

## Enhancing quantum state tomography via resource-efficient attention-based neural networks

Adriano Macarone Palmieri,<sup>1,\*</sup> Guillem Müller-Rigat<sup>1,\*</sup>,<sup>‡</sup> Anubhav Kumar Srivastava<sup>1</sup>,<sup>‡</sup> Maciej Lewenstein<sup>1,2</sup>,  
Grzegorz Rajchel-Mieldzióć<sup>1,3</sup> and Marcin Płodzień<sup>1,§</sup>

<sup>1</sup>ICFO - Institut de Ciències Fotòniques, The Barcelona Institute of Science and Technology, 08860 Castelldefels, Barcelona, Spain

<sup>2</sup>ICREA, Pg. Lluís Companys 23, 08010 Barcelona, Spain

<sup>3</sup>NASK National Research Institute, ul. Kolska 12, 01-045 Warszawa, Poland



(Received 3 October 2023; accepted 14 August 2024; published 4 September 2024)

In this paper, we propose a method for denoising experimental density matrices that combines standard quantum state tomography with an attention-based neural network architecture. The algorithm learns the noise from the data itself, without *a priori* knowledge of its sources. Firstly, we show how the proposed protocol can improve the averaged fidelity of reconstruction over linear inversion and maximum likelihood estimation in the finite-statistics regime, reducing at least by an order of magnitude the amount of necessary training data. Next, we demonstrate its use for out-of-distribution data in realistic scenarios. In particular, we consider squeezed states of few spins in the presence of depolarizing noise and measurement/calibration errors and certify its metrologically useful entanglement content. The protocol introduced here targets experiments involving few degrees of freedom and afflicted by a significant amount of unspecified noise. These include NISQ devices and platforms such as trapped ions or photonic qudits.

DOI: [10.1103/PhysRevResearch.6.033248](https://doi.org/10.1103/PhysRevResearch.6.033248)

### I. INTRODUCTION

Modern quantum technologies are fueled by resources such as coherence, quantum entanglement, and Bell nonlocality. Thus, a necessary step to assess the advantage they may provide is the certification of the above features [1–12]. The resource content of a preparation is revealed from the statistics (e.g., correlations) the device is able to generate. Within the quantum formalism, such data is encoded in the density matrix. In a real-world scenario, the density matrix can be reconstructed based on finite information from experimentally available probes, a process known as quantum state tomography (QST) [13–23]. Hence, QST is a prerequisite for any verification task. On the other hand, the second quantum revolution brought new experimental techniques to control and manipulate quantum systems [24–28], challenging established QST protocols. Both reasons elicited an extremely active field of research that over the years offered a plethora of algorithmic techniques [17, 19, 29–38].

Any conventional QST protocol suffers inevitably from a plethora of noise sources; ranging from measurement calibration errors, dark counts, losses, or technical noises, to name

a few. Such effects are eminently challenging to a model and will eventually decohere the system and wash out any quantum resource.

In recent years, machine learning (ML), artificial neural networks, and deep learning have entered the field of quantum technologies [39], offering many new solutions to QST [40–50]. These approaches learn the noise from the experimental data itself and are unaware of its sources or models [51–55]. Thus, not only shot noise, which is inherent to the QST task, but also other disturbances are susceptible to be mitigated by such methods. As only minimal assumptions about the system are required, they are especially suited for the certification task.

Despite their initial success, there are still setbacks in the application of neural network-based methods for QST. In this paper, we will focus on two of them: (i) the learning ability of a network with a reduced training set size [56], and (ii) the possibility of out-of-distribution (OOD) use of this class of methods [57]. OOD is a subfield of ML that analyzes how models perform on new data that do not belong to the training data distribution, with the latter called the *in-distribution* dataset (ID). To this end, we offer a computationally fast general protocol that combines established QST protocols and a supervised architecture trained to denoise the density matrix.

To begin with, we assess the learning and generalization abilities of the proposed method with generic states, as produced, e.g., from a random emitter. Then, we use our denoising network, trained exclusively on density matrices affected by shot noise only, to reconstruct new ones obtained from a simulated real-case scenario. In particular, we consider squeezed states of few spins under depolarizing and measurement calibration noise, and we certify its entanglement depth and usefulness for metrology applications.

\*These authors contributed equally to this work.

<sup>†</sup>Contact author: [adriano.macarone@icfo.eu](mailto:adriano.macarone@icfo.eu)

<sup>‡</sup>Contact author: [guillem.muller@icfo.eu](mailto:guillem.muller@icfo.eu)

<sup>§</sup>Contact author: [marcin.plodzien@icfo.eu](mailto:marcin.plodzien@icfo.eu)

Our protocol aims to certify quantum resources in low-dimensional systems afflicted by a significant level of unknown noise [51,58,59], where complete tomography is required.

The paper is organized as follows: In Sec. II, we introduce the main concepts behind the QST and in Sec. III we introduce the data generation protocol and neural network architecture, as well as define QST as a denoising task. Section IV is devoted to benchmarking our method against known approaches, and we test it on quantum states of physical interest. In Sec. V, we provide practical instructions to implement our algorithm in an experimental setting. We conclude in Sec. VI with several possible future research directions.

## II. PRELIMINARIES

Consider the  $d$ -dimensional Hilbert space. A set of informationally complete (IC) measurement operators  $\hat{\pi} = \{\hat{\pi}_i\}$ ,  $i = 1, \dots, d^2$ , in principle, allows unequivocally reconstructing the underlying target quantum state  $\hat{\tau} \in \mathbb{C}^{d \times d}$  within the limit of an infinite number of ideal measurements [60,61]. After infinitely many measurements, one can infer the mean values

$$p_i = \text{Tr}[\hat{\tau} \hat{\pi}_i], \quad (1)$$

and construct a valid vector of probabilities  $\mathbf{p} = \{p_i\}$  for any proper state  $\hat{\tau} \in \mathcal{S}$ , where by  $\mathcal{S}$  we denote the set of  $d$ -dimensional quantum states, i.e., containing all unit-trace, positive semi-definite (PSD)  $d \times d$  Hermitian matrices. Alternatively,  $\hat{\pi}$  can form a set of operators that spans the space of Hermitian matrices. In such a case,  $\mathbf{p}$  can be evaluated from multiple measurement settings (e.g., Pauli basis) and is generally no longer a probability distribution. In any case, there exists a one-to-one mapping  $Q$  from the mean values  $\mathbf{p}$  to the target density matrix  $\hat{\tau}$ ,

$$Q: \mathcal{F}_S \longrightarrow \mathcal{S} \\ \mathbf{p} \longmapsto Q[\mathbf{p}] = \hat{\tau}, \quad (2)$$

where  $\mathcal{F}_S$  is the space of accessible probability vectors. In particular, by inverting the Born's rule, Eq. (1), elementary linear algebra allows us to describe the map  $Q$  as

$$Q[\mathbf{p}] = \mathbf{p}^T \hat{G}^{-1} \hat{\pi}, \quad (3)$$

where  $\hat{G}$  is the Gram matrix of the measurements settings, with components  $G_{ij} = \text{Tr}(\hat{\pi}_i \hat{\pi}_j)$ .

The inference of the mean values  $\mathbf{p}$  is only perfect in the limit of an infinite number of measurement shots,  $N \rightarrow \infty$ .

In a realistic scenario, with a finite number of experimental runs  $N$ , we have access to frequencies of relative occurrence  $\mathbf{f} = \{f_i := n_i/N\}$ , where  $n_i$  is the number of times the outcome  $i$  is observed. Such counts allow us to estimate  $\mathbf{p}$  within an unavoidable error dictated by the shot noise, whose amplitude typically scales as  $1/\sqrt{N}$  [62]. With only frequencies  $\mathbf{f}$  available, we can use mapping  $Q$  for estimation  $\hat{\rho}$  of the target density matrix  $\hat{\tau}$ , i.e.,

$$\hat{\rho} = Q[\mathbf{f}]. \quad (4)$$

In the infinite number of trials  $N \rightarrow \infty$ ,  $f_i = p_i$ , and  $\hat{\rho} = \hat{\tau}$ . Yet, in the finite statistics regime, as considered in this paper,

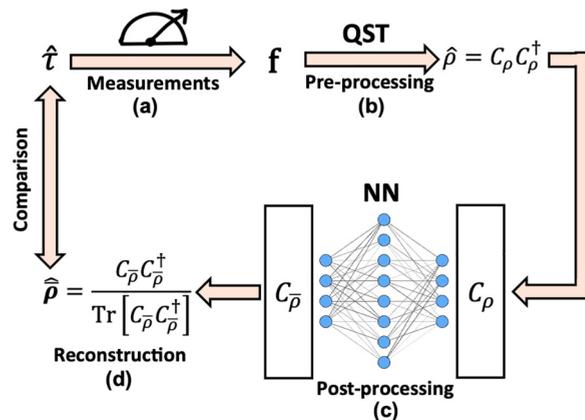


FIG. 1. Schematic representation of the data pipeline of our QST hybrid protocol. Panel (a) shows data acquisition from a generic experimental set-up, during which the frequencies  $\mathbf{f}$  are collected. Next, panel (b) presents standard density matrix reconstruction; in our paper, we test the computationally cheap LI method together with the expensive MLE, to better analyze the network reconstruction behavior and ability. Panel (c) depicts the matrix-to-matrix deep-learning strategy for Cholesky matrices reconstruction. The architecture herein considered combines convolutional layers for input and output and a transformer model in between. Finally, we compare the reconstructed state  $\hat{\rho}$  with the target  $\hat{\tau}$ .

the application of the mapping as defined in Eq. (3) to the frequency vector  $\mathbf{f}$  will generally lead to nonphysical results (i.e.,  $\hat{\rho}$  not PSD). In such case, as an example of proper mapping  $Q$  we can consider different methods for standard tomography tasks, such as linear inversion (LI), or maximum likelihood estimation (MLE); see Appendix A. As operators  $\hat{\pi}$ , we consider positive operator-valued measures (POVMs) and a more experimentally appealing Pauli basis (see Appendix B).

## III. METHODS

This section describes our density matrix reconstruction protocol, data generation, neural network training, and inference procedure. In Fig. 1, we show how these elements interact within the data flow. In the following paragraphs, we elaborate on the proposed protocol in detail.

The first step in our density matrix reconstruction protocol, called *preprocessing*, is a reconstruction of density matrix  $\hat{\rho}$  using finite-statistic QST with frequencies  $\mathbf{f}$  obtained from measurement prepared in target state  $\hat{\tau}$ . Next, we feed the reconstructed density matrix  $\hat{\rho}$  through our neural network acting as a noise filter, which we call this stage *postprocessing*. To enforce the positivity of the neural network output, we employ the Cholesky decomposition of the density matrices, that is,  $\hat{\rho} = C_\rho C_\rho^\dagger$  and  $\hat{\tau} = C_\tau C_\tau^\dagger$ , where  $C_{\rho,\tau}$  are lower-triangular matrices. Such decomposition is uniquely provided that  $\hat{\rho}$  and  $\hat{\tau}$  are positive [63]. We treat the Cholesky matrix  $C_\rho$  obtained from the finite-statistic QST protocol as a *noisy* version of the target Cholesky matrix without noise  $C_\tau$  calculated from  $\hat{\tau}$ . With these data, we prepare a supervised training for our architecture.

### A. Data generation

To construct the training data set, we first start with generating  $N_{\text{train}}$  Haar-random  $d$ -dimensional target density matrices  $\{\hat{\tau}_m\}$ , where  $m = 1, \dots, N_{\text{train}}$ . Next, we simulate experimental measurement outcomes  $\mathbf{f}_m$ , for each  $\hat{\tau}_m$ , in one of the two ways:

(1) *Directly*. When the measurement operators  $\hat{\pi}$  form an IC-POVM, we can take into account the noise by simply simulating the experiment and extracting the corresponding frequency vector  $\mathbf{f}_m = \{n_i/N\}_m$ , where  $N$  is the total number of shots (i.i.d. trials) and the counts  $\{n_i\}_m$  are sampled from the multinomial distribution.

(2) *Indirectly*. As introduced in the preliminaries (Sec. II), with projective measurements  $\hat{\pi}$  the  $\mathbf{p}_m$  is no longer a probability distribution, like the Pauli basis (see Appendix B). So, we can insert an amount of noise as the *direct* case, obtaining  $\mathbf{f}_m = \mathbf{p}_m + \delta\mathbf{p}_m$ , where  $\delta\mathbf{p}_m$  is sampled from the multinormal  $\mathcal{N}(\mathbf{0}, \sim \mathbf{1}/(2\sqrt{N}))$  of mean zero and isotropic variance, saturating the shot-noise limit.

Upon preparing the frequency vectors  $\{\mathbf{f}_m\}$ , we apply QST by mapping  $Q$ , Eq. (4), obtaining the set of reconstructed density matrices  $\{\hat{\rho}_m\}$ . We employ a rudimentary and scalable method, i.e., linear inversion [64], but other QST methods can also be used. Finally, we construct the training dataset as  $N_{\text{train}}$  pairs  $\{\vec{C}_\rho, \vec{C}_\tau\}$ , whereby  $\vec{C}$  we indicate the vectorization (flattening) of the Cholesky matrix  $C$  (see Appendix C for definition).

