

Neural-network quantum state tomography in a two-qubit experiment

Marcel Neugebauer,¹ Laurin Fischer^{1,*}, Alexander Jäger¹, Stefanie Czischek²,
Selim Jochim,¹ Matthias Weidemüller¹ and Martin Gärttner^{2,1,3,†}

¹Physikalisches Institut, Universität Heidelberg, Im Neuenheimer Feld 226, 69120 Heidelberg, Germany

²Kirchhoff-Institut für Physik, Ruprecht-Karls-Universität Heidelberg, Im Neuenheimer Feld 227, 69120 Heidelberg, Germany

³Institut für Theoretische Physik, Ruprecht-Karls-Universität Heidelberg, Philosophenweg 16, 69120 Heidelberg, Germany



(Received 31 July 2020; accepted 14 September 2020; published 9 October 2020)

We study the performance of efficient quantum state tomography methods based on neural-network quantum states using measured data from a two-photon experiment. Machine-learning-inspired variational methods provide a promising route towards scalable state characterization for quantum simulators. While the power of these methods has been demonstrated on synthetic data, applications to real experimental data remain scarce. We benchmark and compare several such approaches by applying them to measured data from an experiment producing two-qubit entangled states. We find that in the presence of experimental imperfections and noise, confining the variational manifold to physical states, i.e., to positive semidefinite density matrices, greatly improves the quality of the reconstructed states but renders the learning procedure more demanding. Including additional, possibly unjustified, constraints, such as assuming pure states, facilitates learning, but also biases the estimator.

DOI: [10.1103/PhysRevA.102.042604](https://doi.org/10.1103/PhysRevA.102.042604)

I. INTRODUCTION

Quantum simulation experiments allow the investigation of quantum many-body problems beyond the regime accessible with classical computers, with potential applications in fundamental physics, quantum chemistry, and materials science [1]. With quantum simulators reaching sizes of tens and soon hundreds of qubits, the question arises how one can validate such devices to confirm that they faithfully emulate the targeted model. Using full state tomography for characterizing the prepared quantum states is not practical beyond the few-particle regime as the number of measurements required scales exponentially with the system size. Thus, efficient methods for state characterization are needed, meaning that the required resources, namely the number of experimental repetitions and the postprocessing time, should scale economically, i.e., subexponentially, with the system's size. Several approaches have been proposed in recent years to solve this problem, including permutationally invariant tomography [2,3], compressed sensing [4,5], and tomographic schemes based on tensor network states [6–8]. The central idea underlying these methods is to make a variational ansatz for the prepared state and adjust the variational parameters to best fit the measured data. To be efficient, the number of variational parameters should scale polynomially in the system size. Consequently, one searches the best fitting state among a restricted set of quantum states. The justification for imposing such constraints is usually given by *a priori* known

physical properties, e.g., a small amount of entanglement in the case of tensor network states.

A more flexible class of efficient tomography methods, requiring less prior knowledge about the prepared state, uses a data-driven approach inspired by rapid advances in machine learning. These techniques, known as neural-network quantum state (NQS) tomography [9], use generative modeling, a strategy for unsupervised learning, to reconstruct unknown quantum states. NQS exploit the expressivity of neural-network models for finding efficient representations of quantum states. They were introduced by Carleo and Troyer [10] and have subsequently been investigated intensively (see Refs. [9,11,12] for recent reviews). When employed for quantum state tomography, the model parameters of an NQS are adjusted, or learned, to best fit a set of experimentally measured data. If one assumes that the experimentally prepared state is pure, the original NQS ansatz can be employed for this [13]. However, in any experimentally realistic scenario noise and coupling to an uncontrolled environment require the adoption of a mixed state description. Thus, the density matrix of the state has to be parameterized, which can be achieved either by the so-called latent state purification ansatz [14] or by directly modeling the outcome probabilities of a tomographically complete measurement [15]. These methods are facing a number of conceptual and practical challenges, in particular when it comes to applying them to real experimental data [16–18], where the assumption of independent and identically distributed (i.i.d.) experimental runs that underlies all quantum state tomography methods may not be fulfilled. In particular, the measurement itself is prone to errors and noise.

Here we compare different methods for NQS tomography by applying them to real experimental data in the case of a

*lfischer@physi.uni-heidelberg.de

†martin.gaerttner@kip.uni-heidelberg.de

two-qubit system. Specifically, we use a parametric down-conversion source producing entangled photon pairs. In such a small system the fitted NQS typically overparameterizes the state, such that the finite representational power of the employed variational ansatz does not limit the quality of the reconstruction. Nevertheless, we find that the learning procedure can be cumbersome. We benchmark the employed NQS tomography methods with maximum likelihood estimation (MLE) using the full density matrix and with direct measurements of Bell correlations. We evaluate the performance of these methods using measured as well as synthetic data, and for different choices of tomographically complete measurements. Our study will provide guidance when taking the next steps towards applying NQS tomography methods in regimes beyond the reach of traditional methods.

II. NQS TOMOGRAPHY METHODS

In this section we briefly review the quantum state tomography methods that we employed. More details can be found in the original publications [13–15].

