



Quantum State Tomography with Conditional Generative Adversarial Networks ()

Downloaded from: <https://research.chalmers.se>, 2022-01-28 18:52 UTC

Citation for the original published paper (version of record):

Ahmed, S., Sánchez Muñoz, C., Nori, F. et al (2021)

Quantum State Tomography with Conditional Generative Adversarial Networks ()

Physical Review Letters, 127(14)

<http://dx.doi.org/10.1103/PhysRevLett.127.140502>

N.B. When citing this work, cite the original published paper.

Quantum State Tomography with Conditional Generative Adversarial Networks

Shahnawaz Ahmed^{1,*}, Carlos Sánchez Muñoz², Franco Nori^{3,4}, and Anton Frisk Kockum^{1,†}

¹*Department of Microtechnology and Nanoscience, Chalmers University of Technology, 412 96 Gothenburg, Sweden*

²*Departamento de Física Teórica de la Materia Condensada and Condensed Matter Physics Center (IFIMAC), Universidad Autónoma de Madrid, Madrid 28049, Spain*

³*Theoretical Quantum Physics Laboratory, RIKEN Cluster for Pioneering Research, Wako-shi, Saitama 351-0198, Japan*

⁴*Department of Physics, University of Michigan, Ann Arbor, Michigan 48109-1040, USA*



(Received 14 December 2020; revised 21 May 2021; accepted 10 June 2021; published 27 September 2021)

Quantum state tomography (QST) is a challenging task in intermediate-scale quantum devices. Here, we apply conditional generative adversarial networks (CGANs) to QST. In the CGAN framework, two dueling neural networks, a generator and a discriminator, learn multimodal models from data. We augment a CGAN with custom neural-network layers that enable conversion of output from any standard neural network into a physical density matrix. To reconstruct the density matrix, the generator and discriminator networks train each other on data using standard gradient-based methods. We demonstrate that our QST-CGAN reconstructs optical quantum states with high fidelity, using orders of magnitude fewer iterative steps, and less data, than both accelerated projected-gradient-based and iterative maximum-likelihood estimation. We also show that the QST-CGAN can reconstruct a quantum state in a single evaluation of the generator network if it has been pretrained on similar quantum states.

DOI: [10.1103/PhysRevLett.127.140502](https://doi.org/10.1103/PhysRevLett.127.140502)

Introduction.—The ability to manipulate and control small quantum systems opens up promising directions for research and technological applications: quantum information processing and computation [1–5], simulations of quantum chemistry [6–10], secure communication [11,12], and much more [6,13–22]. A prominent example is the recent demonstration of a 53-qubit quantum computer performing a computational task in a few hundred seconds that was anticipated to take much longer on a classical supercomputer [5]. Such speedup is possible partly due to the exponentially large state space that can be used for storage and manipulation of information in quantum systems [23–25]. However, this large size of the state space also brings challenges for the characterization and description of these systems.

The process of reconstructing a full description of a quantum state by measuring its properties is called quantum state tomography (QST) [26–28]. Tomography is fundamentally a data processing problem, trying to extract meaningful information from as few (noisy) measurements as possible [29–39]. There exist many cleverly crafted QST techniques apart from general maximum likelihood estimation (MLE) [40,41], e.g., diluted MLE [42], compressed

sensing [43], Bayesian tomography [44,45], projected-gradient descent [46], matrix-product-state and tensor-network tomography [31,47,48], and permutationally invariant tomography [49,50]. However, these techniques are often restricted to specific types of quantum states, lacking versatility [51].

Recently, machine-learning methods have been applied to QST, yielding promising results [52–64]. In particular, generative models [51,65–67], usually restricted Boltzmann machines (RBMs), have been used as *Ansätze* with few parameters to represent a quantum state and learn the probability distribution of outputs expected from that state [48,52,68]. There are also examples of deep neural networks being used for QST [57,69–73], enabling physicists to take advantage of the rapid progress in such machine-learning techniques.

One interesting recent development in machine learning is generative adversarial networks (GANs) [74,75]. Such networks have led to an explosion of new results that were previously thought futuristic: generation of photorealistic images [76–78], conversion of sketches to images [76], text generation in different styles [79,80], text-to-image generation [81], generating and defending against fake news [82,83], and even game design learned from observing video [84]. An improvement of standard GANs that led to many of these results is conditional generative adversarial learning [85], which enabled increased control of the output of generative models. Recently, such GANs have been applied to tomography of materials structure with synchrotron

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by Bibsam.

radiation [86,87] and computed tomography of soft tissue in medicine [88].

In this Letter, we introduce QST with conditional GANs (QST-CGAN). Leveraging a CGAN architecture, complemented by custom layers for representing a quantum state in the form of a density matrix, we show that adversarial learning can be a powerful tool for QST. The QST-CGAN is different from RBM-based methods since it learns a map between the data and the quantum state instead of a probability distribution. The custom layers we introduce bridge a gap between machine learning and quantum information processing; they enable many further applications beyond the QST-CGAN presented here. We benchmark the QST-CGAN on reconstruction of optical quantum states from simulated data and show an example with real experimental data. The QST-CGAN performance is superior to that of maximum-likelihood reconstruction methods in terms of reconstruction fidelity, number of iterative steps, and amount of measurement data required. We also show that a QST-CGAN can reconstruct quantum states in a single pass through the network if it has been pretrained on simulated data.