### B. Neural network architecture

Our proposed architecture is inspired by other recent models [55,65,66], combining convolutional layers with a transformer layer that implements a self-attention mechanism [67,68]. The convolutional layer extracts local features from the data, while the transformer seizes global ones. By combining them, we aim at taking the best of both approaches. The self-attention mechanism utilizes the representation of the input data as nodes within a graph [69] and aggregates the relationships between the nodes.

*Architecture*. The neural network action can be described as a mapping  $h_\theta$  that transforms the input—vectorized Cholesky matrix  $\hat{C}_\rho$  into an output  $h_\theta(\hat{C}_\rho)$ . The symbol  $\theta$  denotes all the variational parameters as weights and biases to be optimized during the training phase. The choice of architecture considered here contains two convolutional layers  $h_{\text{cnn}}$ , a transformer layer  $h_{\text{tr}}$  between, and a final linear layer  $h_1$ , i.e.,

$$h_\theta[\vec{C}_\rho] = \tanh(h_1) \circ h_{\text{tr}} \circ \gamma(h_{\text{cnn}})[\vec{C}_\rho], \quad (5)$$

where  $\gamma(y) = 1/2y(1 + \text{Erf}(y)/\sqrt{2})$ ,  $y \in \mathbb{R}$ , is the Gaussian error linear unit (GELU) activation function [70], broadly used in the modern transformers architectures, and  $\tanh(y)$  is the hyperbolic tangent, acting elementwise on neural network nodes. A detailed explication of the model is offered in Appendix D.

### C. Neural network training

The neural network training process relies on minimizing the cost function defined as a mean-squared error (MSE) of

the network output with respect to the target density matrix  $\hat{\tau}$ ,

$$\mathcal{L}(\theta) = \|\vec{C}_\tau - \vec{C}_\theta\|^2 + \text{Tr}[C_\theta C_\theta^\dagger], \quad (6)$$

with  $\text{Tr}[C_\theta C_\theta^\dagger]$  a regularization term, cf. Chap. 7 of Ref. [71] for detail. We train the model with a dataset containing  $N_{\text{train}}$  training samples  $\{\hat{\rho}_l\}$ . The equivalence between MSE and the Hilbert-Schmidt (HS) distance is discussed in detail in Appendix C, where we also demonstrate that the mean-squared error used in the cost function, Eq. (6), is a natural upper bound of the quantum fidelity. Hence, the choice of the cost function, Eq. (6) approximates the target state in a proper quantum metric.

By minimizing Eq. (6), we obtain the set of optimal parameters  $\hat{\theta}$  for our model  $h_{\hat{\theta}}$ . Finally, the neural network allows for the reconstruction of the target density matrix  $\hat{\tau}$  via Cholesky matrix  $C_{\hat{\rho}}$  [72], i.e.,

$$\hat{\rho} = \frac{C_{\hat{\rho}} C_{\hat{\rho}}^\dagger}{\text{Tr}[C_{\hat{\rho}} C_{\hat{\rho}}^\dagger]} \simeq \hat{\tau}, \quad (7)$$

where  $C_{\hat{\rho}}$  is reshaped from  $\vec{C}_{\hat{\rho}} = h_{\hat{\theta}}[\vec{C}_\rho]$ .

## IV. RESULTS AND DISCUSSION

Following the presentation of our QST protocol, we demonstrate its advantages in scenarios of both computational and physical interest. To this aim, we consider two examples.

As the first example, we study an idealized random quantum emitter (see e.g., Refs. [73,74] for recent experimental proposals) that samples high-dimensional mixed states from the Hilbert-Schmidt distribution. After probing the system using single-setting square-root POVM, we can show the usefulness of our neural network by improving the preprocessed LI and MLE states. This example enables us to assess the learning ability and expressivity of the QST neural network introduced in this paper.

In the second example, we focus on a specific class of multiqubit pure states of special physical relevance, i.e., with metrological resource, as quantified by the quantum Fisher information (QFI). Such states are generated via one-axis twisting dynamics (OAT) [75,76]. We simulate realistic data by considering the experiment in the presence of depolarizing noise, which mixes the OAT state  $|\Psi\rangle$  according to

$$\hat{\sigma} = (1 - p)|\Psi\rangle\langle\Psi| + \frac{p\mathbb{I}}{d}, \quad (8)$$

where  $\mathbb{I}$  is the identity operator. Furthermore, we reproduce miscalibration of the measurements by adding a bias error to the inferred expectation values. This example allows us to evaluate our protocol with OOD data.

### A. Reconstructing high-dimensional random quantum states

*Scenario*. Let us consider states  $\{\hat{\tau}_j\}$  sampled from distribution uniform with respect to the Hilbert-Schmidt measure (see Appendix E) on the Hilbert space of dimension  $d = 9$ . Each target state  $\hat{\tau}_j$  is measured  $N_{\text{trial}}$  times using its copies. This scenario allows us to benchmark our algorithm against the protocol offered in Ref. [53].

We prepare measurements on each trial state  $\hat{\tau}_j$  using information-complete square-root POVM (IC), as defined in

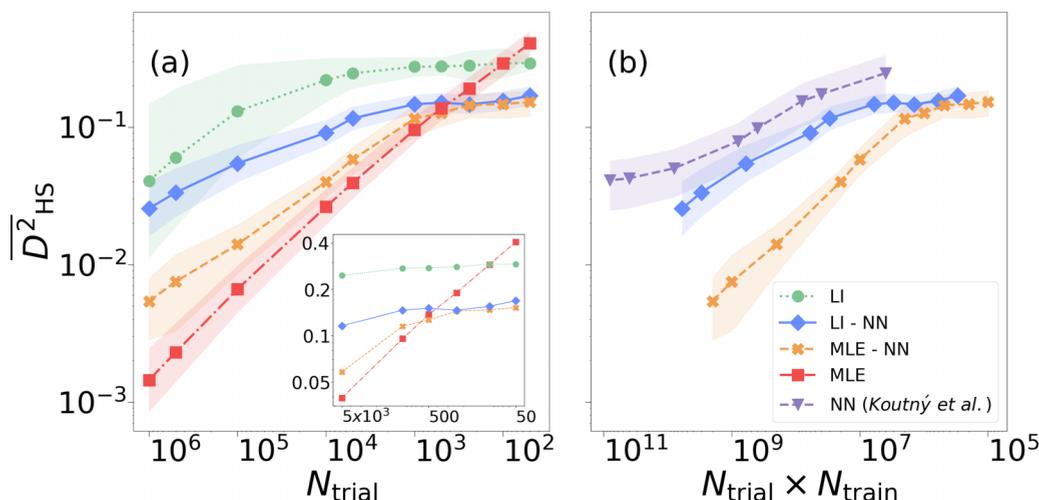


FIG. 2. Evaluation of the QST reconstruction quality measured by the mean value of the Hilbert-Schmidt distance square  $\overline{D^2_{\text{HS}}}$  between the target and the reconstructed state for different QST protocols, averaged over 1000 target states. In both panels, the best-performing setups are those that are as far right (better quality) and bottom (less costly) as possible. Panel (a) uses the number of measurements  $N_{\text{trial}}$  to compare four QST protocols: linear inversion (LI, green dots), neural network enhanced MLE (MLE-NN, orange crosses), neural network enhanced LI (NN-LI, blue diamonds), and maximal likelihood estimation (MLE, red squares). We add an inset focusing in the undersampled regime,  $N_{\text{trial}} \leq 5 \times 10^3$ . Panel (b) shows the quality of reconstruction as a function of product  $N_{\text{trial}} \times N_{\text{train}}$  for the latter two protocols and the network model proposed in Ref. [53] (violet triangles). Both panels depict resource costs on horizontal axes in different scenarios: in (a), the cost is the number of performed measurements, while in (b), the training phase is additionally counted as a cost. Our proposed protocol achieves competitive averaged HS reconstruction for the size of training data of an order of magnitude smaller than the method proposed in Ref. [53]. During models' training, we used  $N_{\text{train}} = 2000$  random pure states for the MLE-NN protocol, and  $N_{\text{train}} = 5000$  for the LI-NN. Lines are to guide the eye; shadow areas represent one standard deviation.

Eq. (B1). This allows obtaining the state reconstruction  $\hat{\rho}_j$  via two standard QST protocols, i.e., linear inversion (LI) and maximum likelihood estimation (MLE) algorithms, as well as by our neural network enhanced protocols denoted LI-NN and MLE-NN, see Fig. 1. Finally, we evaluate the quality of the reconstruction using the square of the Hilbert-Schmidt distance between the target and the reconstructed state  $D_{\text{HS}}^2(\hat{\rho}_j, \hat{\tau}_j) = \text{Tr}[(\hat{\rho}_j - \hat{\tau}_j)^2]$ .

**Benchmarking.** Figure 2(a) presents the averaged Hilbert-Schmidt distance square  $D^2_{\text{HS}}$  as a function of the number of trials states  $N_{\text{trial}}$ . To obtain a reconstruction for a given averaged HS distance, our neural-network-enhanced protocols require a much lower number of  $N_{\text{trial}}$  copies of states, compared to linear inversion and maximum likelihood estimation alone. We note that the proposed protocols improve MLE for the relatively small number of trial states  $N_{\text{trial}} < 10^3$  (see inset), which is important from an experimental point of view. As expected, the lowest HS distance is obtained for many trials of the MLE algorithm.

Recently, a state-of-the-art QST neural network protocol was proposed by Koutný *et al.* [53]. The authors report better performance than the MLE and LI algorithms with  $N_{\text{train}} = 8 \times 10^5$  training samples for qutrit as well as larger systems,  $d \geq 3$ . In Fig. 2(b), we show that our protocol requires an order of magnitude smaller training data, achieving a comparable level of reconstruction.

### B. Certifying metrologically useful entanglement in noisy spin-squeezed states

In this simulation, we reconstruct a class of physically relevant multiqubit pure states in a realistic scenario.

Specifically, we consider a chain of  $L = 4$  spins-1/2 (Hilbert space of dimension  $d = 16$ ). The target quantum states are dynamically generated during the one-axis twisting (OAT) protocol [75,76]

$$|\Psi(t)\rangle = e^{-it\hat{J}_z^2} |+\rangle^{\otimes L}, \quad (9)$$

where  $\hat{J}_z$  is the collective spin operator along  $z$  axis and  $|+\rangle^{\otimes L} = [(|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}]^{\otimes L}$  is the initial state prepared in a coherent spin state along  $x$  axis (orthogonal to  $z$ ). The OAT protocol generates spin-squeezed states useful for high-precision metrology, allowing to overcome the shot-noise limit [76–79], as well as many-body entangled and the many-body Bell correlated states [80–92]. OAT states have been extensively studied theoretically [75,93–101], and can be realized with a variety of ultra-cold systems, utilizing atom-atom collisions [102–105], and atom-light interactions [106,107]. The recent theoretical proposals for the OAT simulation with ultra-cold atoms in optical lattices effectively simulate Hubbard and Heisenberg models [98,108–115].

**Data preparation.** For this task, we generate our data focusing on the experimentally friendly Pauli operators. The state under consideration are prepared *indirectly*, with a fidelity from the target OAT state of  $\sim 85\%$ . Then we simulate the presence of a noise channel by depolarizing our input state with a strength factor  $p = 0.3$  according to Eq. (8). Lastly, to simulate the presence of a calibration defect in the measurement apparatus, a fixed bias of random values of the order of  $10^{-4}$  is added to the Born values. After applying all these noise sources, the LI reconstructions can obtain an average fidelity of  $75.4 \pm 1.1\%$ . For the test set, we select 100 OAT

TABLE I. Comparison of average fidelity and its standard deviation between the reconstructed and the target states of size  $d = 16$  for various QST methods (rows), with varying size of measurement trials  $N_{\text{trial}} = 10^6, 10^5, 10^4, 10^3$ , as indicated by the consecutive columns. The first row presents the average fidelity reconstruction for linear inversion QST, averaged over OAT states, evenly sampled from  $t = 0$  to  $t = \pi$ . Employing our neural network presents an enhancement over the bare LI, as shown in the second row for the same target set. Finally, the third row also shows data for NN-enhanced LI but averaged over general Haar-random states. All initial Born values are calculated by noiseless SIC-POVM.

	$10^6$ trials	$10^5$ trials	$10^4$ trials	$10^3$ trials
LI-OAT	$98.1 \pm 0.3$	$94.3 \pm 0.4$	$86.5 \pm 0.1$	$67.0 \pm 2.0$
NN-OAT	$99.3 \pm 0.2$	$98.6 \pm 0.4$	$97.8 \pm 0.5$	$87.6 \pm 4.1$
NN-Haar	$99.0 \pm 0.2$	$96.9 \pm 3.0$	$94.2 \pm 3.3$	$81.1 \pm 4.1$

states in evenly spaced times  $t \in (0, \pi)$  and assess the average reconstruction achieved by our neural network.

The description of the model under consideration can be found in Table I.

*Inferring the quantum Fisher information.* Finally, we evaluate the metrological usefulness of the reconstructed states measured by the quantum Fisher information (QFI)  $F_Q[\hat{\rho}, \hat{G}]$ . The QFI is a nonlinear function of the state and quantifies the sensitivity upon rotations generated by  $\hat{G}$ . Notwithstanding its usefulness in metrology, it is a highly nontrivial task to evaluate experimentally because of its great sensitivity to noise disturbances, with state-of-the-art research only up to four qubits [116,117]. We refer the reader to Appendix F for more details.