The quantum state of a system is fully specified by its density matrix ρ , which is a Hermitian positive semidefinite matrix with unit trace. Without taking into account the trace constraint, this matrix is specified by d^2 real parameters, where d is the dimension of the Hilbert space. To infer all these parameters uniquely, a tomographically complete measurement, i.e., a measurement with at least d^2 outcomes needs to be carried out. For a general measurement, each of the outcomes corresponds to a positive operator $M_{\mathbf{a}}$ such that the probability of outcome \mathbf{a} is given by

$$P(\mathbf{a}) = \text{tr}[\rho M_{\mathbf{a}}]. \quad (1)$$

The set of operators $\{M_{\mathbf{a}}\}_{\mathbf{a}}$ constitutes a so-called positive operator valued measure (POVM) and fulfills $\sum_{\mathbf{a}} M_{\mathbf{a}} = \mathbb{1}$ [corresponding to the overall normalization $\sum_{\mathbf{a}} P(\mathbf{a}) = 1$]. The density matrix can be reconstructed uniquely from the probabilities $P(\mathbf{a})$ if the number of constraints imposed by the measurements matches (or exceeds) the number of parameters of the density matrix. This approach to reconstruct the state ρ by inverting the linear system in Eq. (1) is typically referred to as linear reconstruction [15]. In traditional quantum state tomography, all outcome probabilities are measured and the full density matrix is reconstructed. Linear reconstruction has the drawback that for finite experimental statistics the reconstructed ρ will in general not be positive semidefinite, i.e., it will not be a physical density matrix, which can lead to unphysical results such as negative state probabilities or fidelities exceeding unity. To avoid this, one can parametrize the density matrix in a way that ensures positivity and determine the physical state that most closely resembles the measured distribution by maximum likelihood estimation (MLE) [19]. The scaling issues of this method become obvious when we consider, for example, the case of N spin $1/2$ particles, where the number of required measurements is $d^2 = 4^N$, thus scaling exponentially in the system size. In efficient quantum state tomography one tries to overcome this problem by finding a variational ansatz for parametrizing either ρ or alternatively the probability distribution P with a number of variational parameters that scales polynomially in the system size. These

variational parameters are adapted to best match the observed measurement outcomes. Specifically, NQS tomography uses variational approaches that are inspired by generative models used in machine learning.

Before describing the NQS variational approach in detail, we need to specify what type of tomographically complete measurements we use. We are concerned with qubit systems, i.e., arrays of N two-level systems. In this case it is natural to assume that each qubit is read out individually and thus the POVM elements are product operators $M_{\mathbf{a}} = M_{a_1}^{(1)} \otimes \dots \otimes M_{a_N}^{(1)}$. If the single-qubit POVMs $\{M_{a_i}^{(1)}\}_{a_i}$ are tomographically complete, then this is also the case for $\{M_{\mathbf{a}}\}_{\mathbf{a}}$. We implemented two different tomographically complete measurements corresponding to the following single-qubit POVMs. First, the Pauli POVM, or Pauli-6 POVM, consists of the projectors

$$\{M_a^{(1)}\}_a = \bigcup_{\alpha=x,y,z} \left\{ \frac{1}{3} |\uparrow_{\alpha}\rangle\langle\uparrow_{\alpha}|, \frac{1}{3} |\downarrow_{\alpha}\rangle\langle\downarrow_{\alpha}| \right\}, \quad (2)$$

where $|\uparrow_{\alpha}\rangle$ and $|\downarrow_{\alpha}\rangle$ are the eigenstates of the Pauli operator σ_{α} with eigenvalue ± 1 . By summarizing the -1 outcomes into a single POVM element one obtains

$$M_0^{(1)} = \frac{1}{3} |\uparrow_x\rangle\langle\uparrow_x| \quad (3)$$

$$M_1^{(1)} = \frac{1}{3} |\uparrow_y\rangle\langle\uparrow_y| \quad (4)$$

$$M_2^{(1)} = \frac{1}{3} |\uparrow_z\rangle\langle\uparrow_z| \quad (5)$$

$$M_3^{(1)} = \mathbb{1} - M_0 - M_1 - M_2, \quad (6)$$

which we call the Pauli-4 POVM. The second measurement setting we consider is the so-called tetrahedral POVM, which consists of the subnormalized projectors onto states pointing into the corners of a regular tetrahedron on the Bloch sphere, i.e., $M_a^{(1)} = (\mathbb{1} + \mathbf{s}_a \cdot \boldsymbol{\sigma})/4$, with $\mathbf{s}_0 = (0, 0, 1)$, $\mathbf{s}_1 = (2\sqrt{2}/3, 0, -1/3)$, $\mathbf{s}_2 = (-\sqrt{2}/3, \sqrt{2}/3, -1/3)$, and $\mathbf{s}_3 = (-\sqrt{2}/3, -\sqrt{2}/3, -1/3)$ and $\boldsymbol{\sigma}$ the vector of Pauli operators. The Pauli-4 and tetrahedral POVM have the minimal number of four possible outcomes for each qubit required for being tomographically complete. In this case the overlap matrix $T_{\mathbf{a}\mathbf{a}'} = \text{Tr}[M_{\mathbf{a}} M_{\mathbf{a}'}]$ is invertible and density matrix can be reconstructed as

$$\rho = \sum_{\{\mathbf{a}\}} P(\mathbf{a}) \mathcal{Q}_{\mathbf{a}}, \quad (7)$$

where the operators $\mathcal{Q}_{\mathbf{a}}$ are given by $\mathcal{Q}_{\mathbf{a}} = \sum_{\{\mathbf{a}'\}} T_{\mathbf{a}\mathbf{a}'}^{-1} M_{\mathbf{a}'}$. As the overlap matrix factorizes for product POVMs, the inversion can be done on the single-spin level. Analogously to ρ , any Hermitian operator, representing an observable, can be decomposed into the complete set of operators $\{M_{\mathbf{a}}\}_{\mathbf{a}}$, which allows us to directly extract its expectation value from the POVM probability distribution $P(\mathbf{a})$.