Our reconstruction method is versatile, general, and ready to be applied for QST of intermediate-scale quantum systems, which are widely explored in current experiments [4]. In Refs. [89,90], we provide more details on our implementation (including data and code) and also discuss classification of quantum states with neural networks.

Quantum state tomography with maximum likelihood estimation.—Quantum state tomography estimates the quantum state (a state vector $|\psi\rangle$ or a density matrix ρ) from measurements of Hermitian operators \mathcal{O} [28,91]. The operators are usually positive-operator-valued measures (POVMs), a set of positive semidefinite matrices $\{\mathcal{O}_i\}$ that sum to identity, $\sum_i^k \mathcal{O}_i = I$, representing a measurement with k possible outcomes. The probability of each outcome is given by $\text{tr}(\mathcal{O}_i \rho)$. A set of operators that allows for the complete characterization of a quantum state is called informationally complete [92].

In an experiment, single-shot measurements are repeated over an ensemble of identical states to collect statistics: the frequencies d_i of POVM outcomes. These frequencies give an estimate of the expectation values $\text{tr}(\mathcal{O}_i \rho)$, where ρ is the density matrix describing the state. The outcomes of many different POVMs can be combined to form a linear system of equations $\mathbf{d} = A \rho_f$, where ρ_f is the flattened density matrix and A is the “sensing matrix” determined by the choice of POVMs [93]. Solving this system of equations by linear inversion methods to obtain ρ can fail, either due to the statistical nature of the (noisy) measurement or due to a high condition number for A [93].

An alternative to linear inversion methods is maximum likelihood estimation. In MLE, the likelihood function [40,94] $L(\rho'|\mathbf{d}) = \prod_i [\text{tr}(\rho' \mathcal{O}_i)]^{d_i}$ is maximized to find the best estimate ρ' for reproducing the experimental data.

In this Letter, we take a different approach by applying CGANs to find ρ' .

Conditional generative adversarial networks.—In generative adversarial learning, a generator G and a discriminator D compete to learn a mapping from some prior noise distribution to a data distribution [74]. The generator and the discriminator are parametrized nonlinear functions [parameters (θ_D, θ_G)], usually multilayered neural networks. The generator takes an input $\mathbf{z} \sim p_z(\mathbf{z})$ from the noise distribution $p_z(\mathbf{z})$ and generates an output $G(\mathbf{z}; \theta_G)$. The discriminator takes an input \mathbf{q} and outputs a probability $D(\mathbf{q}; \theta_D)$ that it belongs to the data distribution p_{data} .

The parameters of G and D are optimized alternatively such that the generator produces outputs that resemble the data and thus fool the discriminator, and the discriminator becomes better at detecting fake (generated) output. In each optimization step, θ_D is updated to maximize the expectation value

$$E_{\mathbf{y} \sim p_{\text{data}}} [\ln(D(\mathbf{y}; \theta_D))] + E_{\mathbf{z} \sim p_z} \{\ln[1 - D(G(\mathbf{z}; \theta_G); \theta_D)]\}, \quad (1)$$

where \mathbf{y} denotes samples from the data. Then, θ_G is updated to minimize

$$E_{\mathbf{z} \sim p_z} \{\ln[1 - D(G(\mathbf{z}; \theta_G); \theta_D)]\}. \quad (2)$$

In this way, the generator learns to map elements from a noise distribution to data as $G: \mathbf{z} \rightarrow \mathbf{y}$ [74,76].

However, since the generator input is random, we have no control over the output. This issue is solved by using a conditional generative adversarial network [76,85]. In a CGAN, the generator and discriminator output is conditioned on some variable \mathbf{x} . This conditioning allows the generator to learn the mapping $G: \mathbf{x}, \mathbf{z} \rightarrow \mathbf{y}$ [76]. The optimization of parameters for the CGAN is done as before, by maximizing Eq. (1) and minimizing Eq. (2); the only difference is that the outputs now are $D(\mathbf{x}, \mathbf{y}; \theta_D)$ and $G(\mathbf{x}, \mathbf{z}; \theta_G)$. This CGAN approach is very flexible and can be used to find complex maps between inputs and outputs. The flexibility stems from using the discriminator network for evaluation instead of, or in addition to, a simpler loss function.

Quantum state tomography using conditional generative adversarial networks.—We now adapt the CGAN framework to the problem of QST. In our approach, illustrated in Fig. 1, the conditioning input to the generator is the measurement statistics and the measurement operators ($\mathbf{x} \rightarrow \mathbf{d}, \{\mathcal{O}_i\}$). The generated output is a density matrix ρ_G . We find that we do not need to provide any input noise \mathbf{z} , consistent with the results in Ref. [76].

The discriminator takes as input the experimental measurement statistics \mathbf{d} (as the conditioning variable) and generated measurement statistics calculated from $\text{tr}(\mathcal{O}_i \rho_G)$. The output from the discriminator is a set of numbers




FIG. 1. Illustration of the CGAN architecture for QST. Data \mathbf{d} sampled from measurements of a set of measurement operators $\{O_i\}$ on a quantum state is fed into both the generator G and the discriminator D . The other input to D is the generated statistics from G . The next to last layer of G outputs a physical density matrix and the last layer computes measurement statistics using this density matrix. The discriminator compares the measurement data and the generated data for each measurement operator and outputs a probability that they match.

describing how well the generated measurement statistics match the data. This partitioning of the evaluation of the generated statistics is inspired by the PATCHGAN architecture of Ref. [76]. If the generator has managed to learn the correct density matrix, the discriminator will not be able to distinguish the generated statistics from the true data.