In this context, we consider the collective spin component  $\hat{J}_v$  as the generator  $\hat{G} = \hat{J}_v$ , with the orientation  $v \in \mathbb{R}^3$  selected to achieve maximal sensitivity. Quantum Fisher information (QFI) related to collective rotations can also serve to verify quantum entanglement [79], specifically the entanglement depth  $k$ , which is the smallest number of genuinely entangled particles required to describe the state. If  $F_Q[\hat{\rho}, \hat{J}_v] > kL$ , then the quantum state  $\hat{\rho}$  possesses an entanglement depth of at least  $k + 1$  [118,119]. In particular, for states with depth  $k = 1$  (i.e., separable), the absence of detected entanglement implies that the metrological capability is limited to the shot-noise threshold [120]. This limit is reached by coherent spin states, such as our initial ( $t = 0$ ) state for the evolution of the one-axis twisting (OAT),  $|+\rangle^{\otimes L}$ .

*OOD results.* In Fig. 3, we present the evolution of the QFI (normalized by the coherent limit,  $L = 4$ ) for the OAT target states (top of the solid blue lines). For this numerical experiment, we make full use of the OOD approach; using a network trained only for tackling sampling noise, we feed it in inference with a dataset that also considers the depolarization and measurement noise. As Fig. 3 shows, for two different realizations of calibration noise, our network can highly improve the LI reconstructions obtaining a fidelity of  $88.7 \pm 2.3\%$  and  $91.0 \pm 1.9\%$ , on the left and right panel, respectively. Thereby, we surpass the three-body bound (QFI/ $L = 3$ ), thus revealing a genuine four-body entanglement, which is the highest depth possible in this system (since it is of size

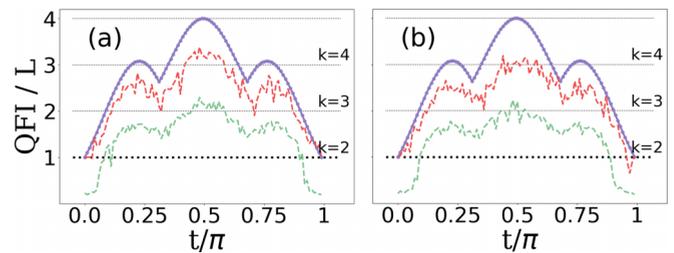


FIG. 3. Two different simulations for out-of-distribution (OOD) inference. In each panel, we evaluate the normalized quantum Fisher information (QFI) for 100 four-qubit states as validation metric. The target, noiseless states are evolved according to the OAT dynamics given in (9), and depicted by the purple dotted line. For this OOD tests, the neural network was trained exclusively to learn statistical sampling noise. During inference, test data are permeated by depolarization and measurement (calibration) errors also. The green line represents the normalized QFI derived from reconstructions via the Linear Inversion (LI) algorithm; the red line illustrates the enhancement provided by the network when supplemented with LI algorithm reconstructions, underscoring the robustness of our protocol in mitigating noise effects.

$L = 4$ ). For example, note that at time  $t = \pi/2$ , the OAT dynamics generates the cat state,  $|\Psi(t = \pi/2)\rangle = (e^{-i\pi/4}|+\rangle^{\otimes 4} + e^{+i\pi/4}|-\rangle^{\otimes 4})/\sqrt{2}$ , which is genuinely  $L$ -body entangled, and so it is certified. For completeness, two complexity analyses are offered. First, in Appendix G, an analysis of the QFI time evolution for OAT states for different sampling noise values only, using Pauli and tomographically optimal SIC-POVM operators in data generation. Next, in Appendix H, a benchmark of our model is shown against two different convolutional architectures to assess the advantage offered by our transformer-based model on quantum data. Lastly, we propose an alternative method to incorporate the statistical noise as a depolarizing channel in the Appendix I.

## V. CONCRETE EXPERIMENTAL IMPLEMENTATION

To recapitulate this contribution, as a complement to Fig. 1 and our repository provided in Ref. [121], we summarize the practical implementation of the protocol introduced in this paper.

(1) *Scenario.* We consider a finite-dimensional quantum system prepared in a target state  $\hat{\tau}$ . Here, our objective is to verify the preparation of a quantum state  $\hat{\tau}$  via QST. To this end, we set a particular measurement basis  $\hat{\pi}$  to probe the system.

(2) *Experiment.* After a finite number of experimental runs, we construct the frequency vector  $\mathbf{f}$  from the countings.

(3) *Preprocessed quantum state tomography.* From the frequency vector  $\mathbf{f}$  and the basis  $\hat{\pi}$ , we infer the first approximation of the state  $\hat{\rho}$  via the desired QST protocol (e.g., one of those introduced in Appendix A).

(4) *Assessing pre-reconstruction.* We evaluate the quality of the reconstruction by, for example, computing  $D_{\text{HS}}^2(\hat{\tau}, \hat{\rho})$ , quantum fidelity, or any other meaningful quantum metric. To improve such a score, we resort to our neural network solution to complete a denoising task. As with any deep-learning method, training is required.

(5) *Training strategies.* Different training strategies can be implemented:

(1) Train over uniform ensembles (e.g., Haar, HS, Bures etc.) if  $\hat{\tau}$  is a typical state or we do not have information about it.

(2) Train over a subspace of states of interest. For example, if we reconstruct OAT states (Sec. IV B), we may train only in the permutation-invariant sector.

(3) Due to the greater generalization ability demonstrated when the case of mixed state was considered, we can perform transfer learning to tailor a pretrained model on a specific apparatus with less computational resources; in this way, the model will have knowledge of that specific apparatus noise. For example, if we have a quantum random source to characterize (Sec. IV A), an amount of experimental data corresponding to the 10–15% of the training dataset can be used to refine the model.

(6) *Feeding the neural network.* We feed the preprocessed state  $\hat{\rho}$  into our trained matrix-to-matrix neural network to recover the enhanced quantum state  $\hat{\rho}$ .

(7) *Assessing the neural network.* We compute the updated reconstruction metric on the post-processed state  $D_{\text{HS}}^2(\hat{\tau}, \hat{\rho})$ . Finally, we assess the usefulness of the neural network by comparing how small such a value is compared to the pre-processed score  $D_{\text{HS}}^2(\hat{\tau}, \hat{\rho})$ .

The strength of our proposed protocol lies in its broad applicability, as the choice of the basis  $\hat{\tau}$  and the QST pre-processing method is arbitrary.

## VI. CONCLUSIONS

We proposed a deep learning protocol that improves standard quantum state tomography methods, such as linear inversion and maximum likelihood estimation. Based on a combination of transformer and convolutional layers, we greatly reduce the dataset dimension for training, and we can perform denoising on new unseen data, when also depolarization and measurement noise are accounted for. First, the proposed method reduces the number of necessary measurements in the target density matrix by at least an order of magnitude compared to other QST protocols supported by finite-statistic neural networks.

Secondly, for four-qubits and Bell-correlated few-spin states generated with OAT, the inference stage was on OOD. We tested our model, pretrained only for statistical sampling noise, on data accounting also for depolarization and measurement noise, achieving an average fidelity reconstruction  $88.7 \pm 2.3\%$  and  $91.0 \pm 1.9\%$  for two different realizations of noise in the measurement setup. The superior learning ability demonstrated for mixed states makes our architecture an optimal candidate for transfer learning, when further refinement is desired on a pretrained model using fewer experimental data. On the other hand, the OOD demonstrate the potential of our protocol for more plug-and-measure applications, given its high resilience to unknown levels of noise because of different physical sources. Thus, it paves the way for the use of these methods in current quantum computers, NISQ devices, and quantum simulators based on spin arrays [19,122].

Another persistent challenge in this field concerns the scalability of the algorithms with the number of subsystems.

In this regard, several strategies are proposed based on *incomplete* tomography like the well known generative NN applications [41]. A limitation of the protocol is its scalability, which is limited by the use of Cholesky decomposition in data preprocessing, which limits our method to applications of modest Hilbert space dimensions. An extension to a Cholesky-free approach is left for future explorations.

Data and code are available on github [121].

## ACKNOWLEDGMENTS

The authors thank L. Zambrano, F. Bianchi, and E. Witkowska for the fruitful discussions. ICFO group acknowledges support from: Europea Research Council AdG NOQIA; MCIN/AEI (PGC2018-0910.13039/501100011033, CEX2019-000910-S/10.13039/501100011033, Plan National FIDEUA PID2019-106901GB-I00, Plan National STAMEENA PID2022-139099NB, I00, project funded by MCIN/AEI/10.13039/501100011033 and by the “European Union NextGenerationEU/PRTR” (PRTR-C17.11), FPI, QUANTERA MAQS PCI2019-111828-2); QUANTERA DYNAMITE PCI2022-132919, QuantERA II Programme co-funded by European Union’s Horizon 2020 program under Grant Agreement No. 101017733; Ministry for Digital Transformation and of Civil Service of the Spanish Government through the QUANTUM ENIA project call – Quantum Spain project, and by the European Union through the Recovery, Transformation and Resilience Plan – NextGenerationEU within the framework of the Digital Spain 2026 Agenda; Fundació Cellex; Fundació Mir-Puig; Generalitat de Catalunya (European Social FundFEDER and CERCA program, AGAUR Grant No. 2021 SGR 01452, QuantumCAT U16-011424, co-funded by ERDF Operational Program of Catalonia 2014-2020); Barcelona Supercomputing Center MareNostrum (FI-2023-3-0024); Funded by the European Union. ICFO Internal “QuantumGaudi” project; European Union’s Horizon 2020 program under the Marie Skłodowska-Curie grant agreement No 847648; “La Caixa” Junior Leaders fellowships, La Caixa-Foundation (ID 100010434) CF/BQ/PR23/11980043.

Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union, European Commission, European Climate, Infrastructure and Environment Executive Agency (CINEA), or any other granting authority. Neither the European Union nor any granting authority can be held responsible for them (HORIZON-CL4-2022-QUANTUM-02-SGA PASQuanS2.1, 101113690, EU Horizon 2020 FET-OPEN OPTologic, Grant No. 899794), EU Horizon Europe Program (This project has received funding from the European Union’s Horizon Europe research and innovation program under grant agreement No 101080086 NeQSTGrant Agreement 101080086 – NeQST).

## APPENDIX A: ESTABLISHED APPROACHES FOR STATE TOMOGRAPHY

We review four of the most well-known approaches to quantum state reconstruction that could be potentially improved with the protocol proposed in this paper.

**1. Linear inversion (LI)**

By inverting Born’s rule Eq. (1) (main text) we can express the state dependence on the mean values  $\mathbf{p} = \{p_i\}$ ,

$$\hat{\tau} = \mathbf{p}^T \text{Tr}[\hat{\pi} \hat{\pi}^T]^{-1} \hat{\pi}. \tag{A1}$$

Note that the inverse of the Gram matrix  $\text{Tr}(\hat{\pi} \hat{\pi}^T)$  exists as the basis is informationally complete (IC). If it is (informationally) overcomplete, one needs to replace the inverse with the pseudoinverse. Finally, if it is undercomplete (only partial information is available), it will determine the state up to a linear subspace.

The LI method infers  $\hat{\rho}$  by replacing in Eq. (A1) the ideal expectation values  $\mathbf{p}$  with the vector  $\mathbf{f}$  of the experimental frequencies (counts). This naive substitution generally leads to a negative matrix,  $\hat{\rho} \not\geq 0$ . An *optimal* way to tame its negative eigenvalues was given in [123] by finding the physical state closest to  $\hat{\rho}$  in the two-norm. The drawback of LI is the fact that it can be affected by any type of noise.

**2. Least-squares estimation (LSE)**

Here, the reconstructed state is chosen to minimize the mean square error between the experimental frequencies  $\mathbf{f}$  and the state probability distribution  $\text{Tr}(\hat{\pi} \hat{\rho})$ . The resulting problem can be expressed as

$$\hat{\rho}_{\text{LSE}} = \arg \min_{\hat{\rho} \geq 0} |\mathbf{f} - \text{Tr}(\hat{\pi} \hat{\rho})|^2. \tag{A2}$$

The problem Eq. (A2) is convex and, in addition, it is a disciplined convex problem (DCP), i.e., a class of convex optimizations that are efficiently addressed with commercially available solvers such as MOSEK [124].

**3. Maximum likelihood estimation (MLE)**

In this case, the reconstructed state maximizes the likelihood of having produced the observed experimental outcomes,

$$\hat{\rho}_{\text{MLE}} = \arg \max_{\hat{\rho} \geq 0} P(\hat{\rho} | \mathbf{f}). \tag{A3}$$

Our counting experiment is modeled as a multinomial. Consequently, the log likelihood is  $\log P(\hat{\rho} | \mathbf{f}) = \mathbf{f} \cdot \log(\text{Tr}[\hat{\pi} \hat{\rho}])$ , which is a concave function of the state, but the resulting task is not a DCP. Therefore, solving it can be expensive, especially for a large Hilbert dimension space  $d$ . The MLE is a robust estimator against noise; however, it is computationally demanding, suffering from the exponential scaling of the inputs.

**APPENDIX B: INFORMATIONALLY COMPLETE MEASUREMENT OPERATORS**

To obtain our set of initial Born values, we use three informationally complete (IC) sets of Hermitian operators  $\pi$  defined on the Hilbert space of  $d$ .