The probability distribution $P(\mathbf{a})$ resulting from a tomographically complete measurement, can be approximated using generative modeling approaches from machine learning [15], which are known to be universal function approximators [20]. Once such a model has been trained to represent the measured distribution $P(\mathbf{a})$, expectation values of observables can be calculated efficiently by drawing samples from the model distribution. Here we focus on a restricted Boltzmann

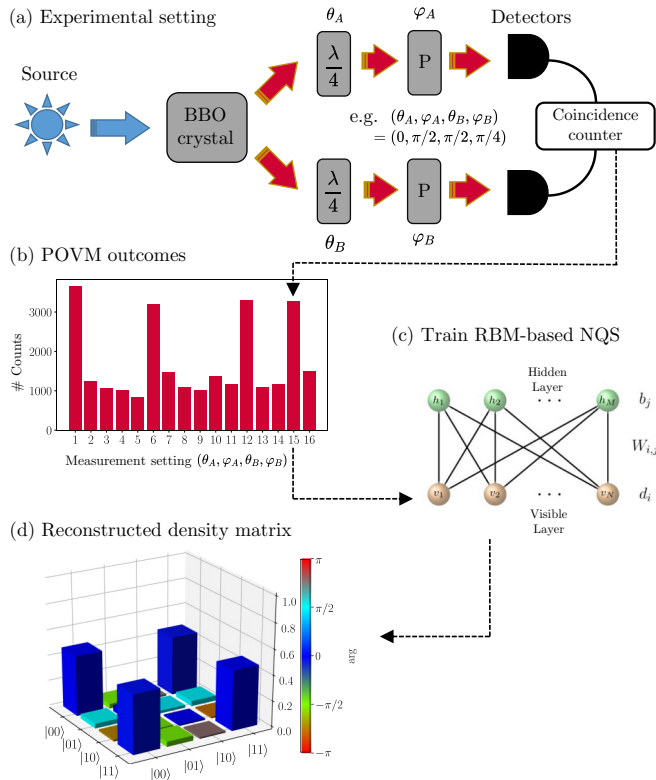


FIG. 1. Neural-network quantum state tomography scheme: (a) Entangled photonic Bell pairs from an SPDC process are measured by coincidence count rates. Arbitrary single-qubit projectors can be selected by adjusting the angles $\theta_A, \varphi_A, \theta_B, \varphi_B$ in the $\frac{\lambda}{4}$ -plates and polarizers P . (b) This is repeated for all elements of the POVM $\{M_{\mathbf{a}}\}_{\mathbf{a}}$ yielding a probability distribution $P(\mathbf{a})$ describing the state. (c) This forms the training data for different neural-network quantum state architectures based on restricted Boltzmann machines (figure adapted from Ref. [24]). (d) After training, the network-encoded states are used to reconstruct the density matrix or to predict observables.

machine (RBM) ansatz, which is a rather veteran model in machine learning but has proven large representational power making it suitable for encoding large classes of quantum many-body states [21]. The RBM ansatz function is

$$Q(\mathbf{a}; \mathcal{W}) = \frac{1}{Z(\mathcal{W})} \sum_{\mathbf{h}} \exp[-E(\mathbf{a}, \mathbf{h}; \mathcal{W})], \quad (8)$$

which is the marginal over all hidden states $\mathbf{h} = (h_1, \dots, h_M)$, $h_j \in \{0, 1\}$, of the joint Boltzmann distribution $\exp[-E(\mathbf{a}, \mathbf{h}; \mathcal{W})]/Z(\mathcal{W})$. The network energy $E(\mathbf{a}, \mathbf{h}; \mathcal{W}) = -\sum_{i,j} a_i W_{i,j} h_j - \sum_i a_i d_i - \sum_j h_j b_j$ depends on the weights $W_{i,j}$ and biases b_j and d_i [see Fig. 1(c)], which are summarized as the set of variational parameters \mathcal{W} . $Z(\mathcal{W}) = \sum_{\{\mathbf{a}, \mathbf{h}\}} \exp[-E(\mathbf{a}, \mathbf{h}; \mathcal{W})]$ ensures normalization. Here we consider specifically the tetrahedral and Pauli-4 POVMs, which have the minimal amount of four possible measurement outcomes, so the multinomial visible units a_i can take four different values. We use a one-hot encoding of their states [15]. The model distribution $Q(\mathbf{a}; \mathcal{W})$ is fitted to the measured distribution of outcomes, which follows $P(\mathbf{a})$, by contrastive divergence learning [22]. In this approach, the Kullback-Leibler (KL) divergence from the model to the

target (data) distribution,

$$D_{\text{KL}}(P\|Q) = \sum_{\{\mathbf{a}\}} P(\mathbf{a}) \ln \left[\frac{P(\mathbf{a})}{Q(\mathbf{a}; \mathcal{W})} \right], \quad (9)$$

is considered. This divergence can be interpreted as a measure for the distance between the two distributions and hence the model distribution can be optimized to approximate the target distribution by minimizing $D_{\text{KL}}(P\|Q)$. The optimization is done by updating the model parameters via gradient descent, following the negative gradient of the KL divergence.