The adaption of the CGAN architecture to QST requires us to introduce two custom layers at the end of the generator neural network. First, we add a “DensityMatrix” layer, which takes the unconstrained intermediate output of the generator, molds it into a lower triangular complex-valued matrix T_G with real entries on the diagonal, constructs $T_G^\dagger T_G$, and normalizes the resulting matrix to have unit trace. This method is inspired by the Cholesky decomposition [40]. It ensures that the output ρ_G is a valid density matrix: Hermitian, positive, and having unit trace. A similar idea was found independently in Ref. [73].

Second, we add an “Expectation” layer that combines the output ρ_G with the given measurement operators $\{O_i\}$ to compute the generated measurement statistics for each measurement outcome as $\text{tr}(O_i \rho_G)$. These two custom layers do not have any trainable parameters. They are only present to enforce the rules of quantum mechanics in the neural networks. This is akin to regularization [95] and normalization [96] in neural networks. We note that our two custom layers could be used to augment any deep-learning neural-network architecture for QST, e.g., Refs. [72,73].

We train the QST-CGAN using standard gradient-based optimization techniques, e.g., Adam [97] with learning-rate scheduling, starting from random initial values for the parameters (θ_D, θ_G) . In this way, data from one experiment can be used to estimate the density matrix of the state in that experiment. However, when reconstructing ρ from another experiment, the QST-CGAN must start from zero again. We can avoid this reset by pretraining on simulated data corresponding to the type of state(s) and noise that is expected to be present in the experiment. The

reconstruction from experimental data then requires less additional training; it even becomes possible to do “single-shot reconstruction” with a single evaluation by the pretrained generator.

We note that adding L_1 loss to Eq. (2) as suggested in Ref. [76] proved helpful in training the QST-CGAN [89] and was used for all results displayed below, but was not necessary to obtain good results. Similarly, adding a gradient penalty [98] to Eq. (1) improved results for single-shot reconstruction.

Benchmarking CGAN quantum state tomography.—To benchmark the QST-CGAN method, we test it on reconstruction of optical quantum states and compare its performance to two MLE methods—iterative MLE (iMLE) [41] and accelerated projected-gradient-based MLE (APG-MLE) [94]. In iMLE, projection operators determined by the measurement statistics are iteratively applied to a random initial density matrix until convergence. The final result is an estimated density matrix ρ' that maximizes the likelihood function $L(\rho'|\mathbf{d})$. In the APG-MLE method, ideas from convex optimization are used to enable faster convergence.

Optical quantum states describe quantized single-mode electromagnetic fields (harmonic oscillators). Our choice of optical quantum states for testing the QST-CGAN was motivated by the existence of visual representations, e.g., Wigner functions [99], for these states, seeing how CGANs have mainly been applied to image processing. However, we stress that the QST-CGAN approach is general and can be applied to any type of quantum system with any type of observable [89].

Some of the common observables for optical quantum states are instances of a displace-and-measure technique. For example, the photon-number distribution obtained after applying a displacement β is the generalized Q function [100]: $Q_n^\beta = \text{tr}[|n\rangle\langle n|D(-\beta)\rho D^\dagger(-\beta)]$, where $|n\rangle$ is the Fock state with n photons, $D(\beta) = e^{\beta a^\dagger - \beta^* a}$ is the displacement operator, and $a(a^\dagger)$ is the bosonic creation (annihilation) operator of the electromagnetic mode. The Husimi Q function (photon field quadratures) is $(1/\pi)Q_0^\beta$ and the Wigner function (photon parity) is $W(\beta) = (2/\pi) \sum_n (-1)^n Q_n^\beta$. The measurement data we consider in the following are samples of Q_0^β and $W(\beta)$ at certain β , as illustrated in Fig. 2.

A state ρ in a truncated Hilbert space with size N is specified by up to $N^2 - 1$ real numbers [93,101] (we use $N = 32$). Thus, in general, informational completeness requires displacements and measurements to be carried out such that \mathbf{d} has at least $N^2 - 1$ elements. However, note that the required number of elements in \mathbf{d} for reconstruction can be lower, $\propto rN$, if ρ has low rank r [102].

Results.—In Fig. 3(a), we compare the reconstruction fidelity for the QST-CGAN and MLE methods as a function of the number of iterations. One iteration is one update of all the weights (θ_D, θ_G) for the QST-CGAN




FIG. 2. Observables for an optical quantum state, the “cat state” $|\alpha\rangle + |-\alpha\rangle$ (up to a normalization), with coherent amplitude $\alpha = 2$. (a) The Wigner function. (b) The Husimi Q function. The stars mark specific β used to sample the data used as input to the QST-CGAN.

(a single gradient-descent step), one application of the projection operators in iMLE, and one update of the density matrix for APG-MLE. We find that the QST-CGAN converges to a fidelity >0.999 in about 2 orders of magnitude fewer iterations than the MLE methods. Note that the choice of network architecture and training parameters will affect the speed of convergence and the computational cost of one iteration for the QST-CGAN.