**1. Square-root POVM**

The first set  $\pi$  of measurement operators consists of POVM generated by the square-root measurements, defined as

$$\pi_i = \hat{H}^{-1/2} |\phi_i\rangle \langle \phi_i| \hat{H}^{-1/2} \quad \hat{H} = \sum_{i \in [d^2]} |\phi_i\rangle \langle \phi_i|, \tag{B1}$$

with  $\{|\phi_i\rangle \in \mathcal{H}\}_{i \in [d^2]}$  are randomly generated Haar distributed pure states (see Appendix E).

**2. SIC-POVM**

The second set  $\pi$  of measurement basis operates in an  $L$ -qubit system. The basis consists of the tensor product of local sic-POVM, constructed by using the local vectors,

$$\begin{aligned} \mathbf{s}_1 &= (0, 0, 1), \\ \mathbf{s}_2 &= \left(\frac{2\sqrt{2}}{3}, 0, -\frac{1}{3}\right), \\ \mathbf{s}_3 &= \left(-\frac{\sqrt{2}}{3}, \sqrt{\frac{2}{3}}, -\frac{1}{3}\right), \\ \mathbf{s}_4 &= \left(-\frac{\sqrt{2}}{3}, -\sqrt{\frac{2}{3}}, -\frac{1}{3}\right). \end{aligned} \tag{B2}$$

The space of Hermitian operators acting in the global Hilbert space  $\mathcal{H} = [\mathbb{C}^2]^{\otimes L}$  can then be spanned by

$$\hat{\pi} = \left\{ \hat{\pi}_{(a,b) \cong i} = \bigotimes_{b \in [L]} \frac{\hat{\sigma}_0 + \mathbf{s}_{ab} \cdot \hat{\boldsymbol{\sigma}}}{4} \right\}_{\{a_b\}_{b \in [L]} \in [4]^L}, \tag{B3}$$

where  $\hat{\boldsymbol{\sigma}} = (\hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z)$  is the Pauli vector and  $\hat{\sigma}_0 = \mathbb{I}_2$  is the identity acting in the local space  $\mathbb{C}^2$ .

Note that for any properly normalized state  $\hat{\tau}$ ,  $\mathbf{p} := \text{Tr}(\hat{\pi} \hat{\tau})$  constitutes a valid probability distribution

$$\forall i \in [2^{2L}], \quad p_i \geq 0 \quad \text{and} \quad \sum_{i \in [2^{2L}]} p_i = 1. \tag{B4}$$

This observation is also true for the previous basis (square-root POVM) that we reviewed.

**3. Pauli basis**

The last IC basis  $\hat{\pi}$  in  $L$ -qubit systems is the Pauli basis constructed as

$$\hat{\pi} = \left\{ \bigotimes_{b \in [L]} \hat{\sigma}_{a_b} \right\}_{\{a_b\}_{b \in [L]} \in \{0,x,y,z\}^L}. \tag{B5}$$

With respect to such a basis, expectation values can be evaluated experimentally by rotation of each qubit individually. This is also true for the SIC-POVM if evaluated with multiple settings. Such expectation values  $\mathbf{p}$  no longer form a probability distribution (note that, in particular, such mean values can be negative). The reason why it will not lead to a probability distribution is that the basis does not form a POVM (that is, its elements are not PSD and do not sum up to  $\mathbb{I}$ ). However, it covers the whole space of Hermitian matrices supported in  $[\mathbb{C}^2]^{\otimes L} = \mathcal{H}$  as any basis specified in this Appendix.

## APPENDIX C: FROM QUANTUM FIDELITY TO MEAN-SQUARED ERROR

### 1. Upper bound on the Bures distance

The Bures distance between two states  $\hat{\rho}$  and  $\hat{\tau}$  is defined as

$$D_B(\hat{\rho}, \hat{\tau}) = 2 - 2\sqrt{F(\hat{\rho}, \hat{\tau})}, \quad (\text{C1})$$

where

$$\sqrt{F(\hat{\rho}, \hat{\tau})} = \text{Tr}[\sqrt{\sqrt{\hat{\rho}}\hat{\tau}\sqrt{\hat{\rho}}}], \quad (\text{C2})$$

is the square root of quantum fidelity between  $\hat{\rho}$ , and  $\hat{\tau}$ . The square root of the fidelity  $\sqrt{F(\hat{\rho}, \hat{\tau})}$  can be expressed as [[125], Eq. (9.30)]

$$\sqrt{F(\hat{\rho}, \hat{\tau})} = \max_{\substack{\hat{\rho}=AA^\dagger \\ \hat{\tau}=BB^\dagger}} \text{Tr}[AB^\dagger + BA^\dagger]/2, \quad (\text{C3})$$

where the maximization is over the *complex amplitudes*  $\{A, B\}$ , which constitute a polar decomposition of  $\{\hat{\rho}, \hat{\tau}\}$  respectively. Equation (C3) is actually the original definition of quantum fidelity motivated by the concept of transition probability. In fact, if both states are pure  $\{\hat{\rho} = |\psi\rangle\langle\psi|, \hat{\tau} = |\phi\rangle\langle\phi|\}$ , then  $\{A = |\psi\rangle, B = |\phi\rangle\}$  and the right-hand side (RHS) of Eq. (C3) amounts to overlap  $|\langle\psi|\phi\rangle|^2$ . Note that the decomposition admits a gauge degree of freedom  $A \mapsto AU$  for  $U$  unitary (and similarly for  $B$ ). Our study resolves redundancy using the *Cholesky decomposition* defined in the main text.

From Eq. (C3) we see that for any polar decomposition (and in particular the Cholesky as canonical one,  $\{\hat{\rho} = C_\rho C_\rho^\dagger, \hat{\tau} = C_\tau C_\tau^\dagger\}$ ), the following inequality always holds:

$$\sqrt{F(\hat{\rho}, \hat{\tau})} \geq \text{Tr}[C_\rho C_\tau^\dagger + C_\tau C_\rho^\dagger]/2. \quad (\text{C4})$$

Finally, rewriting 2 as  $1 + 1 = \text{Tr}(\hat{\rho} + \hat{\tau}) = \text{Tr}(C_\rho C_\rho^\dagger + C_\tau C_\tau^\dagger)$  we arrive at

$$\begin{aligned} D_B(\hat{\rho}, \hat{\tau}) &\leq \text{Tr}[C_\rho C_\rho^\dagger + C_\tau C_\tau^\dagger - C_\rho C_\tau^\dagger - C_\tau C_\rho^\dagger] \\ &= \text{Tr}[(C_\rho - C_\tau)(C_\rho - C_\tau)^\dagger] \\ &= D_{\text{HS}}^2(C_\rho, C_\tau), \end{aligned} \quad (\text{C5})$$

where the HS distance defined in the main text is extended to complex matrices (not necessarily Hermitian) as

$$D_{\text{HS}}^2(C_\rho, C_\tau) = \text{Tr}[(C_\rho - C_\tau)(C_\rho - C_\tau)^\dagger]. \quad (\text{C6})$$

### 2. Hilbert-Schmidt distance as mean-squared error (MSE)

In the following we connect the Hilbert-Schmidt distance [Eq. (C6)] between two Cholesky matrices  $\{C_\rho, C_\tau\}$  associated to quantum states  $\{\hat{\rho}, \hat{\tau}\}$ , with the mean-squared error of the matrix elements.

First, consider a  $d \times d$  complex matrix  $K$ ,  $\{K_{\alpha\beta}\}_{\alpha, \beta \in [d] \times [d]}$ . Next, let us introduce the vectorization  $\vec{K}$  of its matrix elements as

$$\vec{K} = \Re\mathbb{K} \oplus \Im\mathbb{K}, \quad (\text{C7})$$

where  $\mathbb{K}$  is the flattening of the matrix, i.e.,  $\mathbb{K} = (K_{11}, K_{12}, \dots, K_{dd})$ , and  $\oplus$  the direct sum of vectors,  $\vec{v} \oplus \vec{u} = (\vec{v}, \vec{u})$ . Since  $\text{diag}(K) \in \mathbb{R}$  is sufficient for neural network

output, the length of the real vector  $\vec{K}$  is  $d^2$  [ $d$  for the diagonal and  $d(d-1)/2 \times 2$  for the lower triangle ( $\times 2$  for the real and imaginary part); the remaining elements are zero].

Let  $K = C_\rho - C_\tau$ , then the square HS distance, Eq. (C6), reads

$$D_{\text{HS}}^2(C_\rho, C_\tau) = \text{Tr}(KK^\dagger) = \sum_{\alpha, \beta} |K_{\alpha\beta}|^2 \quad (\text{C8})$$

$$= \sum_{\alpha, \beta} \Re K_{\alpha\beta}^2 + \Im K_{\alpha\beta}^2 \quad (\text{C9})$$

$$= (\Re\mathbb{K} \oplus \Im\mathbb{K}) \cdot (\Re\mathbb{K} \oplus \Im\mathbb{K}) \quad (\text{C10})$$

$$= \|\vec{C}_\rho - \vec{C}_\tau\|^2 = \text{MSE}(\vec{C}_\rho, \vec{C}_\tau). \quad (\text{C11})$$

Finally, we observe that the MSE is the natural cost function of a standard feedforward neural network.

## APPENDIX D: ARCHITECTURE DETAILS

The first layer  $h_{\text{cnn}}$  applies a set of  $K$  fixed-size trainable one-dimensional convolutional kernels to  $\vec{C}_\rho$  followed by a nonlinear activation function, i.e.,  $\gamma(h_{\text{cnn}}(\vec{C}_\rho)) \rightarrow \{\mathbf{F}_{\text{cnn}}^1, \dots, \mathbf{F}_{\text{cnn}}^K\}$ . During the training process, the convolutional kernels learn different features of the dataset, which are then fed to the transformer block  $h_{\text{tr}}$ . The transformer block  $h_{\text{tr}}$  distills the correlations between the features extracted from the kernels through the self-attention mechanism, providing a new set of vectors, that is,  $h_{\text{tr}}(\mathbf{F}_{\text{cnn}}^1, \dots, \mathbf{F}_{\text{cnn}}^K) \rightarrow \{\mathbf{F}_{\text{tr}}^1, \dots, \mathbf{F}_{\text{tr}}^K\}$ . In the last step, the outputs of the convolutional kernel of the layer  $h_{\text{cnn}}$  are added and form an output  $\vec{C}_\theta$ ,  $\tanh(h_{\text{cnn}}(\mathbf{F}_{\text{tr}}^1, \dots, \mathbf{F}_{\text{tr}}^K)) \rightarrow \vec{C}_\theta$ . The output is combined in  $C_\theta$  as in Eq. (7), and a custom cost function  $\eta$  is used to apply the REIU activation on the diagonal elements of the reconstructed matrix and the tanh on the off-diagonal elements. The role of the attention mechanism is explored in Appendix H, where we benchmark the transformer block against CNNs.

The training data and the considered architecture allow interpreting the trained neural network as a conditional debiased (for details, see the Appendix J). Although the proposed protocol cannot improve the predictions of unbiased estimators, any estimator that outputs valid quantum states (e.g., LI, MLE) must be biased because of boundary effects. In the given framework, the task of the neural network is to learn such skewness and drift the distribution towards the true mean.

### Computational details

The architecture is trained on 10 000 data in training and 1500 in validation, with a batch size of 1500. The code is run on a Nvidia A100 GPU card with 80 GB of memory (Cuda version 12.1). The total training time amounts to  $\sim 25$  minutes. A similar time can be obtained on a commercially available GPU virtual machine.

## APPENDIX E: SAMPLING RANDOM STATES

We want the algorithm to learn how noise affects generic states that cover the whole space  $\mathcal{S}$  of proper quantum states. In doing so, we will have a flexible solution that can potentially improve any state. To this end, we will train our neural network with random states sampled from different measures

that uniformly cover the volume of interest. In particular, for mixed states, we will use the Hilbert-Schmidt ensemble, while for pure states, the Haar measure; in the following, we review how to generate such random states.

### 1. Hilbert-Schmidt states

The HS measure is defined by the infinitesimal line element

$$ds_{\text{HS}}^2 = \text{Tr}[(d\hat{\tau}_{\text{HS}})^2], \quad (\text{E1})$$

which indeed induces the HS distance  $D_{\text{HS}}$ , as defined in the main text, once integrated.

States uniformly distributed in the space endowed with the corresponding metric, Eq. (E1), may be sampled as

$$\hat{\tau}_{\text{HS}} = \frac{AA^\dagger}{\text{Tr}[AA^\dagger]}, \quad (\text{E2})$$

where  $A$  is a generic complex square matrix whose real and imaginary matrix elements are i.i.d. sampled normal random variables  $\mathcal{N}(0, 1)$ .

### 2. Haar vectors

Such ensemble is defined on the pure-state space  $\hat{\tau} = |\Psi\rangle\langle\Psi|$  and may be used to infer states with high purity, as the metrologically useful states explored in the main text. Formally, it takes advantage of the link between states  $|\Psi\rangle \in \mathcal{H}$ , and the unitary group. The fundamental property of such measure is that it is invariant under unitary transformations, thus making any vector in  $\mathcal{H}$  equiprobable.

