optica

Experimental quantum homodyne tomography via machine learning

E. S. Tiunov,[1,2,†] **V. V. Tiunova (Vyborova),**[1,†] **A. E. Ulanov,**[1] **A. I. Lvovsky,**[1,3,*] **and A. K. Fedorov**[1,2,4]

[1]*Russian Quantum Center, Skolkovo, Moscow 143025, Russia*
[2]*Moscow Institute of Physics and Technology, Dolgoprudny, Moscow Region 141700, Russia*
[3]*Department of Physics, University of Oxford, Oxford OX1 3PG, UK*
[4]*e-mail: akf@rqc.ru*
[*]*Corresponding author: alex.lvovsky@physics.ox.ac.uk*

Complete characterization of states and processes that occur within quantum devices is crucial for understanding and testing their potential to outperform classical technologies for communications and computing. However, solving this task with current state-of-the-art techniques becomes unwieldy for large and complex quantum systems. Here we realize and experimentally demonstrate a method for complete characterization of a quantum harmonic oscillator based on an artificial neural network known as the restricted Boltzmann machine. We apply the method to optical homodyne tomography and show it to allow full estimation of quantum states based on a smaller amount of experimental data compared to state-of-the-art methods. We link this advantage to reduced overfitting. Although our experiment is in the optical domain, our method provides a way of exploring quantum resources in a broad class of large-scale physical systems, such as superconducting circuits, atomic and molecular ensembles, and optomechanical systems. © 2020 Optical Society of America under the terms of the [OSA Open Access Publishing Agreement](#)

https://doi.org/10.1364/OPTICA.389482

1. INTRODUCTION

Exploiting the full potential of quantum technologies involves the challenge of "quantum volume": keeping a high degree of control over a complex many-body quantum system in spite of its growing size [1]. This important challenge concerns, in particular, methods for complete characterization of quantum states and processes. Quantum state tomography (QST), the reconstruction of quantum states from measurement statistics in multiple bases [2,3], is routinely performed in quantum physics experiments of various nature. Nevertheless, because the number of parameters describing a state of a quantum system grows exponentially with its size, tomography becomes increasingly demanding in application to large-scale quantum systems that are now engineered in experiments with ultracold atoms [4–7], ions [8–10], superconducting devices [11], and quantum light [12].

This problem manifests itself in two aspects. First, full quantum tomography of multi-dimensional quantum systems requires large portions of data, which are typically difficult to acquire experimentally. Second, even if such data are available, they are quite difficult to process with reasonable computational resources. Fortunately, it often happens that the physical setting being studied imposes certain *a priori* restrictions on the quantum states that can be prepared in it. As a result, the states can be described using a set of parameters that grows polynomially, rather than exponentially, with the size of the system. This observation gave rise to alternative approaches

such as permutationally invariant tomography [13], quantum compressed sensing [14], and tensor networks [15–17]. Each of these approaches makes particular assumptions about the physical restrictions imposed upon the state in question.

In the absence of knowledge about the physics of the system, one can use a universal approach based on generative artificial neural networks. Generally, neural networks are known to be capable of finding the best fit to arbitrarily complex data patterns with a limited number of parameters available [18]. In the context of quantum physics, this capability has been exploited in the context of neural networks known as the restricted Boltzmann machine (RBM). Such a neural net is proven to be a universal approximator for any discrete distribution [19]. RBMs are capable to encode the information about exponentially many terms of a quantum state in a polynomial number of units [20]. This feature makes RBMs attractive for a variety of quantum variational optimization problems [21], which require finding a quantum state that best satisfies a certain criterion. Examples of such problems, in addition to quantum tomography [22], include searching ground states of Hamiltonians in quantum chemistry tasks [23], investigating tensor network states [24] and topological states [25], and simulating open quantum many-body systems [26–30].

In the original theoretical proposal [22], RBM-based QST has been applied to simulated pure states of interacting many-qubit systems. A subsequent work [31] has generalized this approach

to mixed states and applied it to perform QST of a two-qubit system associated with a polarization-entangled photon pair. Very recently, the method was used in application to an experimental Rydberg-atom simulator with eight and nine atoms, using a pure-state, constant-phase approximation and measurements in a single basis [32]. Neural network techniques in the context of QST were also employed, albeit in a very different setting, to preprocess the data, thereby reducing the effect of state preparation and measurement errors [33].

However, all existing work on the subject has been applied to sets of natural qubits, such as fermion spins. This excludes a large class of "continuous-variable" physical systems whose Hamiltonian is identical to that of the harmonic oscillator. These include light, superconducting cavities, atomic and molecular ensembles, and optomechanical arrangements. Many of these systems are promising candidates for quantum information processing [34,35] and hence the challenge of quantum volume applies to them to the full extent. This necessitates the extension of neural-network QST methods to these systems.