Next, we investigate, in Fig. 3(b), how many data points are required as input to reach high reconstruction fidelity. We find that the QST-CGAN approach starts outperforming the MLE methods around $N = 32$ data points and reaches fidelities close to unity already with <100 data points, while the MLE methods require ~ 1000 data points to attain good fidelity (an RBM-based reconstruction of a similar state also requires thousands of data points to reach




FIG. 3. QST-CGAN performance. The data are the Husimi Q function of the cat state in Fig. 2(b). (a) Reconstruction fidelity $F(\rho, \rho') = \text{tr}(\sqrt{\sqrt{\rho}\rho'\sqrt{\rho}})^2$ as a function of iterations for the QST-CGAN (red), iMLE (blue), and APG-MLE (dashed black). We use 1024 displacements β in a 32×32 grid. The weights of the QST-CGAN and the starting density matrix of the iMLE are randomly initialized. The APG-MLE runs 13 iterations of conjugate-gradient line search from the maximally mixed state before switching to APG. The solid lines show the mean F for 100 runs; the shaded areas show one standard deviation from the mean. (b) Average F as a function of the number of β . For each number, 10 sets of displacements are randomly selected from within a disk with $|\beta| \leq 5$ for the state in Fig. 2. We show the average F reached after 1000 iterations for QST-CGAN and iMLE and 10 000 iterations for APG-MLE.




FIG. 4. (a) Reconstruction of a Wigner-negative state by a QST-CGAN from (b) noisy experimental data. Inset: the target state. The reconstruction uses 4281 data points of the Wigner function measured for β inside the dashed circle. The data outside the circle, e.g., the Wigner-negative region in the top left, are not as reliable due to measurement calibration problems at higher photon numbers. We also attempt reconstruction with a subset of the data points inside the circle and find that ~ 600 data points are enough to achieve a fidelity ~ 0.9 with the full reconstruction.

high fidelity [68]). Note that the rank $r = 1$, since ρ is a pure state.

Experimental state reconstruction from parity measurements.—The benchmarking of the QST-CGAN so far has been on simulated data. We now demonstrate, in Fig. 4, that our QST-CGAN can reconstruct a noisy state from experimental data. In this particular experiment, a superconducting transmon qubit was used to generate a Wigner-negative state in a resonator [103], by applying a selective number-dependent arbitrary phase [104,105] of π to $|0\rangle$ and $|1\rangle$ of a coherent state $|\alpha=1\rangle$. Despite significant state-preparation-and-measurement noise, the QST-CGAN still manages to reconstruct the data well from measurements of the Wigner function, even when using only $\sim 15\%$ of the measurement data.

Single-shot reconstruction with pretraining.—We now pretrain the QST-CGAN on a dataset with several thousand cat states similar to Fig. 2 by selecting $|\alpha| \in [1, 3]$ randomly




FIG. 5. Single-shot reconstructions of 200 cat states (cf. Fig. 2, $|\alpha| \in [1, 3]$, up to six coherent states in superposition), using a pretrained QST-CGAN. (a) Fidelity distribution of the reconstructions after training on a 32×32 grid of data points. (b) Average fidelity (solid line) within one standard deviation (shaded region) after further iterations.

with up to six coherent states in superposition. As shown in Fig. 5(a), this pretrained network is then able to perform single-shot reconstructions for different cat states with a high average fidelity ~ 0.98 . It turned out to be difficult to find a learning strategy enabling further improvement of the fidelity with just a few more iterations for each state, but with tens of iterations a clear improvement is observed [Fig. 2(b)]. The pretrained network thus does not have to iterate many times from an initial random guess for each state, as is the case for the results in Fig. 3 and most other reconstruction methods in use today, resulting in a 4 orders of magnitude faster reconstruction than in Fig. 3(a).

Conclusion and outlook.—In this Letter, we have adapted the CGAN architecture for use in quantum state tomography. The adaption relies on the introduction of two custom layers, which enforce the properties of a density matrix and allows calculation of expectation values of measurements. We showed that our QST-CGAN clearly outperforms MLE reconstruction methods: the QST-CGAN consistently reconstructs states with higher fidelity, needing $\sim 100\times$ fewer iterations and $\sim 10\times$ fewer data points to do so in the examples we showed. Furthermore, we showed that we can pretrain the QST-CGAN on classes of quantum states and achieve high fidelity for single-shot reconstruction.

Looking to the future, we note that the custom layers we introduced could be included into other types of neural networks, e.g., transformers [72], for both QST and other applications in quantum information processing. Any QST method that reconstructs the full density matrix will scale poorly since we must determine exponentially many density matrix elements. Future work could therefore use custom density matrix layers to output reduced quantum state representations with the CGAN approach. The CGAN approach has potential for denoising measurement data by pretraining on simulated noisy data. We further envisage the application of QST-CGAN for adaptive tomography [30,59], by choosing next measurements around the points where the discriminator finds that the reconstructed data do not match the experimental data well [67].