Interestingly, Haar vectors (or equivalently unitaries), can be drawn again from i.i.d. random normal variables  $\mathcal{N}(0, 1)$  [126]. Let  $A$  be a complex square matrix sampled as before. Then, we apply the so-called QR decomposition,

$$A = QR, \quad (\text{E3})$$

where  $Q$  is unitary and  $R$  is upper triangular. Then  $Q' = Q\Lambda$ , where  $\Lambda = \text{diag}(\{|R_{\alpha\alpha}|/|R_{\alpha\alpha}|\}_\alpha)$ , is Haar random.

In practice, we use the python package *qutip* to compute those. For further information on random states, we refer to Ref. [125].

## APPENDIX F: CAPTURING THE METROLOGICAL USEFULNESS OF OAT EVOLVED STATES

Given a quantum state  $\hat{\rho}$  with spectral decomposition  $\hat{\rho} = \sum_k p_k |k\rangle\langle k|$ , we can evaluate its sensitivity upon rotations generated by  $\hat{G}$ ,  $\hat{\rho}(\theta) = e^{-i\theta\hat{G}}\hat{\rho}e^{+i\theta\hat{G}}$  as quantified by the QFI,

$$F_Q[\hat{\rho}, \hat{G}] = 2 \sum_{\substack{k,l \\ p_k + p_l > 0}} \frac{(p_k - p_l)^2}{p_k + p_l} | \langle k | \hat{G} | l \rangle |^2. \quad (\text{F1})$$

In the main text, we consider an ensemble of  $L$  spins-1/2, or two-level atoms. In such system and in the context of magnetometry (e.g., via Ramsey interferometry), the phase is encoded the same way in every constituent via the collective generator  $\hat{J}_{\mathbf{v}} = \sum_{i \in [L]} \mathbf{v} \cdot \hat{\sigma}_i$  where  $\mathbf{v} = (v_x, v_y, v_z)$  is the encoding orientation.

For a generic state  $\hat{\rho}$ , it is not direct to find the optimal spatial direction  $\mathbf{v}$  to exploit the maximal sensitivity. However,

for pure states  $\hat{\rho} = |\Psi\rangle\langle\Psi|$  (like the OAT evolution that we consider), the QFI is (four times) the variance of the generator,  $F_Q[|\Psi\rangle, \hat{G}] = 4[\langle\Psi|\hat{G}^2|\Psi\rangle - \langle\Psi|\hat{G}|\Psi\rangle^2]$  and the best direction is then yielded by the maximal eigenvalue  $\mathbf{v}_{\text{max}}$  of the covariance matrix  $\mathcal{C}$ ,

$$\mathcal{C} = \begin{pmatrix} \langle \hat{J}_x^2 \rangle & \text{Re}\{\langle \hat{J}_x \hat{J}_y \rangle\} & \text{Re}\{\langle \hat{J}_x \hat{J}_z \rangle\} \\ \text{Re}\{\langle \hat{J}_x \hat{J}_y \rangle\} & \langle \hat{J}_y^2 \rangle & \text{Re}\{\langle \hat{J}_y \hat{J}_z \rangle\} \\ \text{Re}\{\langle \hat{J}_x \hat{J}_z \rangle\} & \text{Re}\{\langle \hat{J}_y \hat{J}_z \rangle\} & \langle \hat{J}_z^2 \rangle \end{pmatrix} - \begin{pmatrix} \langle \hat{J}_x \rangle \\ \langle \hat{J}_y \rangle \\ \langle \hat{J}_z \rangle \end{pmatrix} (\langle \hat{J}_x \rangle \quad \langle \hat{J}_y \rangle \quad \langle \hat{J}_z \rangle), \quad (\text{F2})$$

where the expectation value is taken against pure states,  $\langle \cdot \rangle := \langle\Psi| \cdot |\Psi\rangle$ . The maximal value of the QFI achieved is consequently  $F_Q[|\Psi\rangle, \mathbf{J}_{\mathbf{v}_{\text{best}=\mathbf{v}_{\text{max}}}}] = 4\mathbf{v}_{\text{max}}^T \mathcal{C} \mathbf{v}_{\text{max}} = 4\lambda_{\text{max}}(\mathcal{C})$ , where  $\lambda_{\text{max}}$  indicates maximal eigenvalue. Here, we will introduce two basic examples of such results:

(i) A coherent spin state pointing along  $x$  of length  $J = L/2$ ,  $|+\rangle^{\otimes L}$ . This initial state was chosen to start OAT dynamics, Eq. (9) (in the main text). In such case, the optimal axis is any orientation orthogonal to  $x$ , i.e., contained in the  $yz$  plane. The QFI achieved is exactly  $L$ , which is the maximal value that can be reached by separable states (also known as shot-noise limit) [120].

(ii) The GHZ or cat state aligned along  $x$ ,  $(|+\rangle^{\otimes L} + e^{i\phi}|-\rangle^{\otimes L})/\sqrt{2}$ , which is realized at time  $t = \pi/2$  during the OAT dynamics. Now, the optimal generator points in the  $x$  direction, and the corresponding QFI (variance times four) is  $L^2$ . Such value actually is the maximal QFI achievable within the quantum framework and requires genuine  $L$ -partite entanglement [118,119].

Since the OAT reconstructed states  $\hat{\rho}$  are of high purity, the same procedure can approximate the optimal orientation  $\mathbf{v}$ . However, the QFI results are evaluated exactly as per Eq. (F1). For further study, we refer the reader to the excellent review of Ref. [77].

## APPENDIX G: OAT STATE ANALYSIS FOR ZERO CHANNEL AND ZERO MEASUREMENT NOISE

In this section, we consider an optimal scenario, when only statistical sampling noise is being considered. To prepare our data, we also consider SIC-POVM operators this time, given their higher performance in tomography reconstruction tasks. The quality of reconstruction is shown in Table I. First, we verify that the neural network can improve substantially for the OAT states, even though no examples of such states were given in the training phase, which relied only on Haar-random states. Moreover, the OAT-averaged reconstruction values exceed the Haar reconstruction ones. We conjecture that this stems from the bosonic symmetry exhibited by the OAT states. This symmetry introduces redundancies in the density matrix, which might help the neural network detect errors produced by statistical noise. Finally, let us highlight that the network also displays good robustness to noise. In fact, when we feed the same network with states prepared for  $N_{\text{trial}} = 10^3$  trials, we increase the fidelity of the reconstruction from 67% to 87%.

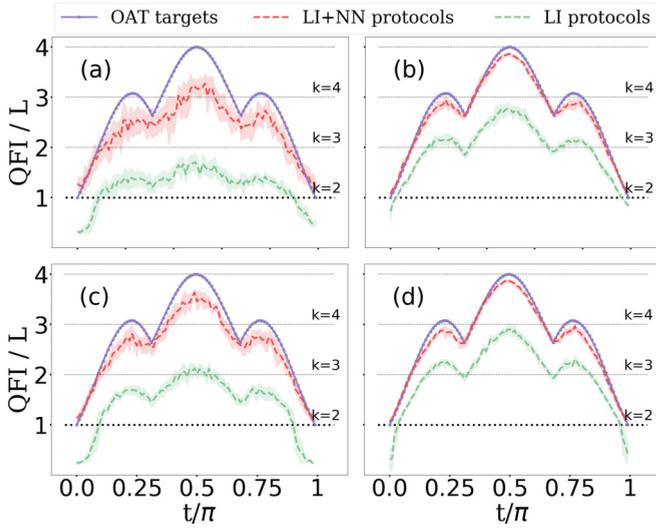


FIG. 4. Time evolution of the normalized QFI during the OAT protocol for  $L = 4$  qubits system. Solid blue lines represent QFI calculated for target quantum states. The mean values of QFI calculated from tomographically reconstructed density matrices are denoted by green-dashed (reconstruction via LI), and red-dotted lines (reconstruction via neural network postprocessed LI outputs). Shaded areas mark one standard deviation after averaging over 10 reconstructions. Panels (a) and (b) correspond to LI protocol with SIC-POVM data, whereas (c) and (d) denote LI reconstruction inferred from Pauli measurements. In the upper row, the left (right) column corresponds to  $N_{\text{trial}} = 10^3$  ( $10^4$ ) trials; in the lower row, the left (right) column reproduces the LI initial fidelity reconstruction of  $\sim 74\%$  ( $\sim 86\%$ ). The red lines represent the whole setup with neural network postprocessing of data from corresponding green lines, indicating improvement over the LI method. The neural network advantage over the bare LI method can be characterized by entanglement depth certification, as shown by the horizontal lines denoting the entanglement depth bounds ranging from the separable limit (bottom line, bold) to the genuine  $L$ -body limit (top line). In particular, the presence of entanglement,  $k \geq 2$ , is witnessed by  $\text{QFI} > L$ , as shown by the violation of the separable bound (bold horizontal line).

## APPENDIX H: BENCHMARKING THE TRANSFORMER LAYER

The transformer architecture represented a breakthrough in deep learning for sequential learning problems, such as natural language processing [127]. It has been used for QST problems that involve an unsupervised approach based on the POVM ansatz [42], and more complex architectures to extract diffusion parameters of Brownian motion [66]. Here, we show the usefulness of the transformer architecture in improving our QST *matrix-to-matrix* protocol. This becomes especially pronounced in the regime of highly correlated, that is, entangled quantum states, as shown in the latter part of this Appendix. To test its role, we explore and compare our transformer model with two different convolutional-only architectures (CNN) based on their reconstruction ability. In particular, we benchmark against two setups: (i) a 4-layer convolutional neural network, where the transformer layer is replaced with two convolutional ones,

$$l_{\theta}^{(4)}[\vec{C}_{\rho}] = (\gamma(h_{\text{cnn}}))^{\circ 4}[\vec{C}_{\rho}] \quad (\text{H1})$$

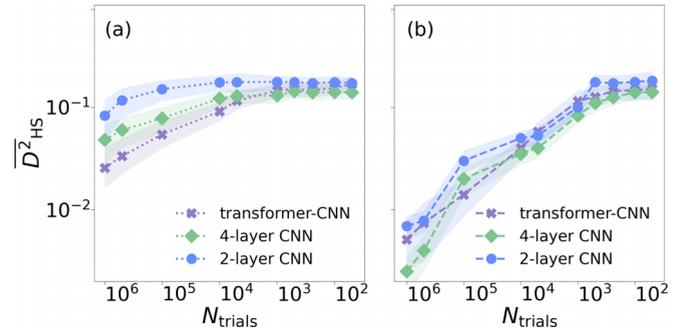


FIG. 5. Comparison of the efficiency of QST reconstruction schemes evaluated using Hilbert-Schmidt distance square  $\overline{D}_{\text{HS}}^2$  for transformer-based, 2-layer, and 4-layer attention-free CNN models, averaged in 1000 mixed states. All models share an equivalent number of training parameters. (a) Average reconstruction values for the 10 different LI preprocessed test datasets. Similarly to Fig. 2, we vary the number of trials  $N_{\text{trials}}$  to analyze the reconstruction efficiency and also use states of dimension 9 for a direct comparison. (b) The same analysis applied to the models trained on the MLE preprocessed data. To summarize, only for the MLE preprocessed data, the 4-layer CNN model can outperform the transformer-based for  $N_{\text{trials}} = 10^6, 10^5$ , while for the LI preprocessed our network shows better outcomes.

and (ii) the simplest CNN model consisting of two convolutional layers, represented as

$$l_{\theta}^{(2)}[\vec{C}_{\rho}] = (\gamma(h_{\text{cnn}}))^{\circ 2}[\vec{C}_{\rho}], \quad (\text{H2})$$

where  $\gamma$  is the GELU activation function in both the cases.

We set an equivalent number of trainable parameters for all three architectures. We tested the two CNNs for the same datasets as previously used for mixed- and pure-state reconstruction. In the following, we analyze the performance of the models for mixed- and pure-state reconstruction.

### 1. Mixed state reconstruction

In Fig. 5, we show the reconstruction of mixed states using the three setups, namely our transformer and the two CNN architectures. Firstly, as shown in panel (a), for the LI preprocessed data, our transformer-based model outperforms CNNs by showing a higher expressivity in terms of better reconstruction ability for a large number of trials  $N_{\text{trials}} > 10^4$ . However, the three models are almost equivalent in the undersampled regime. Next, in panel (b), we show our numerical experiments conducted on the MLE preprocessed dataset, which has a comparatively lesser amount of noise. We can see that the three models are almost equivalent.

### 2. Pure state reconstruction

For Haar-random pure states, using any of the two CNN models for the QST protocol causes an evident drop of  $\sim 10\%$  in the fidelity of reconstruction compared to the our attention-based model for the undersampled regime ( $N_{\text{trial}} = 10^3$ ), as exemplified in Table II.

Finally, we apply our protocol to the pure-state reconstruction task generated from one-axis twisting dynamics. In Fig. 6 we present the quantum Fisher information extracted from the reconstructed states. A significant drop in quantum Fisher

TABLE II. Values of averaged fidelity and its standard deviation between the reconstructed and the target Haar pure states of size  $d = 16$  obtained by the two CNN architectures  $l_\theta^{(2)}, l_\theta^{(4)}$  used inside our matrix-to-matrix protocol. In the second row, the outcomes we obtained by applying the transformed-based model. Similar to Fig. 4, we check for a different number of trials  $N_{\text{trial}} = 10^5, 10^4, 10^3$  to analyze the performance of the different architectures. All the initial Born values are calculated via noiseless SIC-POVM.

