

Precision Bounds on Continuous-Variable State Tomography Using Classical Shadows

Srilekha Gandhari,^{1,*} Victor V. Albert,^{1,2} Thomas Gerrits,² Jacob M. Taylor,^{1,2,3} and Michael J. Gullans^{1,2}

¹*Joint Center for Quantum Information and Computer Science, NIST–University of Maryland, College Park, Maryland 20742, USA*

²*National Institute of Standards and Technology, 100 Bureau Drive, Gaithersburg, Maryland 20899, USA*

³*Joint Quantum Institute, NIST–University of Maryland, College Park, Maryland 20742, USA*



(Received 8 February 2023; revised 15 December 2023; accepted 16 January 2024; published 18 March 2024)

Shadow tomography is a framework for constructing succinct descriptions of quantum states using randomized measurement bases, called “classical shadows,” with powerful methods to bound the estimators used. We recast existing experimental protocols for continuous-variable quantum state tomography in the classical-shadow framework, obtaining rigorous bounds on the number of independent measurements needed for estimating density matrices from these protocols. We analyze the efficiency of homodyne, heterodyne, photon-number-resolving, and photon-parity protocols. To reach a desired precision on the classical shadow of an N -photon density matrix with high probability, we show that homodyne detection requires order $\mathcal{O}(N^{4+1/3})$ measurements in the worst case, whereas photon-number-resolving and photon-parity detection require $\mathcal{O}(N^4)$ measurements in the worst case (both up to logarithmic corrections). We benchmark these results against numerical simulation as well as experimental data from optical homodyne experiments. We find that numerical and experimental analyses of homodyne tomography match closely with our theoretical predictions. We extend our single-mode results to an efficient construction of multimode shadows based on local measurements.

DOI: [10.1103/PRXQuantum.5.010346](https://doi.org/10.1103/PRXQuantum.5.010346)

I. INTRODUCTION

The ability to estimate states accurately with as few measurements as possible—the primary goal of quantum state tomography—yields an indispensable tool in quantum information processing [1]. State characterization is necessary for realizing quantum technologies in finite-dimensional tensor-product quantum systems governed by discrete variables (DVs), as well as in quantum systems governed by continuous variables (CVs) [2], such as electromagnetic or mechanical modes.

The recent development of classical-shadow tomography yields a succinct way to learn information about a DV quantum state through randomly chosen measurements, in such a way that the learned information can later be used to predict properties of the state [3–5]. A key benefit of shadow tomography is that it comes with

rigorously proven guarantees on the minimum number of samples required to achieve high accuracy with high probability. However, the best guarantees require structures from finite-dimensional DV spaces, such as state and unitary designs [4] or symmetric, informationally complete positive operator-valued measures (SIC-POVMs) [6], thereby obscuring any practical extension to the intrinsically infinite-dimensional CV systems.

In this paper, we apply the shadow-tomography framework to a large family of well-known and well-used CV tomographic protocols. We distill the mathematical tools behind shadow guarantees in such a way that the dependence on strictly DV ingredients is eliminated and we may reformulate established CV protocols in the shadow framework. This reformulation yields accuracy guarantees for expectation values of local observables whose required number of samples scales polynomially with both the number of participating CV modes and the maximum occupation (aka photon) number of each mode.

CV tomography is a long-standing and well-developed field [7]. Focusing on established protocols allows us to boost their credibility, as opposed to developing new protocols that may be equally efficient theoretically but whose practical utility is left as an open question (e.g.,

*gandhari@umd.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

see Sec. VI.A in Ref. [8]). Specifically, our work focuses on tomographic methods that can be easily implemented experimentally with existing quantum optical technology, which underpins fiber-based and free-space quantum communication and key distribution [2]. To illustrate the connection to optics and our commitment to the mantra of classical shadows “measure first, ask questions later” [5], we apply our framework retroactively to experimental quantum-optical data published in 2010 [9]. Naturally, our theoretical guarantees apply equally well if the corresponding protocols are performed in other CV platforms such as microwave cavities coupled to superconducting qubits [10], motional degrees of freedom of trapped ions [11], and optomechanical [12] and nanoscale acoustic [13] resonators.