The QST-CGAN was written in TENSORFLOW [106]. Visualization and generation of quantum states was done in QUTIP [107,108]. We thank Marina Kudra for providing the experimental data for Fig. 4. We acknowledge useful discussions with Yong Lu, Marina Kudra, Florian Marquardt, Steven Girvin, Simone Gasparinetti, Fernando Quijandría, and Ingrid Strandberg. We also thank Abhijeet Melkani and Clemens Gneiting for comments on the manuscript. S. A. and A. F. K. acknowledge support from the Knut and Alice Wallenberg Foundation through the Wallenberg Centre for Quantum Technology (WACQT). C. S. M. acknowledges that the project that gave rise to these results received the support of a fellowship from “la Caixa” Foundation (ID 100010434) and from the European Union’s Horizon 2020 Research and Innovation

Programme under the Marie Skłodowska-Curie Grant Agreement No. 847648. The fellowship code is LCF/BQ/PI20/11760026. F.N. is supported in part by: Nippon Telegraph and Telephone Corporation (NTT) Research, the Japan Science and Technology Agency (JST) [via the Quantum Leap Flagship Program (Q-LEAP), the Moonshot R&D Grant No. JPMJMS2061, and the Centers of Research Excellence in Science and Technology (CREST) Grant No. JPMJCR1676], the Japan Society for the Promotion of Science (JSPS) [via the Grants-in-Aid for Scientific Research (KAKENHI) Grant No. JP20H00134 and the JSPS-RFBR Grant No. JPJSBP120194828], the Army Research Office (ARO) (Grant No. W911NF-18-1-0358), the Asian Office of Aerospace Research and Development (AOARD) (via Grant No. FA2386-20-1-4069), and the Foundational Questions Institute Fund (FQXi) via Grant No. FQXi-IAF19-06.

*shahnawaz.ahmed95@gmail.com

†anton.frisk.kockum@chalmers.se

- [1] R. P. Feynman, Simulating physics with computers, *Int. J. Theor. Phys.* **21**, 467 (1982).
- [2] A. Montanaro, Quantum algorithms: An overview, *npj Quantum Inf.* **2**, 15023 (2016).
- [3] G. Wendin, Quantum information processing with superconducting circuits: A review, *Rep. Prog. Phys.* **80**, 106001 (2017).
- [4] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [5] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature (London)* **574**, 505 (2019).
- [6] I. M. Georgescu, S. Ashhab, and F. Nori, Quantum simulation, *Rev. Mod. Phys.* **86**, 153 (2014).
- [7] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature (London)* **549**, 242 (2017).
- [8] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, Toward the first quantum simulation with quantum speedup, *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9456 (2018).
- [9] J. Argüello-Luengo, A. González-Tudela, T. Shi, P. Zoller, and J. I. Cirac, Analogue quantum chemistry simulation, *Nature (London)* **574**, 215 (2019).
- [10] H. Ma, M. Govoni, and G. Galli, Quantum simulations of materials on near-term quantum computers, *npj Comput. Mater.* **6**, 85 (2020).
- [11] H. J. Kimble, The quantum internet, *Nature (London)* **453**, 1023 (2008).
- [12] J. Yin *et al.*, Entanglement-based secure quantum cryptography over 1,120 kilometres, *Nature (London)* **582**, 501 (2020).
- [13] J. Q. You and F. Nori, Atomic physics and quantum optics using superconducting circuits, *Nature (London)* **474**, 589 (2011).