Here we fill this gap by applying the RBM to homodyne tomography of optical states, in which measurements of electromagnetic field quadratures at various phases are performed to reconstruct the state of light in a given mode [2]. We verify our method on experimental data for the cases of optical Schrödinger's cat states and arbitrary Fock-state superpositions up to the two-photon level, where we obtain a high quality of quantum state reconstruction. We perform the universality test for our method via the reconstruction of randomly generated states. We also consider the application of our methods to other relevant quantum states, such as Gottesman–Kitaev–Preskill states [36] and squeezed-displaced vacuum. The approach generally outperforms standard maximum-likelihood-based methods [37], which, as we demonstrate, is deeply linked with reduced overfitting. To our knowledge, this is the first application of neural networks in a continuous-variable quantum setting.

2. NEURAL NETWORK TOMOGRAPHY

An RBM is a neural net containing two layers, visible and hidden, with all-to-all connections between the neurons in different layers and none inside each layer [hence the term "restricted"; see Fig. 1(a)]. The neurons can take on binary values {0,1}. Any set of neuron values, defined by binary vectors $\mathbf{v}$ and $\mathbf{h}$, is associated with the Boltzmann probability

$$p(\mathbf{v}, \mathbf{h}) = \frac{1}{Z} e^{-E(\mathbf{v}, \mathbf{h})}, \tag{1}$$

where Z is the partition function, and $E(\mathbf{v}, \mathbf{h})$ is the Ising-type energy functional

$$E(\mathbf{v}, \mathbf{h}) = -\mathbf{v}^T \hat{W} \mathbf{h} - \mathbf{a}^T \mathbf{v} - \mathbf{b}^T \mathbf{h}, \tag{2}$$

where $\hat{W}$, $\mathbf{a}$, $\mathbf{b}$ are the network parameters: weights and biases, respectively. The conventional RBM is trained to find the parameter set that maximizes the product of marginal distributions,

$$p(\mathbf{v}) = \sum_{\mathbf{h}} p(\mathbf{v}, \mathbf{h}), \tag{3}$$

over the training set $\{\mathbf{v}\}$, i.e., $\prod_{\{\mathbf{v}\}} p(\mathbf{v})$. The RBM trained in this way will produce similarly low energy values for test inputs that are
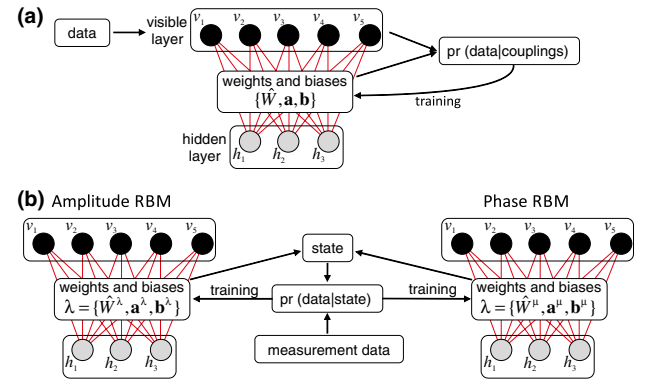

Fig. 1. Architecture of restricted Boltzmann machines for classical pattern recognition tasks (a) and quantum tomography (b).

similar to elements of the training set, which is useful for pattern recognition [38]. Furthermore, by sampling high-probability visible layer vectors, one can use the RBM as a generative neural network [39].

In the classical case, the data (such as the pattern to be recognized) are fed to the RBM through the visible layer. Doing so for quantum tomography would be unimaginable because there are infinitely many quantum states and even more possible measurement data sets. On the other hand, we can take advantage of our *a priori* knowledge of the connection between quantum states and the measurement probabilities associated with different bases.

These important differences dictate a different way that RBMs can be applied for quantum optimization problems. Here we utilize the RBMs to define an *Ansatz* expression for the quantum state $|\Psi\rangle$, which we wish to reconstruct. The neural network parameters are then used as the variational parameters of that Ansatz. We calculate the likelihood function (probability of having acquired the present experimental data set given $|\Psi\rangle$) using the knowledge of quantum mechanics, and optimize the parameters, and therefore $|\Psi\rangle$, to maximize that likelihood. The visible layer no longer plays the role of the container for the data, but only serves to *index* the basis of the Hilbert space: each possible configuration $\mathbf{v}$ of the visible layer is associated with one and only one basis element $|\mathbf{v}\rangle$.

