPDO is properly normalized to the correct value $\mathbb{E}_\theta[\text{Tr}[\Lambda_\theta]] = 2^n$.

In what follows we show how to derive Eq. (E2), with the idea of the proof being diagrammatically depicted in Fig. 14. We first start by computing expectation values of the form $\mathbb{E}[\text{Tr}[A A^\dagger]]$, where A is a random matrix with normally distributed real and imaginary parts, and then proceed to show how the trace of the whole LPDO results in a composition of such quantities.

Let $A \in \mathbb{C}^{2 \times 2}$ be a complex random normal matrix whose entries are identically independently distributed (*iid*) variables according to Eq. (E1). Then it holds

$$\begin{aligned} \mathbb{E}[\text{Tr}[A A^\dagger]] &= \mathbb{E}\left[\text{Tr}\left[\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} a^* & c^* \\ b^* & d^* \end{pmatrix}\right]\right] \\ &= \mathbb{E}[|a|^2 + |b|^2 + |c|^2 + |d|^2] \\ &= 4 \mathbb{E}[|a|^2] = 4 \mathbb{E}[\text{Re}(a)^2 + \text{Im}(a)^2] \quad (\text{E3}) \\ &= 8\sigma^2 \end{aligned}$$

where in the third line we first made use of the fact that the matrix elements are *iid*, and secondly that the real and imaginary parts satisfy $\mathbb{E}[\text{Re}(a)^2] = \mathbb{E}[\text{Im}(a)^2] = \sigma^2$. If instead one considers the product of two different independent random matrices A and B , then it is easy to show that $\mathbb{E}[\text{Tr}[AB]] = 0$.

We now proceed computing the trace of the LPDO tensor network $\text{Tr}[\Lambda_\theta]$, which is diagrammatically shown in Fig. 14(a). Let $A_{\mu, \nu, \kappa}^{[q]} \in \mathbb{C}^{2 \times 2}$ denote the local Kraus tensors acting on each site, then starting from the definition of the LPDO given in Eq. (4) and tracing over the physical indices at each site of the tensor network results in

$$\begin{aligned} \text{Tr}[\Lambda_\theta] &= \sum_{\mu, \nu, \kappa} \text{Tr}\left[\left(A_{\mu_1, \kappa_1}^{[1]} \otimes A_{\mu_1, \mu_2, \kappa_2}^{[2]} \otimes \dots \otimes A_{\mu_{n-1}, \kappa_n}^{[n]}\right)\right. \\ &\quad \left. \times \left(A_{\nu_1, \kappa_1}^{[1]\dagger} \otimes A_{\nu_1, \nu_2, \kappa_2}^{[2]\dagger} \otimes \dots \otimes A_{\nu_{n-1}, \kappa_n}^{[n]\dagger}\right)\right] \\ &= \sum_{\mu, \nu, \kappa} \text{Tr}\left[A_{\mu_1, \kappa_1}^{[1]} A_{\nu_1, \kappa_1}^{[1]\dagger}\right] \text{Tr}\left[A_{\mu_1, \mu_2, \kappa_2}^{[2]} A_{\nu_1, \nu_2, \kappa_2}^{[2]\dagger}\right] \dots \\ &\quad \times \text{Tr}\left[A_{\mu_{n-1}, \kappa_n}^{[n]} A_{\nu_{n-1}, \kappa_n}^{[n]\dagger}\right] \\ &= \sum_{\mu, \nu} A_{\mu_1, \nu_1}^{[1]} A_{\mu_1, \nu_1, \mu_2, \nu_2}^{[2]} \dots A_{\mu_{n-1}, \nu_{n-1}}^{[n]}, \quad (\text{E4}) \end{aligned}$$

where in the last line we introduced local tensors coming from the contraction of the Kraus indices κ and the trace over the




FIG. 11. Coefficients of the sparse Pauli-Lindblad noise model (15) considered in the numerical experiments, defined on a maximum of $n = 20$ qubits with linear connectivity. We remark that these values are realistic noise coefficients sampled according to publicly available data by IBM on noise characterization run on real superconducting quantum hardware.

physical indices at each site, namely

$$\begin{aligned} A_{\mu,v}^{[q]} &:= \sum_{\kappa=1}^{\chi_{\kappa}} \text{Tr}[A_{\mu,\kappa}^{[q]} A_{v,\kappa}^{[q]\dagger}] \quad q \in \{1, n\}, \\ A_{\mu,v,\mu',v'}^{[q]} &:= \sum_{\kappa=1}^{\chi_{\kappa}} \text{Tr}[A_{\mu,\mu',\kappa}^{[q]} A_{v,v',\kappa}^{[q]\dagger}] \quad q \in [2, n-1]. \end{aligned} \quad (\text{E5})$$