For training we adopted the numerical routines provided along with Ref. [15]. We refer to this method of direct fitting of the POVM probability distribution as the POVM ansatz in the following.

Positivity of the density matrix can be enforced by using the latent space purification ansatz [14]. This ansatz exploits the fact that any density matrix can be reexpressed as the partial trace over a pure system with twice as many degrees of freedom. Mathematically

$$\rho_{\sigma\sigma'} = \sum_{\alpha} \Psi_{\sigma\alpha}^* \Psi_{\sigma'\alpha}, \quad (10)$$

where

$$|\Psi\rangle = \sum_{\sigma\alpha} \Psi_{\sigma\alpha} |\sigma\rangle \otimes |\alpha\rangle \quad (11)$$

is the pure state of system qubits $\sigma = (\sigma_1 \dots \sigma_N)$ and ancillary qubits $\alpha = (\alpha_1 \dots \alpha_N)$. The coefficients of this wave function in the given reference basis $(|\sigma\rangle \otimes |\alpha\rangle \in \{|\uparrow_z\rangle, |\downarrow_z\rangle\}^{\otimes 2N})$ can again be parameterized by an RBM-inspired variational ansatz. But note that the wave function coefficients are complex, which requires the RBM ansatz to have complex parameters. Alternatively, one can use one RBM for the encoding of the phase and one for encoding the modulus of the wave function coefficients. We used the latter approach, following Ref. [13]. Training of this ansatz is based on measurements in different bases. In our example of the Pauli POVM, we measured the state in all nine possible choices of measuring each qubit in one of the three Pauli bases. The training algorithm minimizes the sum of the KL divergences between the measured probabilities and the corresponding model prediction over all measured bases. We will refer to this NQS tomography method as the purification ansatz in the following.

If system and ancillary system are in a product state, the density matrix defined as the trace over the ancillary system will be pure. Thus, for restricting the variational manifold to pure states, one can simply remove all interactions between system and ancillary qubits in the variational ansatz. For learning the parameters of this pure-state ansatz, again, measurements in different bases are needed [13], for which we again use the Pauli POVM data.

Further assuming that the wave function $|\psi\rangle = \sum_{\sigma} \psi_{\sigma} |\sigma\rangle$ of the state has real non-negative coefficients ψ_{σ} , one can simply parametrize these coefficients by a real RBM [16]. In this case a measurement in the computational basis $\{|\sigma\rangle\}_{\sigma} = \{|\uparrow_z\rangle, |\downarrow_z\rangle\}^{\otimes N}$ is sufficient for training. We trained the purification ansatz and the pure state and positive-real wave function ansatz using the open source library QuCumber [23].

III. TWO-PHOTON EXPERIMENT

Our experimental test bed is a commercial photon source [25] that produces polarization-entangled photon pairs through type-1 SPDC in a nonlinear BBO crystal. A pump laser of 404 nm impinges with a linear polarization in a 45° angle to the optical axes of the crystal, creating superpositions close to the maximally entangled Bell state $|\psi\rangle = (|\uparrow_z\uparrow_z\rangle + |\downarrow_z\downarrow_z\rangle)/\sqrt{2}$, where $|\uparrow_z\rangle$ and $|\downarrow_z\rangle$ are represented by vertical and horizontal polarization states of the photons. Pre- and postcompensation crystals ensure the spatial and temporal indistinguishability of the nonlinear processes producing $|\uparrow_z\uparrow_z\rangle$ and $|\downarrow_z\downarrow_z\rangle$.

The single-photon polarization states are subsequently analyzed by a series of a quarter-wave plate (QWP), a polarizer and a single-photon detector, respectively, as illustrated in Fig. 1(a). By setting the angles of the QWP and the polarizer, arbitrary outcomes a_i can be selected for each photon. The expectation values of the projection operators $M_{\mathbf{a}} = M_{a_1} \otimes M_{a_2}$ are obtained by measuring the rate of coincidental counts at both detectors over a fixed period of time. This is repeated for each of the 36 (16) elements in $\{M_{\mathbf{a}}\}_{\mathbf{a}}$ of the Pauli-6 (tetrahedral) POVM. Assuming a constant production rate of photon pairs, the numbers of coincidence counts for each polarizer setting normalized by the total number of coincidence counts provides the required probability distribution of outcomes [Fig. 1(b)] from which the density matrix [Fig. 1(d)] can be reconstructed. Measuring each bin of the distribution separately in this fashion is obviously not scalable. In an equivalent scalable scenario, one would beam split each photon into different paths and select one of the single-photon POVM elements in each path. In this setup, detecting coincidence events and recording in which pair of the detectors the photons were observed, is a way of sampling from the full distribution of POVM outcomes in a single measurement setting.