The Carleo and Troyer Ansatz [20], which we utilize here, uses two RBMs of identical architectures [Fig. 1(b)], with the parameter sets $\lambda = \{\hat{W}^\lambda, \mathbf{a}^\lambda, \mathbf{b}^\lambda\}$ and $\mu = \{\hat{W}^\mu, \mathbf{a}^\mu, \mathbf{b}^\mu\}$ to express, respectively, the amplitudes and phases of the state's decomposition into this basis:

$$|\Psi\rangle = \sum_{\mathbf{v}} \sqrt{p_{\mathbf{v}}} e^{i\phi_{\mathbf{v}}/2} |\mathbf{v}\rangle, \tag{4}$$

where

$$p_{\mathbf{v}} = \frac{1}{Z_\lambda} \sum_{\mathbf{h}} e^{-E^\lambda(\mathbf{v}, \mathbf{h})}, \quad \phi_{\mathbf{v}} = \log \sum_{\mathbf{h}} e^{-E^\mu(\mathbf{v}, \mathbf{h})}, \tag{5}$$

and $E^{\lambda, \mu}$ are defined by Eq. (2) for the two corresponding RBMs. We note that the partition function Z is present only in the expression for the amplitudes, but not phases, because the phases have no normalization requirement. The logarithm is included in the phase for mathematical convenience.

In optical homodyne tomography, the basis traditionally used for state reconstruction is the Fock basis, bounded from above by some cut-off photon number N_{ph}. Because an RBM with the

visible layer of size m can represent a Hilbert space of dimension 2^m, the natural choice is to construct the reconstruction basis from photon number states $\{|0\rangle, \ldots, |N_{\mathrm{ph}} = 2^m - 1\rangle\}$. The basis is then encoded in the visible layer in a straightforward fashion, for example, for $m = 2$,

$$|0\rangle \rightarrow \begin{pmatrix} 0 \\ 0 \end{pmatrix} \quad |1\rangle \rightarrow \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad |2\rangle \rightarrow \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad |3\rangle \rightarrow \begin{pmatrix} 1 \\ 1 \end{pmatrix}.$$

The tomography experiment consists of measuring the continuous electromagnetic field quadrature samples X on multiple copies of the state $|\Psi\rangle$ at various phases θ. The log-likelihood functional is then as follows:

$$\Xi = \sum_j \log\langle\theta_j, X_j|\hat{\rho}|\theta_j, X_j\rangle, \tag{6}$$

where $\hat{\rho} = |\Psi\rangle\langle\Psi|$ is the density matrix, and j enumerates measurement outcomes. This is a differentiable function of the RBM parameters, defined through Eqs. (2), (4), and (5). These parameters can therefore be optimized using gradient descent to maximize the log-likelihood.

A general quantum tomography method must be able to work with not only pure states, but also with mixed ones. The method above is readily generalized to mixed states by means of purification: introducing an ancillary "environment" Hilbert space, whose dimension is equal to that of the Hilbert space of interest. The mixed state that needs to be reconstructed can then be written as a partial trace:

$$\hat{\rho} = \mathrm{Tr}_E\left(|\Psi_{SE}\rangle\langle\Psi_{SE}|\right), \tag{7}$$

where the pure state $|\Psi_{SE}\rangle$ is a vector of the tensor product Hilbert space comprising the system and the environment and can be reconstructed from the experimental data as described above (see Supplement 1 for details). We note that, although the dimension of the tensor product space is the square of the dimension of the system, the number of visible units needed to represent that space is only twice as large as that for the system alone.

We emphasize again the difference between the RBM approach to state reconstruction and the conventional quantum expectation-maximization (MaxLik) technique [37,40]. In both cases, we optimize the parameters of the state to maximize the likelihood functional (6). However, in the standard approach, all elements of the density matrix are being optimized, which corresponds to the number of parameters equal to the dimension of the Hilbert space squared. Within the RBM Ansatz, on the other hand, the number of parameters is on the scale of the product of the number of visible and hidden units, i.e., scales logarithmically with the Hilbert space dimension. As discussed previously, this is of great advantage when this dimension is large. Although reducing the number of parameters does restrict the set of states that can be expressed by the RBM Ansatz, we found it to be sufficient to adequately represent the states observed in homodyne tomography experiments.