- [14] X. Gu, A. F. Kockum, A. Miranowicz, Y.-x. Liu, and F. Nori, Microwave photonics with superconducting quantum circuits, *Phys. Rep.* **718–719**, 1 (2017).
- [15] M. Bal, C. Deng, J.-L. Orgiazzi, F. Ong, and A. Lupascu, Ultrasensitive magnetic field detection using a single artificial atom, *Nat. Commun.* **3**, 1324 (2012).
- [16] C. L. Degen, F. Reinhard, and P. Cappellaro, Quantum sensing, *Rev. Mod. Phys.* **89**, 035002 (2017).
- [17] S. Pirandola, B. R. Bardhan, T. Gehring, C. Weedbrook, and S. Lloyd, Advances in photonic quantum sensing, *Nat. Photonics* **12**, 724 (2018).
- [18] T. Gefen, A. Rotem, and A. Retzker, Overcoming resolution limits with quantum sensing, *Nat. Commun.* **10**, 4992 (2019).
- [19] J. Wang, F. Sciarrino, A. Laing, and M. G. Thompson, Integrated photonic quantum technologies, *Nat. Photonics* **14**, 273 (2020).
- [20] S. Lloyd, Enhanced sensitivity of photodetection via quantum illumination, *Science* **321**, 1463 (2008).
- [21] M. Korobko, Y. Ma, Y. Chen, and R. Schnabel, Quantum expander for gravitational-wave observatories, *Light* **8**, 118 (2019).
- [22] S. Barzanjeh, S. Pirandola, D. Vitali, and J. M. Fink, Microwave quantum illumination using a digital receiver, *Sci. Adv.* **6**, eabb0451 (2020).
- [23] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2000).
- [24] C. M. Caves, I. H. Deutsch, and R. Blume-Kohout, Physical-resource requirements and the power of quantum computation, *J. Opt. B* **6**, S801 (2004).
- [25] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature (London)* **567**, 209 (2019).
- [26] G. M. D’Ariano, M. G. Paris, and M. F. Sacchi, Quantum tomography, in *Advances in Imaging and Electron Physics* (Academic Press, 2003), Vol. **128**, pp. 205–308, [arXiv: quant-ph/0302028](https://arxiv.org/abs/quant-ph/0302028).
- [27] Y.-x. Liu, L. F. Wei, and F. Nori, Tomographic measurements on superconducting qubit states, *Phys. Rev. B* **72**, 014547 (2005).
- [28] A. I. Lvovsky and M. G. Raymer, Continuous-variable optical quantum-state tomography, *Rev. Mod. Phys.* **81**, 299 (2009).
- [29] S. Deléglise, I. Dotsenko, C. Sayrin, J. Bernu, M. Brune, J. M. Raimond, and S. Haroche, Reconstruction of non-classical cavity field states with snapshots of their decoherence, *Nature (London)* **455**, 510 (2008).
- [30] G. M. D’Ariano, D. F. Magnani, and P. Perinotti, Adaptive Bayesian and frequentist data processing for quantum tomography, *Phys. Lett. A* **373**, 1111 (2009).
- [31] 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).
- [32] S. T. Flammia and Y.-K. Liu, Direct Fidelity Estimation from Few Pauli Measurements, *Phys. Rev. Lett.* **106**, 230501 (2011).
- [33] D. Petz and L. Ruppert, Optimal quantum-state tomography with known parameters, *J. Phys. A* **45**, 085306 (2012).
- [34] T. Baumgratz, D. Gross, M. Cramer, and M. B. Plenio, Scalable Reconstruction of Density Matrices, *Phys. Rev. Lett.* **111**, 020401 (2013).
- [35] A. Miranowicz, S. K. Ozdemir, J. Bajer, G. Yusa, N. Imoto, Y. Hirayama, and F. Nori, Quantum state tomography of large nuclear spins in a semiconductor quantum well: Optimal robustness against errors as quantified by condition numbers, *Phys. Rev. B* **92**, 075312 (2015).
- [36] Z. Hou, H.-S. Zhong, Y. Tian, D. Dong, B. Qi, L. Li, Y. Wang, F. Nori, G.-Y. Xiang, C.-F. Li, and G.-C. Guo, Full reconstruction of a 14-qubit state within four hours, *New J. Phys.* **18**, 083036 (2016).
- [37] J. G. Titchener, M. Gräfe, R. Heilmann, A. S. Solntsev, A. Szameit, and A. A. Sukhorukov, Scalable on-chip quantum state tomography, *npj Quantum Inf.* **4**, 19 (2018).
- [38] A. Rocchetto, S. Aaronson, S. Severini, G. Carvacho, D. Poderini, I. Agresti, M. Bentivegna, and F. Sciarrino, Experimental learning of quantum states, *Sci. Adv.* **5**, eaau1946 (2019).
- [39] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [40] K. Banaszek, G. M. D’Ariano, M. G. A. Paris, and M. F. Sacchi, Maximum-likelihood estimation of the density matrix, *Phys. Rev. A* **61**, 010304(R) (1999).
- [41] A. I. Lvovsky, Iterative maximum-likelihood reconstruction in quantum homodyne tomography, *J. Opt. B* **6**, S556 (2004).
- [42] J. Reháček, Z. Hradil, E. Knill, and A. I. Lvovsky, Diluted maximum-likelihood algorithm for quantum tomography, *Phys. Rev. A* **75**, 042108 (2007).
- [43] 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).
- [44] R. Blume-Kohout, Optimal, reliable estimation of quantum states, *New J. Phys.* **12**, 043034 (2010).
- [45] C. Granade, J. Combes, and D. G. Cory, Practical Bayesian tomography, *New J. Phys.* **18**, 033024 (2016).
- [46] E. Bolduc, G. C. Knee, E. M. Gauger, and J. Leach, Projected gradient descent algorithms for quantum state tomography, *npj Quantum Inf.* **3**, 44 (2017).
- [47] 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, Efficient tomography of a quantum many-body system, *Nat. Phys.* **13**, 1158 (2017).
- [48] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac, Neural-Network Quantum States, String-Bond States, and Chiral Topological States, *Phys. Rev. X* **8**, 011006 (2018).
- [49] G. Tóth, W. Wieczorek, D. Gross, R. Krischek, C. Schwemmer, and H. Weinfurter, Permutationally Invariant Quantum Tomography, *Phys. Rev. Lett.* **105**, 250403 (2010).
- [50] 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).
- [51] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, Reconstructing quantum states with generative models, *Nat. Mach. Intell.* **1**, 155 (2019).