When each matrix $A_{\mu,v,\kappa}^{[q]} \in \mathbb{C}^{2 \times 2}$ in (E6) is a random matrix with normal complex entries, then one can use Eq. (E3) to compute the expectation values

$$\mathbb{E}[A_{\mu,v}^{[q]}] = \sum_{\kappa=1}^{\chi_{\kappa}} \underbrace{\mathbb{E}[\text{Tr}[A_{\mu,\kappa}^{[q]} A_{v,\kappa}^{[q]\dagger}]]}_{8\sigma^2 \delta_{\mu\nu}} = 8\sigma^2 \chi_{\kappa} \delta_{\mu\nu},$$

and similarly for middle-sites tensors $q \in [2, n-2]$

$$\begin{aligned} \mathbb{E}[A_{\mu,v,\mu',v'}^{[q]}] &= \sum_{\kappa=1}^{\chi_{\kappa}} \underbrace{\mathbb{E}[\text{Tr}[A_{\mu,\mu',\kappa}^{[q]} A_{v,v',\kappa}^{[q]\dagger}]]}_{8\sigma^2 \delta_{\mu\nu} \delta_{\mu'v'}} \\ &= 8\sigma^2 \chi_{\kappa} \delta_{\mu\nu} \delta_{\mu'v'}, \end{aligned}$$




FIG. 12. Same as in Fig. 4 in the main text, but with the strong incoherent depolarizing noise channel introduced in Sec. III B with noise strength $p = 0.1$.

thus having in total

$$\begin{aligned} \mathbb{E}[A_{\mu,v}^{[q]}] &= 8\sigma^2 \chi_{\kappa} \delta_{\mu\nu}, \quad q \in \{0, n-1\}, \\ \mathbb{E}[A_{\mu,v,\mu',v'}^{[q]}] &= 8\sigma^2 \chi_{\kappa} \delta_{\mu\nu} \delta_{\mu'v'}, \quad q \in [1, n-2]. \end{aligned} \quad (\text{E6})$$

Note that we can consider expectation values independently on each local tensor $A^{[k]}$ since each tensor is drawn independently, and the expectation value factorizes over local sites, that is

$$\begin{aligned} &= \mathbb{E}[\text{Tr}[\Lambda_{\theta}]] \\ &= \mathbb{E}\left[\sum_{\mu,v} A_{\mu_1,v_1}^{[1]} A_{\mu_1,v_1,\mu_2,v_2}^{[2]} \cdots A_{\mu_{n-1},v_{n-1}}^{[n]}\right] \\ &= \sum_{\mu,v} \mathbb{E}[A_{\mu_1,v_1}^{[1]}] \mathbb{E}[A_{\mu_1,v_1,\mu_2,v_2}^{[2]}] \cdots \mathbb{E}[A_{\mu_{n-1},v_{n-1}}^{[n]}]. \end{aligned} \quad (\text{E7})$$

Finally, using Eqs. (E6) inside Eq. (E7), one eventually obtains

$$\begin{aligned} &= \mathbb{E}[\text{Tr}[\Lambda_{\theta}]] \\ &= \sum_{\mu,v} (8\sigma^2 \chi_{\kappa})^n \delta_{\mu_1 v_1} \delta_{\mu_1 v_1} \delta_{\mu_2 v_2} \cdots \delta_{\mu_{n-1} v_{n-1}} \\ &= (8\sigma^2 \chi_{\kappa})^n \chi_b^{n-1}, \end{aligned} \quad (\text{E8})$$

where in the last line we used

$$\begin{aligned} &\sum_{\substack{\mu_1, \mu_2, \dots, \mu_{n-1}=1 \\ v_1, v_2, \dots, v_{n-1}=1}}^{\chi_b} \delta_{\mu_1 v_1} \delta_{\mu_1 v_1} \delta_{\mu_2 v_2} \cdots \delta_{\mu_{n-1} v_{n-1}} \\ &= \sum_{\substack{\mu_1, \mu_2, \dots, \mu_{n-1}=1 \\ v_2, \dots, v_{n-1}=1}}^{\chi_b} \delta_{\mu_2 v_2} \cdots \delta_{\mu_{n-1} v_{n-1}} \\ &= \sum_{\mu_1, \dots, \mu_{n-1}=1}^{\chi_b} = \chi_b^{n-1}. \end{aligned} \quad (\text{E9})$$