We test our approach on two sets of experimental data. The first set corresponds to an optical analog of Schrödinger's cat, i.e., the superposition of two opposite-amplitude coherent states. The data have been taken from the experiment [41] and correspond to the cat state of amplitude $\alpha = 1.85$ squeezed by 3 dB along the quadrature axis. The second data set was
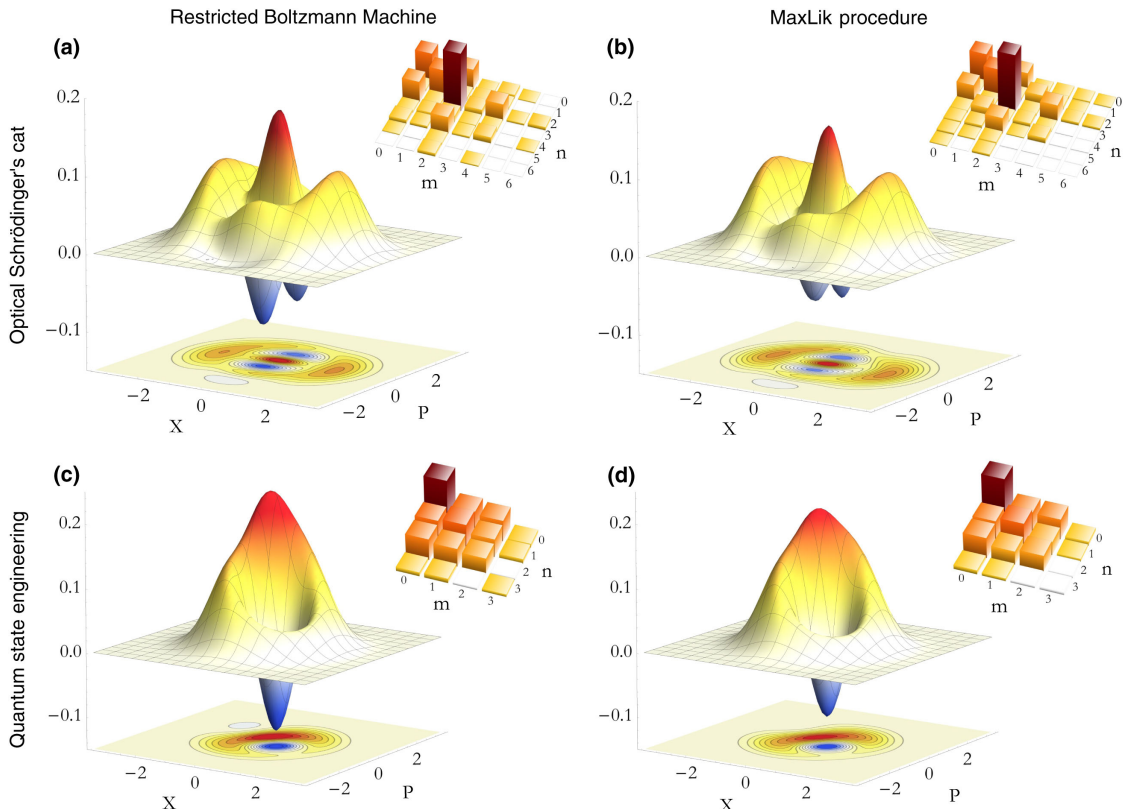

Fig. 2. Experimentally reconstructed Wigner functions and density matrices for optical Schrödinger's cats (a), (b) and engineered Fock superpositions up to the two-photon level (c), (d) using neural network quantum tomography (a), (c) and MaxLik (b), (d). The relative fidelity of the two reconstructed states is about 0.998 in both cases after efficiency correction.

obtained in an experiment on engineering arbitrary superpositions of Fock states $a_0|0\rangle + a_1|1\rangle + a_2|2\rangle$ with the amplitude ratio $a_0 : a_1 : a_2 \sim -0.76 : 0.49 : 0.42$ [42]. We compare our reconstruction results with the standard iterative MaxLik algorithm with efficiency correction. For both methods, we obtain Wigner functions and density matrices of the reconstructed states (Fig. 2).

For the reconstruction of the cat state, we used the cutoff photon number of $N_{ph} = 7$ (i.e., $m = 3$), which corresponds to the amplitude and phase RBMs containing $2m = 6$ visible units each. Additionally, each RBM contained eight hidden units. The reconstruction featured correction for 62% detection efficiency (see Supplement 1). For the Fock state superposition, each RBM had four visible units, four hidden units, $N_{ph} = 3$ ($m = 2$), and efficiency correction of 55%. As we see in Fig. 2, both methods resulted in similar reconstructed states, with the relative fidelity about 0.998 in both cases. In Supplement 1, we present the reconstruction from the same experimental data but without efficiency correction.

3. EFFECTS OF OVERFITTING

Our next goal is to compare the performance of the RBM approach to MaxLik. Using *bona fide* experimental data is suboptimal for this purpose because it is not known what "true" state they correspond to, and hence we cannot tell which method gives better reconstruction.

Therefore we generate a simulated quadrature data set corresponding to the Schrödinger's cat states $|\alpha\rangle - |-\alpha\rangle$ with $\alpha = 4$, reconstruct the state from this set and compare it to the original. The RBM reconstruction was performed without assuming the state to be pure, using an RBM with 10 visible units ($m = 5$) and three hidden units. The cut-off point was at 31 photons both for RBM and MaxLik. The motivation for choosing this relatively large Hilbert space is to explore the case in which the number parameters optimized by the RBM is much less than MaxLik.