	$10^5$ trials	$10^4$ trials	$10^3$ trials
CNNs	$94.2 \pm 3.0\%$	$87 \pm 2.2\%$	$68.0 \pm 4.0\%$
Transformer	$96.9 \pm 3.0\%$	$94.2 \pm 3.3\%$	$81.1 \pm 4.1\%$

information reconstruction is observed in using CNN-based architectures compared to the transformer architecture.

We conclude that our transformer-based architecture works significantly well compared to the only CNN-based architectures in the case of reconstructing pure OAT states, when the QFI is considered a figure of merit. In the other case of mixed states reconstruction, the three architectures are comparable, and the transformer-based approach works better in the LI-preprocessed states in the highly sampled regime.

#### APPENDIX I: IMPROVING THE AVERAGED MLE IN HS METRIC

The MLE is efficient asymptotically, with a number of measurements  $N \rightarrow \infty$ . However, for a few measurements, there is no guarantee that MLE performs best. Such a situation takes place in the undersampled regime. In terms of the HS metric, it is usually more convenient to ignore any knowledge and just take the maximally mixed state  $\mathbb{I}/d$  as an estimation. Here, we propose a simple method to interpolate between the two extreme results by depolarizing the MLE state  $\hat{\rho}$ ,

$$\hat{\rho}_p = p\hat{\rho} + (1-p)\mathbb{I}/d, \quad (11)$$

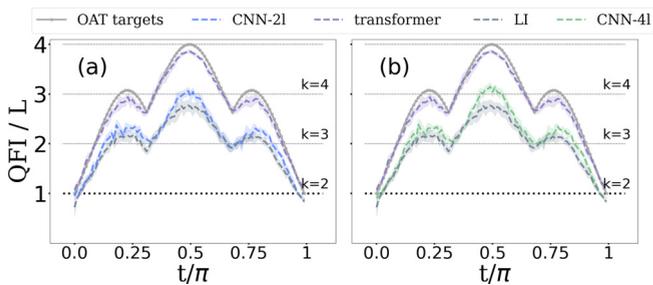


FIG. 6. Time evolution of the normalized QFI during the OAT protocol for four qubits. The dotted dark grey line represents the QFI calculated for the target quantum state, and the light grey dashed line is the QFI upon LI reconstruction (our minimal threshold). Panels (a) and (b) correspond to the QFI obtained for the states reconstructed by the 2-layers and 4-layers CNN respectively. We observe that, firstly, the transformer-based model outperforms the CNN models at all times with reconstruction ability very close to the OAT target states. Secondly, the CNN models perform equivalently irrespective of the number of layers in the architecture as shown in panels (a) and (b) for 2-layers and 4-layers respectively when considering QFI as our reconstruction metric.

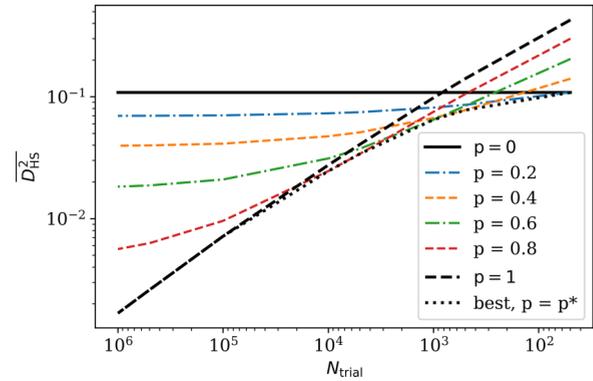


FIG. 7. Averaged HS distance of reconstructed MLE from the HS ensemble with  $d = 9$  as mixed according to Eq. (11) for different values of  $p$  (coloured lines). We highlight the limiting cases, namely  $p = 0$  (solid line); the average with respect  $\mathbb{I}/d$ , and  $p = 1$  (dashed): the MLE result. The envelope of such a family of lines is marked with a dotted line. Such bound can be realized with an optimal  $p^*$ , which depends on the number of trials via the reconstructed  $\{\hat{\rho}_{\text{MLE}}\}$ .

with  $0 \leq p \leq 1$ . The parameter  $p$  would then incorporate the statistical noise inherent to have a finite number of samples.

In Fig. 7 we show the average HS distance as a function of a number of trials  $N_{\text{trial}}$ , for different values of the parameter  $p$ . The first extreme case ( $p = 0$ ), i.e., maximally mixed state is presented as a horizontal solid line, while the second extreme case ( $p = 1$ ), i.e., the MLE reconstructed state, is depicted as a dashed line. All intermediate values of  $p$  form an envelope shape, corresponding to a critical value  $p^*$ , outperforming all other values of  $p$ .

To calculate the critical  $p^*$  let us notice that the average of the squared HS distance with respect to the state Eq. (11),  $D_p^2 = D_{\text{HS}}^2(\hat{\tau}, \hat{\rho}_p)$ , can be expressed as

$$D_p^2 = D_0^2 + (D_1^2 - D_0^2 - D_{01}^2)p + D_{01}^2 p^2, \quad (12)$$

where

$$D_0^2 = D_{\text{HS}}^2(\hat{\tau}, \mathbb{I}/d), \quad (13)$$

$$D_1^2 = D_{\text{HS}}^2(\hat{\tau}, \hat{\rho}_{\text{MLE}}), \quad (14)$$

$$D_{01}^2 = D_{\text{HS}}^2(\mathbb{I}/d, \hat{\rho}_{\text{MLE}}). \quad (15)$$

The Eq. (12) is equally valid in average as it is linear in the  $D^2$ 's, and we consider  $p$  a global parameter. Then,  $\overline{D_0^2}$  can be related to the average purity  $\overline{\text{Tr}(\hat{\tau}^2)}$  as  $D_0^2 = \text{Tr}(\hat{\tau}^2) - d^{-1}$ . Also, as depicted in Fig. 8, it can be interpreted geometrically as a triangle of sides' length  $\{D_0, D_1, D_{01}\}$  in which we want to find the point  $p$  in the segment 01, which is closer to the opposed vertex. Minimization of the same equation leads to the optimal  $p^*$ ,

$$p_* = \begin{cases} 0 & \overline{D_{01}^2} < \overline{D_1^2} - \overline{D_0^2} \\ \frac{\overline{D_{01}^2} + \overline{D_0^2} - \overline{D_1^2}}{2\overline{D_{01}^2}} & \text{otherwise} \\ 1 & \overline{D_{01}^2} < \overline{D_0^2} - \overline{D_1^2} \end{cases}, \quad (16)$$

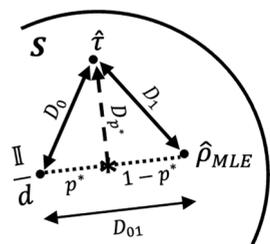


FIG. 8. Geometric interpretation of the optimal depolarization of the MLE state, such as incorporating the statistical noise stemming from a finite number of experimental runs.

yielding an optimal distance

$$\overline{D}_*^2 = \begin{cases} \overline{D}_0^2 & \overline{D}_{01}^2 < \overline{D}_1^2 - \overline{D}_0^2 \\ \overline{D}_0^2 - \frac{\overline{D}_0^2 - \overline{D}_1^2 - \overline{D}_{01}^2}{4\overline{D}_{01}^2} & \text{otherwise} \\ \overline{D}_1^2 & \overline{D}_{01}^2 < \overline{D}_0^2 - \overline{D}_1^2 \end{cases}, \quad (17)$$

In Fig. 9, we observe a realization of the nontrivial solution (otherwise case), which is better than both  $\overline{D}_0^2$  and  $\overline{D}_1^2$ .

Finally, one should check how this method relates to Bayesian approaches [128]. In particular, how to incorporate partial information about the ensemble, e.g., by only assuming as prior knowledge an average purity.

#### APPENDIX J: INTERPRETATION OF THE NEURAL NETWORK AS A CONDITIONAL “DEBIASER”

Our neural network takes as input a reconstructed state  $\hat{\rho}$  from the experimental results  $\mathbf{f}$  through some estimator (e.g., LI, MLE),  $Q[\mathbf{f}] = \hat{\rho}$ . The neural network returns a state  $\hat{\rho}$  that on average over the realizations (in  $\mathbf{f}$ ) better approximates the target  $\hat{\tau}$  than  $\hat{\rho}$ . In the following, we outline the situations in which we expect a poor performance of our algorithm. From the above setting, we immediately see that it is useless for the unbiased estimators.

*Observation 1.* If the estimator  $Q$  is unbiased, i.e.,  $\hat{\rho} = \hat{\tau}$ , no further improvement can be achieved with our approach. In fact, the mean already provides the best estimation.

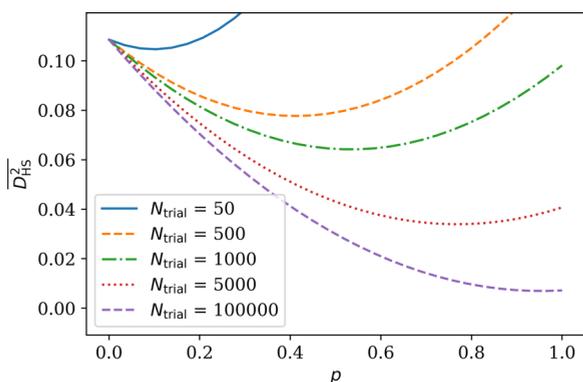


FIG. 9. Average reconstruction distance as a function of the mixing parameter  $p$  for a given set of number of trials. We verify the parabola curves of Eq. (12) and the nontrivial minima.

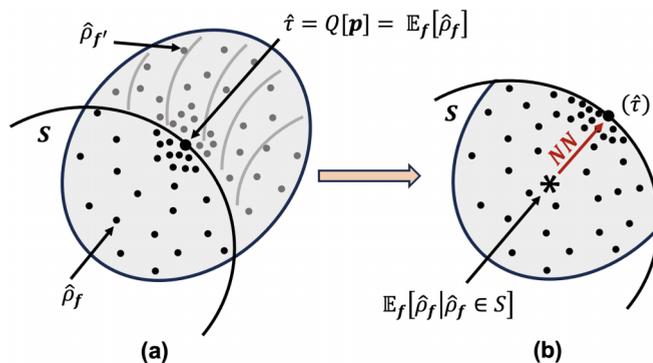


FIG. 10. Action of the neural network as a conditional debiaser. (a) Inference of the state  $\hat{\tau}$  by many finite-size realization  $\{\hat{\rho}_f\}_f$  not necessarily a proper state (i.e., it might be outside  $\mathcal{S}$ ). (b) Disregarding the nonphysical realizations results in a skewed conditional distribution whose mean is displaced from the true state. The action of the neural network is then to shift back the mean to the target state by drifting the distribution.

Consequently, if some enhancement is observed, the inference of the input state must be biased. The bias here comes from the requirement that  $\hat{\rho}$  must be a proper quantum state.

To observe why it is so, let us focus on the simplest case of two projective quantum measurements, i.e., spin measurements of an electron in the  $x$  and  $y$  directions. Although both measurements belong to the set  $[-1/2, 1/2]$ , not all pairs of measurement results are admissible. For example, there are no valid quantum states for which both measurements yield  $1/2$  since then the total spin would be larger than  $1/2$ .

Therefore, if we account for the inevitable noise present for finite statistics, the unphysicality of certain real-world measurements is unavoidable. There are two general QST strategies to overcome this obstacle—either discard the unphysical outcomes (e.g., MLE) or keep them all (e.g., LI). Any of these methods has its drawbacks, and, for finite statistics, one cannot have a strategy that satisfies both the linearity and the physicality of the predicted states [129].

*Observation 2.* Any estimator  $Q$ , which, as required by our neural network architecture, outputs a valid state  $\hat{\rho}_f$  for any physical  $\mathbf{f} \in \mathcal{F}$ , is biased.

This phenomenon becomes more prominent for  $\hat{\tau}$  close to the boundary of the quantum set  $\mathcal{S}$ , i.e., with high purity [130,131]. By convexity, the boundary of  $\mathcal{F}_S$  is attained by extremal elements of  $\mathcal{S}$ . From which we notice that the chance of nonquantum outcomes is higher and  $\mathbb{E}_f[\hat{\rho}_f | \hat{\rho}_f \in \mathcal{S}]$  becomes significantly displaced from  $\hat{\tau}$  (see sketch Fig. 10). In such terms, the task of the neural network can be interpreted as drifting the skewed probability distribution towards the true mean, by means, however, of only quantum outcomes. In other words:

From  $\{\mathbf{f} \in \mathcal{F}_S\}$ , find ensemble  $\{\mathbf{f}_{\text{NN}} \in \mathcal{F}_S\}$  such that the output bias  $|\mathbb{E}_{\mathbf{f}_{\text{NN}}}(\hat{\rho}_{\text{NN}}) - \hat{\tau}|$  is minimal.

As a result, generic mixed states are harder to improve than pure states. More studies are needed to confirm the said behavior. In particular, one has to verify that no improvement is possible if the reconstruction method to obtain  $\hat{\rho}_f$  already incorporates the skewness and how performance depends on the bias of the estimator  $Q$ .