We benchmark the reconstructed states on their predictions of the famous CHSH Bell parameter S [26]. This observable is measured in a Bell test to demonstrate correlations which cannot be reproduced by classical local hidden variable theories [27]. In such an experiment two particles (photons) are distributed to two parties who independently perform one of two possible measurements on their respective particle. The two parties evaluate the correlation coefficients $E = (N_{\uparrow\uparrow} + N_{\downarrow\downarrow} - N_{\uparrow\downarrow} - N_{\downarrow\uparrow})/N_{\text{tot}}$ for different measurement settings parameterized by an angle θ . Genuine quantum correlations are certified if the Bell parameter

$$S(\theta) = E(0, \theta) + E(0, -\theta) + E(2\theta, \theta) - E(2\theta, -\theta) \quad (12)$$

exceeds the classical bound of $|S| \leq 2$. We compare the results obtained for this observable when predicting it from the states obtained with the different tomography schemes to direct measurements of S as a function of θ .

IV. NQS TOMOGRAPHY BENCHMARK

Our experimental results are summarized in Fig. 2. Figures 2(a) and 2(b) show the results of the Pauli POVM measurements and tetrahedral POVM measurements, respectively. In both panels of Fig. 2, the results of the direct measurement of the Bell parameter for various different val-

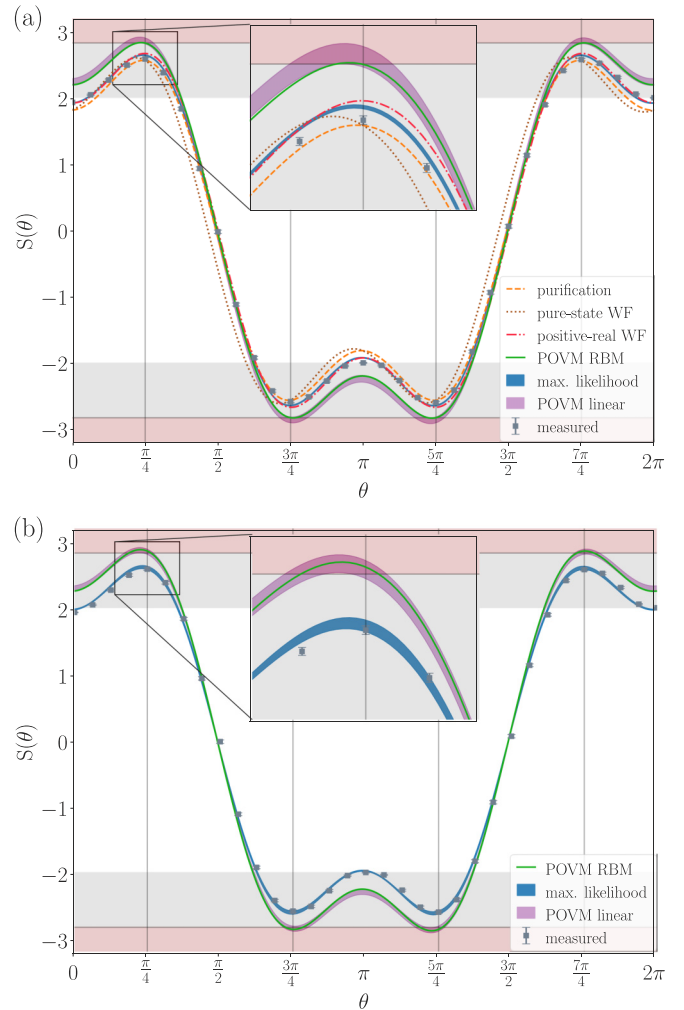


FIG. 2. Comparison of the predicted Bell parameter from the different NQS tomography schemes to direct measurements and maximum likelihood estimation. (a) Training with measurements in the Pauli product bases applicable to all considered NQS tomography schemes (b) Training with measurements of the tetrahedral POVM applicable to the POVM ansatz only. Gray regions indicate results forbidden by local realism, red regions denote results forbidden for physical quantum states. Error regions for linear reconstruction and MLE are the standard deviations after repeating the evaluation with data resampled from the measured distribution.

ues of θ are shown as blue points. We counted photons for $t_{\text{count}} = 3$ s for each measurements setting resulting in 25000 counts on average per point. For the tomographic measurements, we independently measured the probability distributions over the outcomes of the Pauli-6 (and thus also Pauli-4) POVM and the tetrahedral POVM using a total of 60000 and 27000 coincidence counts, respectively. As a first check we used these tomographically complete measurements to perform a classical MLE of the density matrix. From the resulting density matrix the Bell parameter can again be calculated. The result, shown by the blue line, agrees well with the Bell measurements (except for small deviations close to the extremal points, see inset of Fig. 2(a), which may be due to known shortcomings of MLE [28,29]). This confirms that the experimental conditions were largely stable between different

data sets. We also show the result of a direct linear reconstruction of ρ from the measured normalized count rates according to Eq. (7), including the statistical error resulting from the finite number of measurements (purple shadings in Fig. 2). The Bell parameter obtained for the linearly reconstructed ρ deviates from the directly measured data and even exceeds the value of $S = 2\sqrt{2}$ maximally allowed quantum mechanically, which indicates that the reconstruction yields an unphysical density matrix. Indeed one obtains negative eigenvalues for the reconstructed ρ $\{0.985, 0.078, -0.007, -0.055\}$ for the Pauli-4 POVM and $\{0.972, 0.11, -0.022, -0.056\}$ for the tetrahedral POVM).