The first method that we are able to recast in the shadow-tomography framework is homodyne detection [7,14–19]. We show that the number of samples needs to scale at most as the fifth power of the maximum occupation number (up to a logarithmic correction) to yield reliable *homodyne shadow* estimates of a single-mode state. Since we require a finite occupation number cutoff, this bound holds only for portions of states supported on finite-dimensional subspaces of the infinite-dimensional Fock space. Our theoretical guarantees use several important technical bounds proved in earlier work [20]. While there have been studies of the statistical efficiency of homodyne tomography (e.g., Refs. [20–24]), to our knowledge, our bounds are an improvement over previous results as they allow us to analyze the sample complexity for the convergence of the operator norm of the estimated state to its ideal value. These improvements are made possible because of our use of matrix concentration inequalities reviewed in Sec. II (see Refs. [25,26]) that were not used in past work on CV state tomography. A recent paper also considered the statistical efficiency of homodyne tomography using these improved matrix concentration inequalities [27]. Here we present more explicit formulas for the bounds in the case of states with a hard photon number cutoff. Adaptation of our approach to the formulation in Ref. [27] also leads to explicit bounds on their sample complexity that appear to give agreement between the two papers in areas where there is overlap. On a related note, Rosati [19] recently studied the sample complexity of estimating unknown Gaussian or generalized (i.e., linear combinations of) Gaussian CV states, where a risk function is minimized. The distance between the outcomes of the original state and the estimated state, averaged over channels and measurements, was the metric of error. Our work, on the other hand, considers an absolute error via the trace norm of the difference of the original state and the estimated state. Finally, we remark that our results provide upper bounds on the scaling of the number of samples required, i.e., the sample complexity, for estimating the state from random measurements. We leave open the

important problem of obtaining lower bounds on the optimal sample complexity [4], which have been examined, for example, in the case of homodyne tomography [28].

Building on the formalism in Ref. [29], we are also able to recast a large class of protocols using displaced Fock states and make contact with heterodyne, photon-number-resolving (PNR), and photon-parity tomography [7]. For the latter two methods, we show that the number of samples needs to scale at most as the fourth power of the maximum occupation (up to a logarithmic correction) to yield reliable *PNR shadow* estimates. For heterodyne measurements, our upper bound is effectively exponential in the maximum occupation.

Following our theoretical analysis of single-mode systems, we provide a generalization to multimode states. We prove that local CV shadow data can be used to reconstruct k -local reduced density matrices on any subset of the modes with a sample complexity polynomial in the total number of modes and exponential in k .

We support our theoretical analysis with multiple numerical simulations and the analysis of existing experimental homodyne data. We begin by verifying that both shadow methods indeed construct reliable estimates of a target state. We then determine how the number of samples required for both methods scales with the maximum occupation number N . Through simulations involving the reconstruction of the vacuum state, coherent state superposition (CSS) states (i.e., cat states), and random pure states, we observe that the scaling range for both the measurement schemes agree with the theoretical predictions. We also provide a basic demonstration of homodyne tomography on multimode separable and entangled states, finding that the latter require more samples to reach the same precision. Our results are summarized in Table I.

We begin with a general result for estimating density matrix elements of a single-mode CV state using a generic shadow-based protocol in Sec. II. Its application to homodyne measurements and photon-number-resolving measurements is discussed in Sec. III. In Sec. IV, we extend our theory to multimode CV states, and we support our methods by numerical results and analysis of experimental data in Sec. V. We conclude this paper in Sec. VI.

II. SINGLE-MODE SHADOWS: GENERAL CASE

Here we distill the necessary mathematical tools from DV shadow-based accuracy guarantees in a general lemma that we apply to specific CV protocols in later sections. We adapt the slight reformulation from Appendix A in Ref. [30], the original shadow work, in which guarantees are shown exclusively in terms of the system state. This allows us to derive sampling guarantees that are suboptimal relative to the qubit case, but are sufficiently general to be applied to numerous CV tomographic protocols.

TABLE I. Summary of the different protocols that exist, to our knowledge, for studying the sample complexities of CV state tomography. For each protocol, we indicate the measurement basis and a bound on the scaling of sample complexity (with system size N and number of modes M for multimode shadows). The sample complexity of design-based shadows was estimated with use of known variance bounds [4,8]. Additionally, for single-mode measurement methods considered in this paper, we provide the scaling observed in numerical simulations.