Figure 3(a) shows the photon statistics of the state reconstructed using the two methods. Theoretically, we expect this state to show Poisson statistics for odd photon numbers, but zero probability for
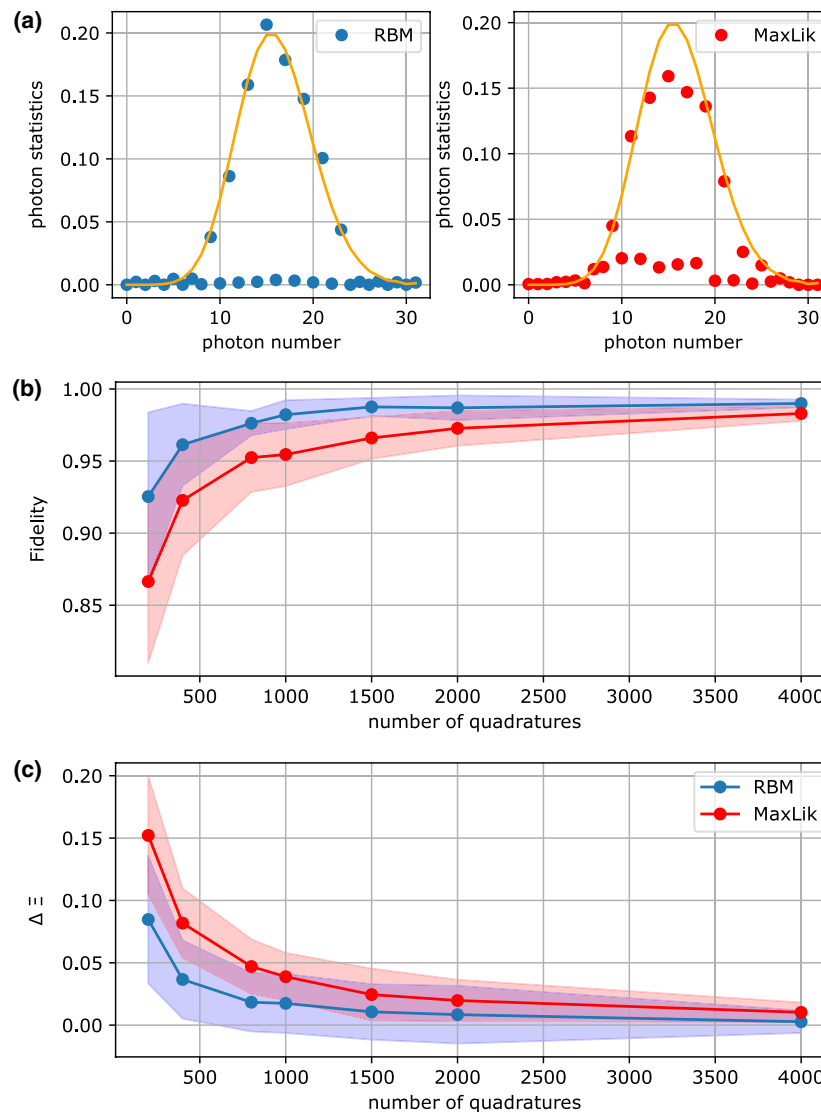

Fig. 3. Reconstruction of the cat state $|\alpha\rangle - |-\alpha\rangle$ with $\alpha = 4$ from bootstrapped data. (a) Photon statistics for the state reconstructed from 800 quadratures using MaxLik (right) and RBMs (left). (b) Reconstruction fidelity as a function of number of quadrature measurements. Each point is averaged over multiple datasets of the same size. Shaded regions show the standard deviation. (c) Cross-validation log-likelihood difference (see text). Higher values correspond to more significant overfitting.
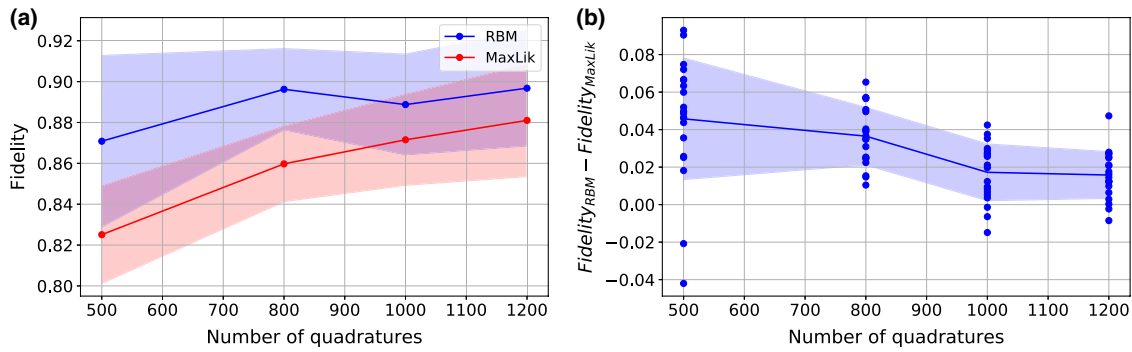
Fig. 4. Reconstruction of random states. (a) Reconstruction fidelity as a function of number of quadrature measurements. Each point is averaged over multiple states, with a single quadrature dataset generated for each state. Shaded regions show the standard deviation. (b) Differences between RBM and MaxLik fidelities for individual random states.

even photon numbers. We see that the state reconstructed using RBMs largely follows this rule, whereas the MaxLik reconstructed state has significant nonzero statistics for even photon numbers. In Fig. 3(b), we plot the fidelity of the reconstructed state with the original one as a function of the data set size and observe that RBM performs significantly better. For example, the RBM reconstructs the state from 1000 quadrature samples with the same fidelity of 98.5% as does MaxLik from 5000 samples. This is of value because complex quantum state engineering experiments typically produce desired states at very low rates [41], so the usage of RBM can greatly reduce the data collection effort.