Turning to NQS tomography, starting with the POVM ansatz, we trained a multinomial RBM with three hidden neurons to approximate both the measured Pauli-4 POVM distribution [Fig. 2(a)] and the tetrahedral POVM distribution [Fig. 2(b)]. In both cases the distribution was learned well by the network. The KL divergence decreased to $\sim 10^{-4}$. This is expected as the number of parameters in the network exceeds the number of bins in the probability distribution, thus overparameterizing the state. The predicted Bell parameter (solid green lines in Fig. 2) agrees with the result of the linear reconstruction and clearly deviates from the directly measured values. This shows that the linear reconstruction method on which the POVM ansatz [15] is based can give unreliable predictions for observables under imperfect measurements and noise, due to the absence of the positivity constraint on the density matrix.

In the case of the Pauli measurement we can also use the obtained data for training the purification ansatz, which does include the positivity constraint. Here we used a network architecture with two neurons for the latent and three neurons in the hidden layer [14]. This yields a prediction of comparable quality to the MLE, see dashed yellow line in Fig. 2(a). This confirms that the constraint of positivity can greatly improve the capabilities of tomography methods for estimating physical observables.

We now restrict the ansatz to encode only pure states by omitting couplings to the ancillary spins (pure-state ansatz). Although the MLE density matrix has a purity of only 0.91 for the Pauli and 0.87 for the tetrahedral POVM measurement this gives a reasonable estimate of the Bell parameter, as indicated by the dotted brown line in Fig. 2(a). However, we observe a shift in θ of the predicted curve with respect to the directly measured Bell correlations.

If we further simplify the ansatz by assuming real non-negative wave function coefficients, the agreement between prediction (dot-dashed red curve) and data is surprisingly good. This has two reasons: First, the prepared state has small complex phases of the coherences between the components of the Bell state. Second, this reconstruction technique only uses the information of the Pauli z -basis measurement, which happens to yield better predictions for the Bell parameter than methods using the full data set. Thus, the good agreement with the directly measured Bell parameter is rather a coincidence and may not hold for other observables as we further discuss below. We also observed that the fidelity to a perfect Bell state $|\psi\rangle = (|\uparrow_z\uparrow_z\rangle + |\downarrow_z\downarrow_z\rangle)/\sqrt{2}$ is highest for the state reconstructed assuming a pure state with positive-real coefficients. Given that the MLE reconstruction can be regarded as the

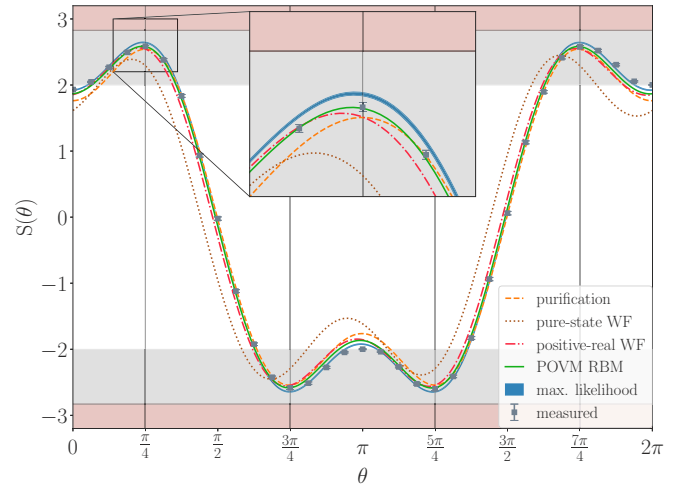


FIG. 3. Bell parameter estimations when applying the NQS tomography schemes to synthetic data in the Pauli bases sampled from the maximum likelihood density matrix from Fig. 2(a).

most reliable estimation of the density matrix, we conclude that the pure state assumption leads to a crude overestimation of the Bell-state fidelity. This shows that making unjustified assumptions about the state can lead to false conclusions about the quality of the prepared state.

We further test our hypothesis that the failure of the POVM ansatz, i.e., the linear reconstruction method, is mainly caused by imperfect measurements. For this, we generated a synthetic data set by exact sampling from the outcome distribution predicted by the reconstructed MLE density matrix. Figure 3 shows the results of this procedure for the Pauli measurement. When learning on this synthetic data set all methods perform well confirming that the failure of the linear reconstruction is indeed due to noisy measurements, i.e., imprecise settings of the chosen detection angles. The shift in θ , observed for the pure-state ansatz with complex wave function coefficients persists. This is understood by observing that during training, the pure-state ansatz faces an impossible task: The MLE density matrix exhibits considerable mixing, so the pure state will not be able to adapt to the sampled training data distribution arbitrarily well. The network instead tries to find a compromise that distributes the inevitable error over all possible bases. The Bell parameter is only sensitive to bases involving the z or x direction, which for the complex wave function training show a considerable error in an attempt to match the data in other bases that are irrelevant to the Bell parameter. In contrast, in the case of the positive-real state the training is only based on z -basis data, which yields a much better agreement with the independently measured Bell parameter.