2. Optimization details

The parameterized LPDO Λ_{θ} (A1) is trained with Adam optimizer [51] combined with an exponential decay of the




FIG. 13. Some coefficients of the MPOs in Pauli transfer matrix of the true (green circles) and reconstructed (orange crosses) noise channels, for the strong incoherent depolarizing noise channel introduced in Sec. III B, with $p = 0.1$. The numbers above the plot indicate the coefficient of the chosen Pauli operator (21). The mismatch between the green circles and orange crosses results from the error in channel tomography, obtained with a training set $N = 10^7$ samples.

learning rate, which was found to stabilize training and ensure good convergence towards the end of the optimization process.

Adam is a variant of stochastic gradient descent very common in machine learning research, and consists of the following update rules:

$$\begin{aligned}
 m_t &\leftarrow \beta_1 m_{t-1} + (1 - \beta_1) g_t \\
 v_t &\leftarrow \beta_2 v_{t-1} + (1 - \beta_2) g_t^2 \\
 \hat{m}_t &\leftarrow m_t / (1 - \beta_1^t) \\
 \hat{v}_t &\leftarrow v_t / (1 - \beta_2^t) \\
 \theta_t &\leftarrow \theta_{t-1} - \eta \hat{m}_t / (\sqrt{\hat{v}_t} + \varepsilon),
 \end{aligned}
 \tag{E10}$$

where $g_t = \nabla_{\theta} f(\theta_{t-1})$ is the gradient of the loss function $f(\theta)$ to be minimized having tunable parameters θ , g_t^2 indicates its element-wise square, and η is the step size (or learning rate). In our simulations we used standard values for the hyperparameters, $\beta_1 = 0.9$ and $\beta_2 = 0.999$, $\varepsilon = 10^{-8}$.

In addition to Adam, we used an exponential decay of the learning rate

$$\eta_t = \eta_0 \gamma^{t/T}
 \tag{E11}$$

where η_0 is the initial learning rate at the start of training, γ is the decay rate, t is the time step, and T is a decay time. In our simulations we used $\eta_0 = 10^{-2}$, $\gamma = 0.9$, and the decay time T was set equal to the number of training batches in an epoch, which depends on the number of tomographic samples N . The exponential decay starts only after a warm-up period of 500 gradient-descent steps.

Importantly, note that in our case the parameters to be optimized are elements of the Kraus operators (A1), and they consist of complex variables. Accordingly, the cost function is minimized by taking steps in the direction of the conjugated gradient [60]. All optimization runs, including Adam and the exponential decay of the learning rate, were implemented as provided by the jax-based optimization library optax [49].

As customary in machine learning, training was run by gradient-descent updates on mini batches of data of size 250 (50 when the number of tomographic samples is scarce $N = 10^3$). Of the whole tomographic dataset consisting of N measurement samples, $\min(N/10, 12500)$ of them were used as a test dataset to estimate the loss function. The stopping criterion used during training was to stop the optimization if the Frobenius distance between the optimized LPDO and the target one did not change more than 10^{-7} over the last five




FIG. 14. (Left) Tensor network representation of the trace of the LPDO parameterized by θ , which represents the noise channel. (Right) Tensor network representation of the expectation value of the trace of the LPDO upon initialization.

training epochs. In the realistic case scenario where one does not have the target LPDO to compare with, one can instead monitor the loss function on the test set and stop training if this stops improving.

APPENDIX F: INVERSION OF THE MPOS

In order to run the tensor error mitigation technique discussed in Sec. V, it is necessary to be able to compute the inverse of MPOs representing the quantum channels (which, however, is not a valid quantum channel [61]). That is, given a matrix product operator Γ , one needs to find another operator Γ^{-1} such that $\Gamma\Gamma^{-1} = \mathbb{I}$. As proposed in [43], this can be done by minimizing the error

$$\Delta_\phi = \|\Gamma \Upsilon_\phi - \mathbb{I}\|_F^2 \tag{F1}$$

where $\|\cdot\|_F$ is the squared Frobenius distance, Γ the MPO to be inverted, and Υ_ϕ is a parameterized MPO the tensor elements ϕ of which are tuned to approach Γ^{-1} . As shown in [43], this minimization problem can be reduced to a quadratic problem in the local tensors, and then solved by sweeping over the sites and solving local systems of linear equations at each of them.