- [52] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [53] Q. Xu and S. Xu, Neural network state estimation for full quantum state tomography, [arXiv:1811.06654](https://arxiv.org/abs/1811.06654).
- [54] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, *Nat. Phys.* **14**, 447 (2018).
- [55] G. Torlai and R. G. Melko, Latent Space Purification via Neural Density Operators, *Phys. Rev. Lett.* **120**, 240503 (2018).
- [56] 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).
- [57] A. M. 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).
- [58] J. Gray, L. Banchi, A. Bayat, and S. Bose, Machine-Learning-Assisted Many-Body Entanglement Measurement, *Phys. Rev. Lett.* **121**, 150503 (2018).
- [59] Y. Quek, S. Fort, and H. K. Ng, Adaptive quantum state tomography with neural networks, [arXiv:1812.06693](https://arxiv.org/abs/1812.06693).
- [60] S. Yu, F. Albarrán-Arriagada, J. C. Retamal, Y.-T. Wang, W. Liu, Z.-J. Ke, Y. Meng, Z.-P. Li, J.-S. Tang, E. Solano, L. Lamata, C.-F. Li, and G.-C. Guo, Reconstruction of a photonic qubit state with reinforcement learning, *Adv. Quantum Technol.* **2**, 1800074 (2019).
- [61] A. Melkani, C. Gneiting, and F. Nori, Eigenstate extraction with neural-network tomography, *Phys. Rev. A* **102**, 022412 (2020).
- [62] T. Weiss and O. Romero-Isart, Quantum motional state tomography with nonquadratic potentials and neural networks, *Phys. Rev. Research* **1**, 033157 (2019).
- [63] Y. Liu, D. Wang, S. Xue, A. Huang, X. Fu, X. Qiang, P. Xu, H.-L. Huang, M. Deng, C. Guo, X. Yang, and J. Wu, Variational quantum circuits for quantum state tomography, *Phys. Rev. A* **101**, 052316 (2020).
- [64] T. Xin, X. Nie, X. Kong, J. Wen, D. Lu, and J. Li, Quantum Pure State Tomography via Variational Hybrid Quantum-Classical Method, *Phys. Rev. Applied* **13**, 024013 (2020).
- [65] S. Lloyd and C. Weedbrook, Quantum Generative Adversarial Learning, *Phys. Rev. Lett.* **121**, 040502 (2018).
- [66] M. Kieferová and N. Wiebe, Tomography and generative training with quantum Boltzmann machines, *Phys. Rev. A* **96**, 062327 (2017).
- [67] L. Hu, S.-H. Wu, W. Cai, Y. Ma, X. Mu, Y. Xu, H. Wang, Y. Song, D.-L. Deng, C.-L. Zou, and L. Sun, Quantum generative adversarial learning in a superconducting quantum circuit, *Sci. Adv.* **5**, eaav2761 (2019).
- [68] E. S. Tiunov, V. V. Tiunova (Vyborova), A. E. Ulanov, A. I. Lvovsky, and A. K. Fedorov, Experimental quantum homodyne tomography via machine learning, *Optica* **7**, 448 (2020).
- [69] G. Carleo, Y. Nomura, and M. Imada, Constructing exact representations of quantum many-body systems with deep neural networks, *Nat. Commun.* **9**, 5322 (2018).
- [70] A. Rocchetto, E. Grant, S. Strelchuk, G. Carleo, and S. Severini, Learning hard quantum distributions with variational autoencoders, *npj Quantum Inf.* **4**, 28 (2018).
- [71] Z. Cai and J. Liu, Approximating quantum many-body wave functions using artificial neural networks, *Phys. Rev. B* **97**, 035116 (2018).
- [72] P. Cha, P. Ginsparg, F. Wu, J. Carrasquilla, P. L. McMahon, and E.-A. Kim, Attention-based quantum tomography, [arXiv:2006.12469](https://arxiv.org/abs/2006.12469).
- [73] S. Lohani, B. T. Kirby, M. Brodsky, O. Danaci, and R. T. Glasser, Machine learning assisted quantum state estimation, *Mach. Learn.* **1**, 035007 (2020).
- [74] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, Generative adversarial nets, in *Advances in Neural Information Processing Systems* (Montréal, 2014), Vol. 3, pp. 2672–2680.
- [75] S. Mahdizadehghadam, A. Panahi, and H. Krim, Sparse generative adversarial network, in *Proceedings of the 2019 International Conference on Computer Vision Workshop, ICCVW 2019* (IEEE, New York, 2019), pp. 3063–3071, [arXiv:1908.08930](https://arxiv.org/abs/1908.08930).
- [76] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, Image-to-image translation with conditional adversarial networks, in *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, New York, 2017), pp. 5967–5976, [arXiv:1611.07004](https://arxiv.org/abs/1611.07004).
- [77] T. Karras, S. Laine, and T. Aila, A style-based generator architecture for generative adversarial networks, in *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, New York, 2019), pp. 4396–4405, [arXiv:1812.04948](https://arxiv.org/abs/1812.04948).
- [78] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, Analyzing and improving the image quality of StyleGAN, [arXiv:1912.04958](https://arxiv.org/abs/1912.04958).
- [79] Z. Yang, Z. Hu, C. Dyer, E. P. Xing, and T. Berg-Kirkpatrick, Unsupervised text style transfer using language models as discriminators, in *Advances in Neural Information Processing Systems* (NeurIPS) (Montréal, 2018), pp. 7287–7298, [arXiv:1805.11749](https://arxiv.org/abs/1805.11749).
- [80] S. Subramanian, S. Rajeswar, A. Sordoni, A. Trischler, A. Courville, and C. Pal, Towards text generation with adversarially learned neural outlines, in *Advances in Neural Information Processing Systems* (NeurIPS) (Montréal, 2018), pp. 7551–7563.
- [81] T. Xu, P. Zhang, Q. Huang, H. Zhang, Z. Gan, X. Huang, and X. He, AttnGAN: Fine-grained text to image generation with attentional generative adversarial networks, in *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* (IEEE, New York, 2018), pp. 1316–1324, [arXiv:1711.10485](https://arxiv.org/abs/1711.10485).
- [82] Y. Mirsky and W. Lee, The creation and detection of deepfakes: A survey, *ACM Comput. Syst.* **54**, 1 (2021).
- [83] R. Zellers, A. Holtzman, H. Rashkin, Y. Bisk, A. Farhadi, F. Roesner, and Y. Choi, Defending against neural fake news, [arXiv:1905.12616](https://arxiv.org/abs/1905.12616).