Having benchmarked the quality of the different NQS tomography methods by their ability to predict a specific observable, namely the Bell parameter $S(\theta)$, we now analyze their ability to approximate the measured data as a more generic performance measure. We study how well they approximate the data in terms of the loss function that is minimized during the training of the purification ansatz (and of the pure-state and positive-real wave function ansatz), i.e., the average KL divergence between measured and




FIG. 4. Average KL divergence in the Pauli bases between measurements and reconstructed states for all employed tomography schemes. The complexity of the methods increases left to right, expressing the scaling of computational resources consumed (see text for details).

predicted outcome distributions of the Pauli-6 measurement, see Fig. 4.

The quality of the pure-state and positive-real wave function ansatz is clearly worse in this measure than estimations of the general density matrix. This is expected as the pure-state assumption is not justified. In the case of the positive-real wave function approach only part of the data, namely the measurements in the z basis are used for training resulting in a poor KL divergence to the full data set. This strengthens our conclusion that the good results of the positive-real wave function approach in reproducing the measured Bell parameter are restricted to this specific observable. The purification ansatz yields results that are not quite as good as the linear reconstruction Pauli-4 RBM and the MLE, cf. Fig 4. This shows that the purification network does not achieve optimal training of its cost function. This could either stem from difficulties of the numerical optimization or from insufficient representational power of the network. The POVM RBM learns the measured Pauli-4 distribution to high accuracy as the learning is not hampered by the positivity constraint. However, it does not resolve some of the details of the full Pauli-6 distribution as some Pauli-6 outcomes are summarized into one Pauli-4 outcome [Eq. (6)] causing a residual KL divergence to the full data set. We emphasize that the good performance of the POVM ansatz in this measure does not mean that it performs well in predicting physical observables, as for this task the positivity constraint on the density matrix can be crucial as discussed above.

We sorted the different approaches according to their complexity, referring to the computational resources required for the optimization, or training procedure. The positive-real wave function and POVM approaches involve only models with real-valued parameters and require no unitary rotations into different bases, which makes optimization easier. The pure-state and the purification ansatz were more challenging to train due to the unconventional network architecture required to encode the complex wave function coefficients. This is in contrast to the POVM ansatz where the POVM probability distribution can be learned using a standard RBM architecture. This is the price one has to pay for including the positivity constraint. The maximum likelihood estimate is computationally easier to obtain in this regime of few qubits

but scales exponentially for larger systems, which is why we classified it as the most complex method.

V. CONCLUSIONS

In summary, we have found that including the positivity of the density matrix as a constraint of the variational ansatz in efficient quantum state tomography methods greatly improves the quality of the reconstructed quantum state in terms of predicting physical observables. This may be expected as the constraint results in a smaller set of variational states among which the optimization is carried out. But interestingly, the constraint is not as important for training on synthetically generated data. From this we conclude that when applying efficient NQS tomography methods to experimental data, imperfections in the measurements can lead to subtle issues, which require a careful choice of the state parametrization. On the other hand, the inclusion of the positivity constraint renders the learning procedure more cumbersome as the purification ansatz imposes a complicated, rather nonstandard network architecture. Making additional assumptions about the state, which are not *a priori* fulfilled, can strongly bias observables predicted from the reconstructed state.

Questions to be addressed in the future concern the scaling properties of neural network state tomography approaches. It has been shown that neural-network quantum states can approximate large classes of states efficiently [10,21,30–34] with favorable generalization properties [35,36]. Our work shows that for real experimental data, there may be a tradeoff between the complexity of the ansatz and the quality of the learned state representations. This is in particular true for the inclusion of the constraint of the positivity of the density matrix, which, in general, leads to an exponential hardness of the optimization problem [14] but also significantly improves the quality of the tomographic reconstruction. It will be crucial to understand whether these issues persist when going to larger system sizes, where the measured distributions are necessarily undersampled and cannot be approximated precisely. In a regime where the generative model only needs to capture the coarse features of the quantum state, intuition from machine learning suggests that it may actually be expected to perform better. Furthermore, it has recently been shown that, in a regime where the Hilbert space dimension is much larger than the number of available measurements, adding the positivity constraint can deteriorate the quality of the reconstructed state [37,38]. Whether such effects also appear in tomography procedures based on artificial neural networks used in this work will be investigated when scaling up system sizes. Regarding representational power and learning performance, the choice of the RBM architecture is certainly not without alternative and we will employ and benchmark other generative models such as autoregressive models, autoencoders, or generative adversarial networks in the future.

ACKNOWLEDGMENTS

This work is supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC 218/1 – 390900948 (the

Heidelberg STRUCTURES Excellence Cluster), by the DFG, Project ID 273811115 – SFB 1225 (ISOQUANT), and by the

European Union Horizon 2020 program under Grant Agreement No. 817482 (PASQuaS).