In addition to such explicit method, the error term Δ_ϕ can also be minimized by variationally tuning the parameters by means of an optimizer. Indeed, in our simulations we noticed that a combined approach of these two methods provides better results, especially when the MPO to be inverted is not sparse and contains many nonzero but small entries, as is the case for the MPOs coming from the noise characterization procedure II B. Specifically, one can use a classical

optimization routine to find

$$\phi_{\text{opt}} = \arg \min \Delta_\phi = \arg \min_\phi \|\Gamma \Upsilon_\phi - \mathbb{I}\|_F^2, \tag{F2}$$

and the optimization task can be performed either *globally* by minimizing all the parameters in Υ_ϕ at the same time, or again in a DMRG-like [17] fashion by dividing it into many local subsequent optimization problems where only the parameters belonging to one single site are optimized at each time, with the others being fixed.

For the tensor error mitigation experiments on $n = 10$ qubits with the sparse Pauli-Lindblad noise reported in Sec. V, the exact noise maps—that is those built explicitly from the definition of the noise channels—were inverted with an MPO with virtual bond dimension $\chi_b = 4$ using the linear algebra based inversion procedure proposed in [43], which was found to converge to negligible inversion error $\Delta_\phi \lesssim 10^{-5}$. Instead, for the MPOs associated with the noise channels coming from the characterization procedure, the explicit inversion method, again with an ansatz MPO of bond dimension $\chi_b = 4$, converged to $\Delta_\phi \approx 6$, and was followed by a round of global and local variational minimization of the error function with optimizer L-BFGS-B as provided by quimb [47], which improves the inversion achieving a final error of $\Delta_\phi \approx 0.6$.

Even though the inversion of the MPOs is not perfect, especially for the characterized noise channels, we note that in our cases the error from the inversion procedure is usually much smaller than the one from the characterization procedure, as one can see by comparing the normalized inversion error $\Delta_{\phi_{\text{opt}}}/2^{2n} \approx 10^{-7}$ with the normalized characterization error $\Delta(\Lambda, \Lambda_{\theta_{\text{opt}}}) \approx 10^{-4}$ (see Fig. 4 with 10^7 shots). We leave a more comprehensive analysis of the inversion error and their impact of noise mitigation as a topic for future studies.

-
- [1] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, Evidence for the utility of quantum computing before fault tolerance, *Nature (London)* **618**, 500 (2023).
 - [2] Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, J. R. McClean, and T. E. O’Brien, Quantum error mitigation, *Rev. Mod. Phys.* **95**, 045005 (2023).
 - [3] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter *et al.*, Logical quantum processor based on reconfigurable atom arrays, *Nature (London)* **626**, 58 (2024).
 - [4] G. M. D’Ariano and P. L. Presti, Quantum tomography for measuring experimentally the matrix elements of an arbitrary quantum operation, *Phys. Rev. Lett.* **86**, 4195 (2001).
 - [5] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, 2010).
 - [6] A. Erhard, J. J. Wallman, L. Postler, M. Meth, R. Stricker, E. A. Martinez, P. Schindler, T. Monz, J. Emerson, and R. Blatt, Characterizing large-scale quantum computers via cycle benchmarking, *Nat. Commun.* **10**, 5347 (2019).
 - [7] E. van den Berg, Z. K. Mineev, A. Kandala, and K. Temme, Probabilistic error cancellation with sparse Pauli-Lindblad models on noisy quantum processors, *Nat. Phys.* **19**, 1116 (2023).
 - [8] J. Kunjummen, M. C. Tran, D. Carney, and J. M. Taylor, Shadow process tomography of quantum channels, *Phys. Rev. A* **107**, 042403 (2023).
 - [9] R. Levy, D. Luo, and B. K. Clark, Classical shadows for quantum process tomography on near-term quantum computers, *Phys. Rev. Res.* **6**, 013029 (2024).
 - [10] G. A. L. White, P. Jurcevic, C. D. Hill, and K. Modi, Unifying non-Markovian characterisation with an efficient and self-consistent framework, [arXiv:2312.08454](https://arxiv.org/abs/2312.08454).
 - [11] 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).
 - [12] 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).
 - [13] S. T. Flammia, D. Gross, Y.-K. Liu, and J. Eisert, Quantum tomography via compressed sensing: Error bounds, sample complexity and efficient estimators, *New J. Phys.* **14**, 095022 (2012).
 - [14] S. Ahmed, F. Quijandria, and A. F. Kockum, Gradient-descent quantum process tomography by learning Kraus operators, *Phys. Rev. Lett.* **130**, 150402 (2023).