The improved performance of the RBM approach for a smaller amount of experimental data is likely associated with lower overfitting [20]. Indeed, the number of parameters in MaxLik is, as discussed, $32^2 - 1 = 1023$, whereas for RBM, it is $2 \times (10 \times 3 + 10 + 3) = 86$. In order to demonstrate that overfitting is indeed the cause of poorer performance of MaxLik, we implement the following cross-validation test. We generate multiple quadrature data sets of the same size and reconstruct the state from one of them. Then we calculate the log-likelihood (6) for the data from each set with respect to the reconstructed state. If overfitting plays a significant role in the reconstruction, the likelihood of the "native" data set (from which the state was reconstructed) is expected to be significantly higher than for other sets. We plot the mean difference of the log-likelihoods for the "native" and "non-native" data sets in Fig. 3(c) and observe this difference to be much higher for MaxLik than for RBM. This confirms our hypothesis.

To test the generality of our conclusions, we applied RBM reconstruction to three states of different nature: Gottesman–Kitaev–Preskill [36], squeezed-displaced vacuum, and random states (see Supplement 1). We observed the same results as for the cat state. This corroborates our hypothesis that the likely reason for RBM's superiority to MaxLik is that the former method is less prone to overfitting.

As a further test, we applied RBM and MaxLik reconstruction to a set of 20 random superpositions of Fock states from zero to 31 photons. To generate these superpositions, a Gaussian random number generator with zero mean and unit variance was used to generate the real and imaginary amplitudes of each Fock component, and the resulting states were subsequently normalized to unity. The performance of the reconstruction (with $N_{ph} = 31$) is shown in Fig. 4, demonstrating the advantage of RBM-based QST for a great majority of states.

4. DISCUSSION AND OUTLOOK

As with any tomography method, a technique for estimating the state reconstruction precision is needed. The standard approach to this task is bootstrapping—that is, generating multiple simulated quadrature sample sets from the reconstructed state and reconstructing a state from each of these sets. The variance of these "secondary" states with respect to the originally reconstructed one gives an estimate for the statistical uncertainty of the reconstruction. This approach, as well as more advanced error estimation methods [43,44], can be successfully applied to RBM tomography.

A related question is the number of hidden units in the RBM. On one hand, increasing this hyperparameter improves the reconstruction quality because of the better expressive capacity of the neural network. On the other hand, it increases the computation complexity. The compromise depends on the specific state being reconstructed. For example, the squeezed-displaced vacuum state, whose wave function is Gaussian and does not have multiple fine features, required less hidden units than the cat and Gottesman–Kitaev–Preskill states. For all states we tested, the reconstruction quality improvement saturated for the number of hidden units being similar to or less than the number of visible units.

Our results demonstrate that the neural network QST approach is a promising way of characterizing the states observed in optical experiments. We found this method to be capable of reliable state reconstruction and much less prone to overfitting compared to the standard MaxLik approach. However, the full capability of our method is expected to be unveiled for very large Hilbert spaces, to which traditional methods become inapplicable. Therefore the natural next step would be to implement a complex multimode entangled state and apply RBM for its reconstruction. Promising sources of such states are multimode parametric oscillators, which have seen rapid development in recent years [45,46].

As stated earlier, the complexity of the QST problem, as well as the number of required measurements, grows exponentially with the system size. The RBM Ansatz appears to circumvent this issue, as the number of RBM parameters is polynomial with respect to the system size. The price to pay is that it may not be possible to efficiently describe all states with this Ansatz. In the discrete-variable domain, there exists a known class of physically interesting quantum states that carry no efficient RBM description [47]. It is important to undertake a similar study for continuous-variable systems to understand the application range of this method—in particular, to which extent it can be used in the multimode case.

To proceed in the direction of large systems, we will also need to change the strategy of RBM training. Presently, our evaluation of the likelihood function relies on exhaustive summation of amplitudes for all elements of the Hilbert space basis (see Supplement 1). However, such a summation will be impossible in large Hilbert spaces. Instead, we will have to rely on approximate methods of RBM training such as contrastive divergence [48] or Gibbs sampling [49] to select the basis elements with largest amplitudes. Alternative neural network architectures should also be explored. In particular, it would be interesting to look for ways to utilize forward-propagating neural networks, rather than RBMs, for QST [50]. Such neural networks are more common in modern machine learning because their training is much more straightforward.