-
- [1] I. M. Georgescu, S. Ashhab, and F. Nori, *Rev. Mod. Phys.* **86**, 153 (2014).
- [2] G. Tóth, W. Wieczorek, D. Gross, R. Krischek, C. Schwemmer, and H. Weinfurter, *Phys. Rev. Lett.* **105**, 250403 (2010).
- [3] T. Moroder, P. Hyllus, G. Tóth, C. Schwemmer, A. Niggebaum, S. Gaile, O. Gühne, and H. Weinfurter, *New J. Phys.* **14**, 105001 (2012).
- [4] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, *Phys. Rev. Lett.* **105**, 150401 (2010).
- [5] C. A. Riofrío, D. Gross, S. T. Flammia, T. Monz, D. Nigg, R. Blatt, and J. Eisert, *Nat. Commun.* **8**, 15305 (2017).
- [6] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, *Nat. Commun.* **1**, 149 (2010).
- [7] T. Baumgratz, D. Gross, M. Cramer, and M. B. Plenio, *Phys. Rev. Lett.* **111**, 020401 (2013).
- [8] B. P. Lanyon, C. Maier, M. Holzäpfel, T. Baumgratz, C. Hempel, P. Jurcevic, I. Dhand, A. S. Buyskikh, A. J. Daley, M. Cramer, M. B. Plenio, R. Blatt, and C. F. Roos, *Nat. Phys.* **13**, 1158 (2017).
- [9] G. Torlai and R. G. Melko, *Annu. Rev. Condens. Matter Phys.* **11**, 325 (2020).
- [10] G. Carleo and M. Troyer, *Science* **355**, 602 (2017).
- [11] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [12] J. Carrasquilla, *Adv. Phys.: X* **5**, 1797528 (2020).
- [13] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, *Nat. Phys.* **14**, 447 (2018).
- [14] G. Torlai and R. G. Melko, *Phys. Rev. Lett.* **120**, 240503 (2018).
- [15] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, *Nat. Mach. Intell.* **1**, 155 (2019).
- [16] G. Torlai, B. Timar, E. P. L. van Nieuwenburg, H. Levine, A. Omran, A. Keesling, H. Bernien, M. Greiner, V. Vuletić, M. D. Lukin, R. G. Melko, and M. Endres, *Phys. Rev. Lett.* **123**, 230504 (2019).
- [17] A. M. Palmieri, E. Kovlakov, F. Bianchi, D. Yudin, S. Straupe, J. D. Biamonte, and S. Kulik, *Npj Quantum Inf.* **6**, 20 (2020).
- [18] P. Cha, P. Ginsparg, F. Wu, J. Carrasquilla, P. L. McMahon, and E.-A. Kim, [arXiv:2006.12469](https://arxiv.org/abs/2006.12469).
- [19] M. Paris and J. Reháček, *Quantum State Estimation*, 1st ed. (Springer, Berlin, 2004).
- [20] N. Le Roux and Y. Bengio, *Neural Comput.* **20**, 1631 (2008).
- [21] R. G. Melko, G. Carleo, J. Carrasquilla, and J. I. Cirac, *Nat. Phys.* **15**, 887 (2019).
- [22] G. E. Hinton, in *Neural Networks: Tricks of the Trade: Second Edition*, edited by G. Montavon, G. B. Orr, and K.-R. Müller (Springer, Berlin, 2012), pp. 599–619.
- [23] M. J. S. Beach, I. D. Vlugt, A. Golubeva, P. Huembeli, B. Kulchytskyy, X. Luo, R. G. Melko, E. Merali, and G. Torlai, *SciPost Phys.* **7**, 9 (2019).
- [24] S. Czischek, *Neural-Network Simulation of Strongly Correlated Quantum Systems* (Springer Nature, Berlin, 2020).
- [25] qutools GmbH, *quED Entanglement Demonstrator Manual*, 1st ed., München, 2017, <https://www.qutools.com/>.
- [26] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, *Phys. Rev. Lett.* **23**, 880 (1969).
- [27] J. S. Bell, *Speakable and Unspeakable in Quantum Mechanics: Collected Papers on Quantum Philosophy*, 2nd ed. (Cambridge University Press, Cambridge, 2004).
- [28] C. Schwemmer, L. Knips, D. Richart, H. Weinfurter, T. Moroder, M. Kleinmann, and O. Gühne, *Phys. Rev. Lett.* **114**, 080403 (2015).
- [29] C. Ferrie and R. Blume-Kohout, [arXiv:1808.01072](https://arxiv.org/abs/1808.01072).
- [30] D.-L. Deng, X. Li, and S. Das Sarma, *Phys. Rev. X* **7**, 021021 (2017).
- [31] X. Gao and L.-M. Duan, *Nat. Commun.* **8**, 662 (2017).
- [32] Z. Cai and J. Liu, *Phys. Rev. B* **97**, 035116 (2018).
- [33] G. Carleo, Y. Nomura, and M. Imada, *Nat. Commun.* **9**, 5322 (2018).
- [34] Y. Levine, O. Sharir, N. Cohen, and A. Shashua, *Phys. Rev. Lett.* **122**, 065301 (2019).
- [35] D. Sehayek, A. Golubeva, M. S. Albergo, B. Kulchytskyy, G. Torlai, and R. G. Melko, *Phys. Rev. B* **100**, 195125 (2019).
- [36] T. Westerhout, N. Astrakhantsev, K. S. Tikhonov, M. I. Katsnelson, and A. A. Bagrov, *Nat. Commun.* **11**, 1593 (2020).
- [37] H.-Y. Huang, R. Kueng, and J. Preskill, *Nat. Phys.* **16**, 1050 (2020).
- [38] G. I. Struchalin, Y. A. Zagorovskii, E. V. Kovlakov, S. S. Straupe, and S. P. Kulik, [arXiv:2008.05234](https://arxiv.org/abs/2008.05234).