- [15] T. Surawy-Stepney, J. Kahn, R. Kueng, and M. Guta, Projected least-squares quantum process tomography, *Quantum* **6**, 844 (2022).
- [16] 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).
- [17] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys. (NY)* **326**, 96 (2011).
- [18] R. Orús, A practical introduction to tensor networks: Matrix product states and projected entangled pair states, *Ann. Phys. (NY)* **349**, 117 (2014).
- [19] J. C. Bridgeman and C. T. Chubb, Hand-waving and interpretive dance: An introductory course on tensor networks, *J. Phys. A: Math. Theor.* **50**, 223001 (2017).
- [20] S. Montangero, *Introduction to Tensor Network Methods: Numerical Simulations of Low-Dimensional Many-Body Quantum Systems* (Springer International Publishing, New York, 2018).
- [21] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, *Phys. Rev. Lett.* **119**, 180509 (2017).
- [22] S. Endo, S. C. Benjamin, and Y. Li, Practical quantum error mitigation for near-future applications, *Phys. Rev. X* **8**, 031027 (2018).
- [23] S. Filippov, M. Leahy, M. A. C. Rossi, and G. García-Pérez, Scalable tensor-network error mitigation for near-term quantum computing, [arXiv:2307.11740](https://arxiv.org/abs/2307.11740).
- [24] M. Cattaneo, M. A. C. Rossi, K. Korhonen, E.-M. Borrelli, G. García-Pérez, Z. Zimborás, and D. Cavalcanti, Self-consistent quantum measurement tomography based on semidefinite programming, *Phys. Rev. Res.* **5**, 033154 (2023).
- [25] J. Watrous, *The Theory of Quantum Information* (Cambridge University Press, Cambridge, 2018).
- [26] M. M. Wilde, *Quantum Information Theory*, 2nd ed. (Cambridge University Press, Cambridge, 2017).
- [27] C. J. Wood, J. D. Biamonte, and D. G. Cory, Tensor networks and graphical calculus for open quantum systems, *Quantum Inf. Comput.* **15**, 759 (2015).
- [28] A. H. Werner, D. Jaschke, P. Silvi, M. Kliesch, T. Calarco, J. Eisert, and S. Montangero, Positive tensor network approach for simulating open quantum many-body systems, *Phys. Rev. Lett.* **116**, 237201 (2016).
- [29] Y. Guo and S. Yang, Quantum error mitigation via matrix product operators, *PRX Quantum* **3**, 040313 (2022).
- [30] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, Matrix product density operators: Simulation of finite-temperature and dissipative systems, *Phys. Rev. Lett.* **93**, 207204 (2004).
- [31] S. Srinivasan, S. Adhikary, J. Miller, G. Rabusseau, and B. Boots, [Towards a trace-preserving tensor network representation of quantum channels](#), in *Second Workshop on Quantum Tensor Networks in Machine Learning: 35th Conference on Neural Information Processing Systems (NeurIPS 2021)* (Curran Associates, Inc., Red Hook, New York, 2021).
- [32] J. Řeháček, B.-G. Englert, and D. Kaszlikowski, Minimal qubit tomography, *Phys. Rev. A* **70**, 052321 (2004).
- [33] S. T. Merkel, J. M. Gambetta, J. A. Smolin, S. Poletto, A. D. Córcoles, B. R. Johnson, C. A. Ryan, and M. Steffen, Self-consistent quantum process tomography, *Phys. Rev. A* **87**, 062119 (2013).
- [34] E. Nielsen, J. K. Gamble, K. Rudinger, T. Scholten, K. Young, and R. Blume-Kohout, Gate set tomography, *Quantum* **5**, 557 (2021).
- [35] R. Brieger, I. Roth, and M. Kliesch, Compressive gate set tomography, *PRX Quantum* **4**, 010325 (2023).
- [36] J. Helsen, M. Ioannou, J. Kitzinger, E. Onorati, A. H. Werner, J. Eisert, and I. Roth, Shadow estimation of gate-set properties from random sequences, *Nat. Commun.* **14**, 5039 (2023).
- [37] IBM Quantum Documentation, <https://docs.quantum.ibm.com/> (2023).
- [38] Qiskit contributors, Qiskit: An open-source framework for quantum computing, <https://www.ibm.com/quantum/qiskit> (2023).
- [39] S. Aaronson, Shadow tomography of quantum states, *SIAM J. Comput.* **49**, STOC18-368 (2020).
- [40] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [41] A. Acharya, S. Saha, and A. M. Sengupta, Shadow tomography based on informationally complete positive operator-valued measure, *Phys. Rev. A* **104**, 052418 (2021).
- [42] A. Acharya, S. Saha, S. Sridharan, Y. Bahrour, and A. M. Sengupta, Learning the eigenstructure of quantum dynamics using classical shadows, [arXiv:2309.12631](https://arxiv.org/abs/2309.12631).
- [43] Y. Guo and S. Yang, Scalable quantum state tomography with locally purified density operators and local measurements, [arXiv:2307.16381](https://arxiv.org/abs/2307.16381).
- [44] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, Efficient quantum state tomography, *Nat. Commun.* **1**, 149 (2010).
- [45] B. P. Lanyon, C. Maier, M. Holzäpfel, T. Baumgratz, C. Hempel, P. Jurcevic, I. Dhand, A. S. Buyskikh, A. J. Daley, M. Cramer *et al.*, Efficient tomography of a quantum many-body system, *Nat. Phys.* **13**, 1158 (2017).
- [46] 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).
- [47] J. Gray, quimb: A python package for quantum information and many-body calculations, *J. Open Source Soft.* **3**, 819 (2018).
- [48] 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, <https://github.com/google/jax> (2018).
- [49] DeepMind, I. Babuschkin, K. Baumli, A. Bell, S. Bhupatiraju, J. Bruce, P. Buchlovsky, D. Budden, T. Cai, A. Clark *et al.*, The DeepMind JAX Ecosystem, <https://github.com/google-deepmind> (2020).
- [50] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, A limited memory algorithm for bound constrained optimization, *SIAM J. Sci. Comput.* **16**, 1190 (1995).
- [51] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [52] S. Chen, Y. Liu, M. Otten, A. Seif, B. Fefferman, and L. Jiang, The learnability of Pauli noise, *Nat. Commun.* **14**, 52 (2023).