- [1] O. Gühne and G. Tóth, Entanglement detection, *Phys. Rep.* **474**, 1 (2009).
- [2] M. P. da Silva, O. Landon-Cardinal, and D. Poulin, Practical characterization of quantum devices without tomography, *Phys. Rev. Lett.* **107**, 210404 (2011).
- [3] M. Kliesch and I. Roth, Theory of quantum system certification, *PRX Quantum* **2**, 010201 (2021).
- [4] A. Tavakoli, Semi-device-independent certification of independent quantum state and measurement devices, *Phys. Rev. Lett.* **125**, 150503 (2020).
- [5] N. Friis, G. Vitagliano, M. Malik, and M. Huber, Entanglement certification from theory to experiment, *Nat. Rev. Phys.* **1**, 72 (2019).
- [6] J. Eisert, D. Hangleiter, N. Walk, I. Roth, D. Markham, R. Parekh, U. Chabaud, and E. Kashefi, Quantum certification and benchmarking, *Nat. Rev. Phys.* **2**, 382 (2020).
- [7] O. M. Sotnikov, I. A. Iakovlev, A. A. Iliasov, M. I. Katsnelson, A. A. Bagrov, and V. V. Mazurenko, Certification of quantum states with hidden structure of their bitstrings, *npj Quantum Inf.* **8**, 41 (2022).
- [8] S. Chen, J. Li, B. Huang, and A. Liu, Tight bounds for quantum state certification with incoherent measurements, in *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, Denver, CO, USA, 2022), pp. 1205–1213.
- [9] A. Gočanin, I. Šupić, and B. Dakić, Sample-efficient device-independent quantum state verification and certification, *PRX Quantum* **3**, 010317 (2022).
- [10] E. C. Boghiu, F. Hirsch, P. S. Lin, M. T. Quintino, and J. Bowles, Device-independent and semi-device-independent entanglement certification in broadcast Bell scenarios, *SciPost Phys. Core* **6**, 028 (2023).
- [11] D. Hangleiter, M. Kliesch, M. Schwarz, and J. Eisert, Direct certification of a class of quantum simulations, *Quantum Sci. Technol.* **2**, 015004 (2017).
- [12] M. Płodzień, J. Chwedeńczuk, M. Lewenstein, and G. Rajchel-Mieldzioc, Entanglement classification and *non-k*-separability certification via Greenberger-Horne-Zeilinger-class fidelity, [arXiv:2406.10662](https://arxiv.org/abs/2406.10662).
- [13] U. Leonhardt, Quantum-state tomography and discrete Wigner function, *Phys. Rev. Lett.* **74**, 4101 (1995).
- [14] A. G. White, D. F. V. James, P. H. Eberhard, and P. G. Kwiat, Nonmaximally entangled states: Production, characterization, and utilization, *Phys. Rev. Lett.* **83**, 3103 (1999).
- [15] C. F. Roos, G. P. T. Lancaster, M. Riebe, H. Häffner, W. Hänsel, S. Gulde, C. Becher, J. Eschner, F. Schmidt-Kaler, and R. Blatt, Bell states of atoms with ultralong lifetimes and their tomographic state analysis, *Phys. Rev. Lett.* **92**, 220402 (2004).
- [16] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Chekhlov, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt *et al.*, Scalable multiparticle entanglement of trapped ions, *Nature (London)* **438**, 643 (2005).
- [17] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, Quantum state tomography via compressed sensing, *Phys. Rev. Lett.* **105**, 150401 (2010).
- [18] D. Gross, Recovering low-rank matrices from few coefficients in any basis, *IEEE Trans. Inf. Theory* **57**, 1548 (2011).
- [19] G. Tóth, W. Wiczecek, D. Gross, R. Kriscsek, C. Schwemmer, and H. Weinfurter, Permutationally invariant quantum tomography, *Phys. Rev. Lett.* **105**, 250403 (2010).
- [20] T. Moroder, P. Hyllus, G. Tóth, C. Schwemmer, A. Niggebaum, S. Gaile, O. Gühne, and H. Weinfurter, Permutationally invariant state reconstruction, *New J. Phys.* **14**, 105001 (2012).
- [21] 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).
- [22] T. Baumgratz, D. Gross, M. Cramer, and M. B. Plenio, Scalable reconstruction of density matrices, *Phys. Rev. Lett.* **111**, 020401 (2013).
- [23] 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).
- [24] A. Acín, I. Bloch, H. Buhrman, T. Calarco, C. Eichler, J. Eisert, D. Esteve, N. Gisin, S. J. Glaser, F. Jelezko *et al.*, The quantum technologies roadmap: A European community view, *New J. Phys.* **20**, 080201 (2018).
- [25] A. Kinos, D. Hunger, R. Kolesov, K. Mølmer, H. de Riedmatten, P. Goldner, A. Tallaire, L. Morvan, P. Berger, S. Welinski *et al.*, Roadmap for rare-earth quantum computing, [arXiv:2103.15743](https://arxiv.org/abs/2103.15743).
- [26] A. Laucht, F. Hohls, N. Ubbelohde, M. F. Gonzalez-Zalba, D. J. Reilly, S. Stobbe, T. Schröder, P. Scarlino, J. V. Koski, A. Dzurak *et al.*, Roadmap on quantum nanotechnologies, *Nanotechnology* **32**, 162003 (2021).
- [27] C. Becher, W. Gao, S. Kar, C. D. Marciniak, T. Monz, J. G. Bartholomew, P. Goldner, H. Loh, E. Marcellina, K. E. J. Goh *et al.*, 2023 roadmap for materials for quantum technologies, *Mater. Quantum Tech.* **3**, 012501 (2023).
- [28] J. Fraxanet, T. Salamon, and M. Lewenstein, The coming decades of quantum simulation, in *Sketches of Physics* (Springer International Publishing, New York, 2023), pp. 85–125.
- [29] 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).
- [30] F. Huszár and N. M. T. Houlshby, Adaptive Bayesian quantum tomography, *Phys. Rev. A* **85**, 052120 (2012).
- [31] K. S. Kravtsov, S. S. Straupe, I. V. Radchenko, N. M. T. Houlshby, F. Huszár, and S. P. Kulik, Experimental adaptive Bayesian tomography, *Phys. Rev. A* **87**, 062122 (2013).
- [32] C. Granade, C. Ferrie, and S. T. Flammia, Practical adaptive quantum tomography, *New J. Phys.* **19**, 113017 (2017).
- [33] C. Schwemmer, G. Tóth, A. Niggebaum, T. Moroder, D. Gross, O. Gühne, and H. Weinfurter, Experimental comparison of efficient tomography schemes for a six-qubit state, *Phys. Rev. Lett.* **113**, 040503 (2014).
- [34] H. Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [35] S. Aaronson, Shadow tomography of quantum states, in *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018* (Association for Computing Machinery, New York, NY, USA, 2018), pp. 325–338.

- [36] P. Yang, M. Yu, R. Betzholz, C. Arenz, and J. Cai, Complete quantum-state tomography with a local random field, *Phys. Rev. Lett.* **124**, 010405 (2020).
- [37] 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).
- [38] H.-Y. Huang, J. Preskill, and M. Soleimanifar, Certifying almost all quantum states with few single-qubit measurements, [arXiv:2404.07281](https://arxiv.org/abs/2404.07281).
- [39] A. Dawid, J. Arnold, B. Requena, A. Gresch, M. Płodzień, K. Donatella, K. A. Nicoli, P. Stornati, R. Koch, M. Büttner *et al.*, Modern applications of machine learning in quantum sciences, [arXiv:2204.04198](https://arxiv.org/abs/2204.04198).
- [40] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, *Nat. Phys.* **14**, 447 (2018).
- [41] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, Reconstructing quantum states with generative models, *Nat. Mach. Intell.* **1**, 155 (2019).
- [42] P. Cha, P. Ginsparg, F. Wu, J. Carrasquilla, P. L. McMahon, and E.-A. Kim, Attention-based quantum tomography, *Mach. Learn.: Sci. Technol.* **3**, 01LT01 (2022).
- [43] A. Melkani, C. Gneiting, and F. Nori, Eigenstate extraction with neural-network tomography, *Phys. Rev. A* **102**, 022412 (2020).
- [44] T. Schmale, M. Reh, and M. Gärttner, Efficient quantum state tomography with convolutional neural networks, *npj Quantum Inf.* **8**, 115 (2022).
- [45] T. Xin, S. Lu, N. Cao, G. Anikeeva, D. Lu, J. Li, G. Long, and B. Zeng, Local-measurement-based quantum state tomography via neural networks, *npj Quantum Inf.* **5**, 109 (2019).
- [46] A. Lidiak, C. Jameson, Z. Qin, G. Tang, M. B. Wakin, Z. Zhu, and Z. Gong, Quantum state tomography with tensor train cross approximation, [arXiv:2207.06397](https://arxiv.org/abs/2207.06397).
- [47] 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).
- [48] A. A. Akhtar, H.-Y. Hu, and Y.-Z. You, Scalable and flexible classical shadow tomography with tensor networks, *Quantum* **7**, 1026 (2023).
- [49] W.-j. Li, K. Xu, H. Fan, S.-j. Ran, and G. Su, Efficient quantum mixed-state tomography with unsupervised tensor network machine learning, [arXiv:2308.06900](https://arxiv.org/abs/2308.06900).
- [50] S. Kuzmin, V. Mikhailova, I. Dyakonov, and S. Straupe, Learning the tensor network model of a quantum state using a few single-qubit measurements, *Phys. Rev. A* **109**, 052616 (2024).
- [51] A. Palmieri, E. Kovlakov, F. Bianchi, D. Yudin, S. Straupe, J. D. Biamonte, and S. Kulik, Experimental neural network enhanced quantum tomography, *npj Quantum Inf.* **6**, 20 (2020).
- [52] C. Pan and J. Zhang, Deep learning-based quantum state tomography with imperfect measurement, *Int. J. Theor. Phys.* **61**, 227 (2022).
- [53] D. Koutný, L. Motka, Z. Hradil, J. Řeháček, and L. L. Sánchez-Soto, Neural-network quantum state tomography, *Phys. Rev. A* **106**, 012409 (2022).
- [54] S. Ahmed, C. S. Muñoz, F. Nori, and A. F. Kockum, Quantum state tomography with conditional generative adversarial networks, *Phys. Rev. Lett.* **127**, 140502 (2021).
- [55] H. Ma, Z. Sun, D. Dong, C. Chen, and H. Rabitz, Tomography of quantum states from structured measurements via quantum-aware transformer, [arXiv:2305.05433](https://arxiv.org/abs/2305.05433).
- [56] T. Westerhout, N. Astrakhantsev, K. S. Tikhonov, M. I. Katsnelson, and A. A. Bagrov, Generalization properties of neural network approximations to frustrated magnet ground states, *Nat. Commun.* **11**, 1593 (2020).
- [57] J. Guerin, K. Delmas, R. Ferreira, and J. Guiochet, **Out-of-distribution detection is not all you need**, in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI, Washington, DC, 2023)*, Vol. 37, pp. 14829–14837.
- [58] N. Bent, H. Qassim, A. A. Tahir, D. Sych, G. Leuchs, L. L. Sánchez-Soto, E. Karimi, and R. W. Boyd, Experimental realization of quantum tomography of photonic qudits via symmetric informationally complete positive operator-valued measures, *Phys. Rev. X* **5**, 041006 (2015).
- [59] R. Stricker, M. Meth, L. Postler, C. Edmunds, C. Ferrie, R. Blatt, P. Schindler, T. Monz, R. Kueng, and M. Ringbauer, Experimental single-setting quantum state tomography, *PRX Quantum* **3**, 040310 (2022).
- [60] C. A. Fuchs and R. Schack, Quantum-Bayesian coherence, *Rev. Mod. Phys.* **85**, 1693 (2013).
- [61] C. Fuchs, M. Hoang, and B. Stacey, The SIC question: History and state of play, *Axioms* **6**, 21 (2017).
- [62] L. Peter and D. Petrosyan, *Fundamentals of Quantum Optics and Quantum Information* (Springer, Berlin, 2007).
- [63] If it is rank deficient (e.g., for pure states), we add a small correction of amplitude order  $10^{-5}$  to ensure state positivity.
- [64] After removal of negative eigenvalues.
- [65] Y. Zhu, Y.-D. Wu, G. Bai, D.-S. Wang, Y. Wang, and G. Chiribella, Flexible learning of quantum states with generative query neural networks, *Nat. Commun.* **13**, 6222 (2022).
- [66] B. Requena, S. Masó-Oriols, J. Bertran, M. Lewenstein, C. Manzo, and G. Muñoz-Gil, Inferring pointwise diffusion properties of single trajectories with deep learning, *Biophys. J.* **122**, 4360 (2023).
- [67] D. Bahdanau, K. Cho, and Y. Bengio, Neural machine translation by jointly learning to align and translate, [arXiv:1409.0473](https://arxiv.org/abs/1409.0473).
- [68] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, Attention is all you need, [arXiv:1706.03762](https://arxiv.org/abs/1706.03762).
- [69] V. P. Dwivedi and X. Bresson, A generalization of transformer networks to graphs, [arXiv:2012.09699](https://arxiv.org/abs/2012.09699).
- [70] D. Hendrycks and K. Gimpel, Gaussian error linear units (GELUs), [arXiv:1606.08415](https://arxiv.org/abs/1606.08415).
- [71] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
- [72] Notice that  $C_{\hat{\rho}}$  is not necessarily a Cholesky matrix in its canonical form (e.g., with diagonal elements positive), as such constraints are not imposed. However, this fact does not pose further issues as in any case, by construction, postprocessing will lead to a proper state  $\hat{\rho}$ .
- [73] T. Liu, S. Liu, H. Li, H. Li, K. Huang, Z. Xiang, X. Song, K. Xu, D. Zheng, and H. Fan, Observation of entanglement transition of pseudo-random mixed states, *Nat. Commun.* **14**, 1971 (2023).