Our approach can be generalized to broader classes of physical problems. First, in addition to light, it is applicable to any physical system that can be mapped to a harmonic oscillator, such as atomic ensembles [51] and nanomechanics [52]. Second, we reiterate that the neural-network-based QST studied here belongs to a larger class of problems in which one looks for a quantum state that best satisfies a certain criterion. A particularly promising field of research, in our opinion, is complex phenomena in condensed matter systems, such as many-body localization, and describing exotic phase transitions. Approaches based on machine learning constitute a new and promising way of tackling them.

Disclosures. The authors declare no conflicts of interest.

See Supplement 1 for supporting content.

[†] These authors contributed equally to this work.

REFERENCES

1. L. S. Bishop, S. Bravyi, A. Cross, J. M. Gambetta, and J. Smolin, "Quantum volume," Technical report (IBM, 2017).
2. A. I. Lvovsky and M. G. Raymer, "Continuous-variable optical quantum-state tomography," Rev. Mod. Phys. **81**, 299 (2009).
3. U. Leonhardt, *Measuring the Quantum State of Light* (Cambridge University, 1997).
4. D. Barredo, S. de Leseleuc, V. Lienhard, T. Lahaye, and A. Browaeys, "An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays," Science **354**, 1021 (2016).
5. M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletić, M. Greiner, and M. D. Lukin, "Atom-by-atom assembly of defect-free one-dimensional cold atom arrays," Science **354**, 1024 (2016).
6. H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, "Probing many-body dynamics on a 51-atom quantum simulator," Nature **551**, 579 (2017).
7. D. Barredo, V. Lienhard, S. de Léséleuc, T. Lahaye, and A. Browaeys, "Synthetic three-dimensional atomic structures assembled atom by atom," Nature **561**, 79 (2018).
8. J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe, "Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator," Nature **551**, 601 (2017).
9. N. Friis, O. Marty, C. Maier, C. Hempel, M. Holzäpfel, P. Jurcevic, M. B. Plenio, M. Huber, C. Roos, R. Blatt, and B. Lanyon, "Observation of entangled states of a fully controlled 20-qubit system," Phys. Rev. X **8**, 021012 (2018).
10. C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, and P. Zoller, "Self-verifying variational quantum simulation of lattice models," Nature **569**, 355 (2019).
11. C. Neill, P. Roushan, K. Kechedzhi, S. Boixo, S. V. Isakov, V. Smelyanskiy, R. Barends, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. Fowler, B. Foxen, R. Graff, E. Jeffrey, J. Kelly, E. Lucero, A. Megrant, J. Mutus, M. Neeley, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, H. Neven, and J. M. Martinis, "A blueprint for demonstrating quantum supremacy with superconducting qubits," Science **360**, 195 (2018).
12. U. L. Andersen, J. S. Neergaard-Nielsen, P. van Loock, and A. Furusawa, "Hybrid discrete- and continuous-variable quantum information," Nat. Phys. **11**, 713 (2015).
13. G. Toth, W. Wieczorek, D. Gross, R. Krischek, C. Schwemmer, and H. Weinfurter, "Permutationally invariant quantum tomography," Phys. Rev. Lett. **105**, 250403 (2010).
14. 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).
15. M. Cramer, M. B. Plenio, S. T. Flammia, D. Gross, S. D. Bartlett, R. Somma, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, "Efficient quantum state tomography," Nat. Commun. **1**, 149 (2010).
16. B. P. Lanyon, C. Maier, M. Holzpfel, 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, "Efficient tomography of a quantum many-body system," Nat. Phys. **13**, 1158 (2017).
17. J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, "Reconstructing quantum states with generative models," Nat. Mach. Intell. **1**, 155 (2019).
18. G. Cybenko, "Approximations by superpositions of sigmoidal functions," Math. Control Signals Syst. **2**, 303 (1989).
19. N. Le Roux and Y. Bengio, "Representational power of restricted Boltzmann machines and deep belief networks," Neural Comput. **20**, 1631–1649 (2008).
20. G. Carleo and M. Troyer, "Solving the quantum many-body problem with artificial neural networks," Science **355**, 602 (2017).
21. R. Melko, G. Carleo, J. Carrasquilla, and J. I. Cirac, "Restricted Boltzmann machines in quantum physics," Nat. Phys. **15**, 887–892 (2019).
22. G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, "Neural-network quantum state tomography," Nat. Phys. **14**, 447 (2017).
23. R. Xia and S. Kais, "Quantum machine learning for electronic structure calculations," Nat. Commun. **9**, 4195 (2018).
24. J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, "Equivalence of restricted Boltzmann machines and tensor network states," Phys. Rev. B **97**, 085104 (2018).
25. S. Lu, X. Gao, and L.-M. Duan, "Efficient representation of topologically ordered states with restricted Boltzmann machines," Phys. Rev. B **99**, 155136 (2019).
26. M. Schuld, I. Sinayskiy, and F. Petruccione, "Viewpoint: neural networks take on open quantum systems," Physics **12**, 74 (2019).
27. N. Yoshioka and R. Hamazaki, "Constructing neural stationary states for open quantum many-body systems," Phys. Rev. B **99**, 214306 (2019).
28. A. Nagy and V. Savona, "Variational quantum Monte Carlo method with a neural-network ansatz for open quantum systems," Phys. Rev. Lett. **122**, 250501 (2019).
29. M. J. Hartmann and G. Carleo, "Neural-network approach to dissipative quantum many-body dynamics," Phys. Rev. Lett. **122**, 250502 (2019).
30. F. Vicentini, A. Biella, N. Regnault, and C. Ciuti, "Variational neural-network ansatz for steady states in open quantum systems," Phys. Rev. Lett. **122**, 250503 (2019).
31. G. Torlai and R. G. Melko, "Latent space purification via neural density operators," Phys. Rev. Lett. **120**, 240503 (2018).
32. 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, "Integrating neural networks with a quantum simulator for state reconstruction," Phys. Rev. Lett. **123**, 230504 (2019).
33. A. M. Palmieri, E. Kovlakov, F. Bianchi, D. Yudin, S. Straupe, J. Biamonte, and S. Kulik, "Experimental neural network enhanced quantum tomography," npj Quant. Inf. **6**, 20 (2020).