- [53] J. J. Wallman and J. Emerson, Noise tailoring for scalable quantum computation via randomized compiling, *Phys. Rev. A* **94**, 052325 (2016).
- [54] D. C. McKay, I. Hincks, E. J. Pritchett, M. Carroll, L. C. G. Góvia, and S. T. Merkel, Benchmarking quantum processor performance at scale, [arXiv:2311.05933](https://arxiv.org/abs/2311.05933).
- [55] S. A. Moses, C. H. Baldwin, M. S. Allman, R. Ancona, L. Ascarrunz, C. Barnes, J. Bartolotta, B. Bjork, P. Blanchard, M. Bohn, J. G. Bohnet, N. C. Brown, N. Q. Burdick, W. C. Burton, S. L. Campbell, J. P. Campora, C. Carron, J. Chambers, J. W. Chan, Y. H. Chen *et al.*, A race-track trapped-ion quantum processor, *Phys. Rev. X* **13**, 041052 (2023).
- [56] M. Ozols, How to generate a random unitary matrix, [http://home.lu.lv/~sd20008/papers/essays/Random%20unitary%20\[paper\].pdf](http://home.lu.lv/~sd20008/papers/essays/Random%20unitary%20[paper].pdf) (2009).
- [57] S. N. Filippov, S. Maniscalco, and G. García-Pérez, Scalability of quantum error mitigation techniques: from utility to advantage, [arXiv:2403.13542](https://arxiv.org/abs/2403.13542).
- [58] N. Keenan, N. F. Robertson, T. Murphy, S. Zhuk, and J. Goold, Evidence of Kardar-Parisi-Zhang scaling on a digital quantum simulator, *npj Quantum Inf.* **9**, 72 (2023).
- [59] B. G. M. Araújo, M. M. Taddei, D. Cavalcanti, and A. Acín, Local quantum overlapping tomography, *Phys. Rev. A* **106**, 062441 (2022).
- [60] JAX documentation, the Autodiff Cookbook, https://jax.readthedocs.io/en/latest/notebooks/autodiff_cookbook.html (2023).
- [61] S. Mangini, L. Maccone, and C. Macchiavello, Qubit noise deconvolution, *EPJ Quantum Technol.* **9**, 29 (2022).