- [84] S. W. Kim, Y. Zhou, J. Phillion, A. Torralba, and S. Fidler, Learning to simulate dynamic environments with GameGAN, [arXiv:2005.12126](https://arxiv.org/abs/2005.12126).
- [85] M. Mirza and S. Osindero, Conditional generative adversarial nets, [arXiv:1411.1784](https://arxiv.org/abs/1411.1784).
- [86] X. Yang, M. Kahnt, D. Bruckner, A. Schropp, Y. Fam, J. Becher, J. D. Grunwaldt, T. L. Sheppard, and C. G. Schroer, Tomographic reconstruction with a generative adversarial network, *J. Synchrotron Radiat.* **27**, 486 (2020).
- [87] Z. Liu, T. Bicer, R. Kettimuthu, D. Gursoy, F. De Carlo, and I. Foster, TomoGAN: Low-dose synchrotron x-ray tomography with generative adversarial networks: Discussion, *J. Opt. Soc. Am. A* **37**, 422 (2020).
- [88] M. Selim, J. Zhang, B. Fei, G.-q. Zhang, and J. Chen, STAN-CT: Standardizing CT image using generative adversarial network, [arXiv:2004.01307](https://arxiv.org/abs/2004.01307).
- [89] S. Ahmed, C. Sánchez Muñoz, F. Nori, and A. F. Kockum, companion paper, Classification and reconstruction of optical quantum states with deep neural networks, *Phys. Rev. Research*, **3**, 033278 (2021)..
- [90] Code for quantum state tomography with conditional generative adversarial networks is available at <http://doi.org/10.5281/zenodo.5105469>.
- [91] M. G. Raymer, Measuring the quantum mechanical wave function, *Contemp. Phys.* **50**, 323 (2009).
- [92] G. M. D Ariano, P. Perinotti, and M. F. Sacchi, Informationally complete measurements and group representation, *J. Opt. B* **6**, S487 (2004).
- [93] C. Shen, R. W. Heeres, P. Reinhold, L. Jiang, Y.-K. Liu, R. J. Schoelkopf, and L. Jiang, Optimized tomography of continuous variable systems using excitation counting, *Phys. Rev. A* **94**, 052327 (2016).
- [94] J. Shang, Z. Zhang, and H. K. Ng, Superfast maximum-likelihood reconstruction for quantum tomography, *Phys. Rev. A* **95**, 062336 (2017).
- [95] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov, Improving neural networks by preventing co-adaptation of feature detectors, [arXiv:1207.0580](https://arxiv.org/abs/1207.0580).
- [96] E. Hoffer, R. Banner, I. Golan, and D. Soudry, Norm matters: Efficient and accurate normalization schemes in deep networks, in *Advances in Neural Information Processing Systems (NeurIPS)* (Montréal, 2018), pp. 2160–2170, [arXiv:1803.01814](https://arxiv.org/abs/1803.01814).
- [97] D. P. Kingma and J. L. Ba, Adam: A method for stochastic optimization, in *Proceedings of the 3rd International Conference on Learning Representations, ICLR 2015—Conference Track Proceedings* (San Diego, 2015), [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [98] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, Improved training of Wasserstein GANs, in *Advances in Neural Information Processing Systems* (Montréal, 2017), pp. 5767–5777.
- [99] J. Weinbub and D. K. Ferry, Recent advances in Wigner function approaches, *Appl. Phys. Rev.* **5**, 041104 (2018).
- [100] G. Kirchmair, B. Vlastakis, Z. Leghtas, S. E. Nigg, H. Paik, E. Ginossar, M. Mirrahimi, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Observation of quantum state collapse and revival due to the single-photon Kerr effect, *Nature (London)* **495**, 205 (2013).
- [101] D. Sych, J. Reháček, Z. Hradil, G. Leuchs, and L. L. Sánchez-Soto, Informational completeness of continuous-variable measurements, *Phys. Rev. A* **86**, 052123 (2012).
- [102] A. Kalev, R. L. Kosut, and I. H. Deutsch, Quantum tomography protocols with positivity are compressed sensing protocols, *npj Quantum Inf.* **1**, 15018 (2015).
- [103] M. Kudra, J. Biznářová, F. Roudsari, J. J. Burnett, D. Niepce, S. Gasparinetti, B. Wickman, and P. Delsing, High quality three-dimensional aluminum microwave cavities, *Appl. Phys. Lett.* **117**, 070601 (2020).
- [104] R. W. Heeres, B. Vlastakis, E. Holland, S. Krastanov, V. V. Albert, L. Frunzio, L. Jiang, and R. J. Schoelkopf, Cavity State Manipulation Using Photon-Number Selective Phase Gates, *Phys. Rev. Lett.* **115**, 137002 (2015).
- [105] T. Fösel, S. Krastanov, F. Marquardt, and L. Jiang, Efficient cavity control with SNAP gates, [arXiv:2004.14256](https://arxiv.org/abs/2004.14256).
- [106] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, Tensorflow: A system for large-scale machine learning (2016), <https://doi.org/10.5281/zenodo.5043456>.
- [107] J. R. Johansson, P. D. Nation, and F. Nori, QuTiP: An open-source Python framework for the dynamics of open quantum systems, *Comput. Phys. Commun.* **183**, 1760 (2012).
- [108] J. R. Johansson, P. D. Nation, and F. Nori, QuTiP: An open-source Python framework for the dynamics of open quantum systems, *Comput. Phys. Commun.* **184**, 1234 (2013).