- [74] H. Tang, L. Bianchi, T.-Y. Wang, X.-W. Shang, X. Tan, W.-H. Zhou, Z. Feng, A. Pal, H. Li, C.-Q. Hu, M. S. Kim, and X.-M. Jin, Generating Haar-uniform randomness using stochastic quantum walks on a photonic chip, *Phys. Rev. Lett.* **128**, 050503 (2022).
- [75] M. Kitagawa and M. Ueda, Squeezed spin states, *Phys. Rev. A* **47**, 5138 (1993).
- [76] D. J. Wineland, J. J. Bollinger, W. M. Itano, and D. J. Heinzen, Squeezed atomic states and projection noise in spectroscopy, *Phys. Rev. A* **50**, 67 (1994).
- [77] L. Pezzè, A. Smerzi, M. K. Oberthaler, R. Schmied, and P. Treutlein, Quantum metrology with nonclassical states of atomic ensembles, *Rev. Mod. Phys.* **90**, 035005 (2018).
- [78] F. Wolfgramm, A. Cerè, F. A. Beduini, A. Predojević, M. Koschorreck, and M. W. Mitchell, Squeezed-light optical magnetometry, *Phys. Rev. Lett.* **105**, 053601 (2010).
- [79] G. Müller-Rigat, A. K. Srivastava, S. Kurdziałek, G. Rajchel-Mieldzioc, M. Lewenstein, and I. Frérot, Certifying the quantum Fisher information from a given set of mean values: A semidefinite programming approach, *Quantum* **7**, 1152 (2023).
- [80] J. Tura, R. Augusiak, A. B. Sainz, T. Vértesi, M. Lewenstein, and A. Acín, Detecting nonlocality in many-body quantum states, *Science* **344**, 1256 (2014).
- [81] R. Schmied, J.-D. Bancal, B. Allard, M. Fadel, V. Scarani, P. Treutlein, and N. Sangouard, Bell correlations in a Bose-Einstein condensate, *Science* **352**, 441 (2016).
- [82] A. Aloy, J. Tura, F. Baccari, A. Acín, M. Lewenstein, and R. Augusiak, Device-independent witnesses of entanglement depth from two-body correlators, *Phys. Rev. Lett.* **123**, 100507 (2019).
- [83] F. Baccari, J. Tura, M. Fadel, A. Aloy, J.-D. Bancal, N. Sangouard, M. Lewenstein, A. Acín, and R. Augusiak, Bell correlation depth in many-body systems, *Phys. Rev. A* **100**, 022121 (2019).
- [84] J. Tura, A. Aloy, F. Baccari, A. Acín, M. Lewenstein, and R. Augusiak, Optimization of device-independent witnesses of entanglement depth from two-body correlators, *Phys. Rev. A* **100**, 032307 (2019).
- [85] M. Żukowski and Č. Brukner, Bell's theorem for general  $N$ -qubit states, *Phys. Rev. Lett.* **88**, 210401 (2002).
- [86] E. G. Cavalcanti, C. J. Foster, M. D. Reid, and P. D. Drummond, Bell inequalities for continuous-variable correlations, *Phys. Rev. Lett.* **99**, 210405 (2007).
- [87] Q. Y. He, P. D. Drummond, and M. D. Reid, Entanglement, EPR steering, and Bell-nonlocality criteria for multipartite higher-spin systems, *Phys. Rev. A* **83**, 032120 (2011).
- [88] E. G. Cavalcanti, Q. Y. He, M. D. Reid, and H. M. Wiseman, Unified criteria for multipartite quantum nonlocality, *Phys. Rev. A* **84**, 032115 (2011).
- [89] A. Niezgodą, M. Panfil, and J. Chwedeńczuk, Quantum correlations in spin chains, *Phys. Rev. A* **102**, 042206 (2020).
- [90] A. Niezgodą and J. Chwedeńczuk, Many-body nonlocality as a resource for quantum-enhanced metrology, *Phys. Rev. Lett.* **126**, 210506 (2021).
- [91] J. Chwedeńczuk, Many-body Bell inequalities for bosonic qubits, *SciPost Phys. Core* **5**, 025 (2022).
- [92] G. Müller-Rigat, A. Aloy, M. Lewenstein, and I. Frérot, Inferring nonlinear many-body Bell inequalities from average two-body correlations: Systematic approach for arbitrary spin- $j$  ensembles, *PRX Quantum* **2**, 030329 (2021).
- [93] D. J. Wineland, J. J. Bollinger, W. M. Itano, F. L. Moore, and D. J. Heinzen, Spin squeezing and reduced quantum noise in spectroscopy, *Phys. Rev. A* **46**, R6797 (1992).
- [94] K. Gietka, P. Szańkowski, T. Wasak, and J. Chwedeńczuk, Quantum-enhanced interferometry and the structure of twisted states, *Phys. Rev. A* **92**, 043622 (2015).
- [95] Y. Li, Y. Castin, and A. Sinatra, Optimum spin squeezing in Bose-Einstein condensates with particle losses, *Phys. Rev. Lett.* **100**, 210401 (2008).
- [96] Y. Li, P. Treutlein, J. Reichel, and A. Sinatra, Spin squeezing in a bimodal condensate: Spatial dynamics and particle losses, *Eur. Phys. J. B* **68**, 365 (2009).
- [97] M. Wang, W. Qu, P. Li, H. Bao, V. Vuletić, and Y. Xiao, Two-axis-twisting spin squeezing by multipass quantum erasure, *Phys. Rev. A* **96**, 013823 (2017).
- [98] D. Kajtoch, E. Witkowska, and A. Sinatra, Spin-squeezed atomic crystal, *EPL (Europhys. Lett.)* **123**, 20012 (2018).
- [99] M. Schulte, C. Lisdat, P. O. Schmidt, U. Sterr, and K. Hammerer, Prospects and challenges for squeezing-enhanced optical atomic clocks, *Nat. Commun.* **11**, 5955 (2020).
- [100] K. Gietka, A. Usui, J. Deng, and T. Busch, Simulating the same physics with two distinct Hamiltonians, *Phys. Rev. Lett.* **126**, 160402 (2021).
- [101] T. Comparin, F. Mezzacapo, and T. Roscilde, Robust spin squeezing from the tower of states of  $U(1)$ -symmetric spin Hamiltonians, *Phys. Rev. A* **105**, 022625 (2022).
- [102] M. F. Riedel, P. Böhi, Y. Li, T. W. Hänsch, A. Sinatra, and P. Treutlein, Atom-chip-based generation of entanglement for quantum metrology, *Nature (London)* **464**, 1170 (2010).
- [103] C. Gross, T. Zibold, E. Nicklas, J. Estève, and M. K. Oberthaler, Nonlinear atom interferometer surpasses classical precision limit, *Nature (London)* **464**, 1165 (2010).
- [104] C. D. Hamley, C. S. Gerving, T. M. Hoang, E. M. Bookjans, and M. S. Chapman, Spin-nematic squeezed vacuum in a quantum gas, *Nat. Phys.* **8**, 305 (2012).
- [105] A. Qu, B. Evrard, J. Dalibard, and F. Gerbier, Probing spin correlations in a Bose-Einstein condensate near the single-atom level, *Phys. Rev. Lett.* **125**, 033401 (2020).
- [106] I. D. Leroux, M. H. Schleier-Smith, and V. Vuletić, Implementation of cavity squeezing of a collective atomic spin, *Phys. Rev. Lett.* **104**, 073602 (2010).
- [107] K. Maussang, G. E. Marti, T. Schneider, P. Treutlein, Y. Li, A. Sinatra, R. Long, J. Estève, and J. Reichel, Enhanced and reduced atom number fluctuations in a BEC splitter, *Phys. Rev. Lett.* **105**, 080403 (2010).
- [108] P. He, M. A. Perlin, S. R. Muleady, R. J. Lewis-Swan, R. B. Hutson, J. Ye, and A. M. Rey, Engineering spin squeezing in a 3D optical lattice with interacting spin-orbit-coupled fermions, *Phys. Rev. Res.* **1**, 033075 (2019).
- [109] M. Płodzień, M. Kościelski, E. Witkowska, and A. Sinatra, Producing and storing spin-squeezed states and Greenberger-Horne-Zeilinger states in a one-dimensional optical lattice, *Phys. Rev. A* **102**, 013328 (2020).
- [110] M. Płodzień, M. Lewenstein, E. Witkowska, and J. Chwedeńczuk, One-axis twisting as a method of generating

- many-body Bell correlations, *Phys. Rev. Lett.* **129**, 250402 (2022).
- [111] M. Płodzień, T. Wasak, E. Witkowska, M. Lewenstein, and J. Chwedeńczuk, Generation of scalable many-body bell correlations in spin chains with short-range two-body interactions, *Phys. Rev. Res.* **6**, 023050 (2024).
- [112] M. Mamaev, I. Kimchi, R. M. Nandkishore, and A. M. Rey, Tunable-spin-model generation with spin-orbit-coupled fermions in optical lattices, *Phys. Rev. Res.* **3**, 013178 (2021).
- [113] T. H. Yanes, M. Płodzień, M. M. Sinkevičienė, G. Žlabys, G. Juzeliūnas, and E. Witkowska, One- and two-axis squeezing via laser coupling in an atomic Fermi-Hubbard model, *Phys. Rev. Lett.* **129**, 090403 (2022).
- [114] M. Dziurawiec, T. H. Yanes, M. Płodzień, M. Gajda, M. Lewenstein, and E. Witkowska, Accelerating many-body entanglement generation by dipolar interactions in the Bose-Hubbard model, *Phys. Rev. A* **107**, 013311 (2023).
- [115] T. H. Yanes, G. Žlabys, M. Płodzień, D. Burba, M. M. Sinkevičienė, E. Witkowska, and G. Juzeliūnas, Spin squeezing in open Heisenberg spin chains, *Phys. Rev. B* **108**, 104301 (2023).
- [116] M. Yu, D. Li, J. Wang, Y. Chu, P. Yang, M. Gong, N. Goldman, and J. Cai, Experimental estimation of the quantum Fisher information from randomized measurements, *Phys. Rev. Res.* **3**, 043122 (2021).
- [117] M. Yu, Y. Liu, P. Yang, M. Gong, Q. Cao, S. Zhang, H. Liu, M. Heyl, T. Ozawa, N. Goldman, and J. Cai, Quantum Fisher information measurement and verification of the quantum Cramér–Rao bound in a solid-state qubit, *npj Quantum Inf.* **8**, 56 (2022).
- [118] P. Hyllus, W. Laskowski, R. Krischek, C. Schwemmer, W. Wieczorek, H. Weinfurter, L. Pezzé, and A. Smerzi, Fisher information and multiparticle entanglement, *Phys. Rev. A* **85**, 022321 (2012).
- [119] G. Tóth, Multipartite entanglement and high-precision metrology, *Phys. Rev. A* **85**, 022322 (2012).
- [120] L. Pezzé and A. Smerzi, Entanglement, nonlinear dynamics, and the Heisenberg limit, *Phys. Rev. Lett.* **102**, 100401 (2009).
- [121] A. Macarone-Palmieri, Data generation, test-sets, and model training codes for attention-based-QST, available at <https://github.com/AdriQD/attentionQST> (2023).
- [122] M. Oszmaniec, R. Augusiak, C. Gogolin, J. Kołodyński, A. Acín, and M. Lewenstein, Random bosonic states for robust quantum metrology, *Phys. Rev. X* **6**, 041044 (2016).
- [123] J. A. Smolin, J. M. Gambetta, and G. Smith, Efficient method for computing the maximum-likelihood quantum state from measurements with additive Gaussian noise, *Phys. Rev. Lett.* **108**, 070502 (2012).
- [124] M. ApS, The MOSEK optimization toolbox for MATLAB manual, Version 9.0, available at: <https://docs.mosek.com/latest/toolbox/index.html> (2019).
- [125] I. Bengtsson and K. Życzkowski, *Geometry of Quantum States: An Introduction to Quantum Entanglement*, 2nd ed. (Cambridge University Press, Cambridge, 2017).
- [126] F. Mezzadri, How to generate random matrices from the classical compact groups, [arXiv:math-ph/0609050](https://arxiv.org/abs/math-ph/0609050).
- [127] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. U. Kaiser, and I. Polosukhin, Attention is all you need, in *Advances in Neural Information Processing Systems*, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., New York, 2017), Vol. 30.
- [128] C. Granade, J. Combes, and D. G. Cory, Practical Bayesian tomography, *New J. Phys.* **18**, 033024 (2016).
- [129] G. B. Silva, S. Glancy, and H. M. Vasconcelos, Investigating bias in maximum-likelihood quantum-state tomography, *Phys. Rev. A* **95**, 022107 (2017).
- [130] C. Ferrie and R. Blume-Kohout, Maximum likelihood quantum state tomography is inadmissible, [arXiv:1808.01072](https://arxiv.org/abs/1808.01072).
- [131] T. L. Scholten and R. Blume-Kohout, Behavior of the maximum likelihood in quantum state tomography, *New J. Phys.* **20**, 023050 (2018).