34. M. V. Larsen, X. Guo, C. R. Breum, J. S. Neergaard-Nielsen, and U. L. Andersen, "Deterministic generation of a two-dimensional cluster state," Science **366**, 369 (2019).

35. W. Asavanant, Y. Shiozawa, S. Yokoyama, B. Charoensombutamon, H. Emura, R. N. Alexander, S. Takeda, J. Yoshikawa, N. C. Menicucci, H. Yonezawa, and A. Furusawa, "Generation of time-domain-multiplexed two-dimensional cluster state," Science **366**, 373 (2019).

36. D. Gottesman, A. Kitaev, and J. Preskill, "Encoding a qubit in an oscillator," Phys. Rev. A **64**, 012310 (2001).

37. A. I. Lvovsky, "Iterative maximum-likelihood reconstruction in quantum homodyne tomography," J. Opt. B **6**, S556 (2004).

38. H. Larochelle and R. Bengio, "Classification using discriminative restricted Boltzmann machines," in *Proceedings of the 25th International Conference on Machine Learning* (ACM, 2008), pp. 536–543.

39. H. Hu, L. Gao, and Q. Ma, "Deep restricted Boltzmann networks," arXiv:1611.07917 (2016).

40. Z. Hradil, J. Rehacek, J. Fiurášek, and M. Ježek, "Maximum-likelihood methods in quantum mechanics," Lect. Notes Phys. **649**, 59 (2004).

41. D. V. Sychev, A. E. Ulanov, A. A. Pushkina, M. W. Richards, I. A. Fedorov, and A. I. Lvovsky, "Enlargement of optical Schrödinger's cat states," Nat. Photonics **11**, 379 (2017).

42. E. Bimbard, N. Jain, A. MacRae, and A. I. Lvovsky, "Quantum-optical state engineering up to the two-photon level," Nat. Photonics **4**, 243 (2010).

43. P. Faist and R. Renner, "Practical and reliable error bars in quantum tomography," Phys. Rev. Lett. **117**, 010404 (2016).

44. J. Wang, V. B. Scholz, and R. Renner, "Confidence polytopes in quantum state tomography," Phys. Rev. Lett. **122**, 190401 (2019).

45. J. Yoshikawa, S. Yokoyama, T. Kaji, C. Sornphiphatphong, Y. Shiozawa, K. Makino, and A. Furusawa, "Generation of one-million-mode continuous-variable cluster state by unlimited time-domain multiplexing," APL Photon. **1**, 060801 (2016).

46. Y. Cai, J. Roslund, G. Ferrini, F. Arzani, X. Xu, C. Fabre, and N. Treps, "Multimode entanglement in reconfigurable graph states using optical frequency combs," Nat. Commun. **8**, 15645 (2017).

47. X. Gao and L. M. Duan, "Efficient representation of quantum many-body states with deep neural networks," Nat. Commun. **8**, 662 (2017).

48. M. A. Carreira-Perpinan and G. E. Hinton, "On contrastive divergence learning," in *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics (AISTATS 2005)* (The Society for Artificial Intelligence and Statistics, 2005), p. 33.

49. S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," IEEE Trans. Pattern Anal. Mach. Intell. **PAMI-6**, 721–741 (1984).

50. O. Sharir, Y. Levine, N. Wies, G. Carleo, and A. Shashua, "Deep autoregressive models for the efficient variational simulation of many-body quantum systems," Phys. Rev. Lett. **124**, 020503 (2020).

51. K. Hammerer, A. S. Sørensen, and E. S. Polzik, "Quantum interface between light and atomic ensembles," Rev. Mod. Phys. **82**, 1041 (2010).

52. M. Aspelmeyer, T. J. Kippenberg, and F. Marquardt, *Cavity Optomechanics* (Springer, 2014).