Protocol	Basis expansion	Upper bound scaling	Upper bound on scaling from numerical simulations
Single-mode homodyne shadows	Position states/pattern functions	$N^{4+1/3}$	N^4
Single-mode PNR shadows	Displaced Fock states/ T operator	N^4	N^6
Design-based CV tomography [8]	3-design	N^4	$\dots$
Multimode shadows	Local shadows	$N^{4M} - N^{(6+1/3)M}$	$\dots$

We assume loss-free propagation and detection for our theoretical derivations.

A single-mode CV density matrix can be represented in various ways, e.g., using Wigner, Q , P [31,32], or more general [29] representations, moment expansions [33], homodyne measurements [7,16,17], or outer products of occupation number (aka Fock) or other simple basis states. All these ways effectively express the density matrix as a linear combination of some basis set of operators $\{\sigma(\mu)\}$ with corresponding coefficients $p(\mu)$, where μ parameterizes some generally continuous and noncompact space $\mathcal{X}$,

$$\rho = \int_{\mathcal{X}} d\mu p(\mu) \sigma(\mu). \quad (1)$$

Assuming a non-negative and normalized weight function $p(\mu)$ and sufficiently well-behaved space $\mathcal{X}$ allows one to re-express the above formula in a statistical fashion, i.e., as an expectation value over samples σ , sampled from $\mathcal{X}$ according to the probability distribution governed by $p(\mu)$,

$$\rho = \mathbb{E}_{\mu \sim \mathcal{X}} \sigma(\mu) = \mathbb{E}_{\mu} \sigma(\mu). \quad (2)$$

Classical-shadow tomography in our formulation refers to tomographic protocols with the restriction that $p(\mu)$ is a probability distribution over random measurement settings and their outcomes. A general quantum state tomography protocol based solely on Eq. (1) may not have a formulation as a type of classical-shadow tomography. Crucially, though, we see below that there is a natural analysis of homodyne tomography that fits into the classical-shadow framework.

The above formulation allows a simple application of matrix concentration inequalities (loosely speaking, the matrix version of the law of large numbers; see Refs. [25,26]). If we are given a set of T independent samples $T = \{\mu_1, \dots, \mu_T\}$, we can construct an estimator for the state by averaging the basis operators corresponding to the

outcomes in the set,

$$\sigma_T = \frac{1}{T} \sum_{i=1}^T \sigma(\mu_i). \quad (3)$$

We define the above estimator as a *shadow of size T* . Such shadows recover the state by taking their expectation value over all sample sets T ,

$$\mathbb{E}_T \sigma_T = \frac{1}{T} \sum_{i=1}^T \mathbb{E}_{\mu_i} \sigma(\mu_i) = \mathbb{E}_{\mu} \sigma(\mu) = \rho. \quad (4)$$

This convergence property is a necessary starting point for shadow estimation, but it says nothing about *how fast* the convergence happens as the number of samples increases [34].

To bound the rate of convergence of the above estimate, we restrict ρ to finite-dimensional truncated versions ρ^N , with the dimension cutoff N assumed to be sufficiently large so that the truncated operators capture the essence of the original ones. The resulting matrices have nonzero entries only for row and column indices 0 through $N - 1$. A “regularized” shadow is constructed with snapshots projected into the $(N - 1)$ -photon subspace, with use of $P_N = \sum_{i=0}^{N-1} |i\rangle\langle i|$,

$$\sigma_T^N \equiv P_N \sigma_T P_N = \frac{1}{T} \sum_{i=1}^T P_N \sigma(\mu_i) P_N. \quad (5)$$

Our analysis can be readily extended to more general soft cutoffs by replacement of the projectors onto fixed photon number spaces by operators with smooth cutoffs [35]. Note also that we do not renormalize the above regularized density matrices or shadows. Doing so introduces additional factors that would complicate our calculations. Hence, for ease and uniformity, we chose to work with the regularized matrices as they are.

We now bound the difference between the sample average $\sigma_T^{(N)}$ and the actual average $\rho^{(N)} = P_N \rho P_N$, backing out the minimum shadow size T required for an accurate

estimate. We use the trace norm or 1-norm of the difference $\sigma_T^N - \rho^N$ to quantify the error.

Lemma 1. Fix $\epsilon, \delta \in (0, 1)$, and let σ_T^N be an N -dimensional classical shadow of a single-mode continuous-variable quantum state ρ . If the size of the classical shadow is at least T such that

$$T = \frac{2N^2 (v^2 + R\epsilon/2N)}{\epsilon^2} (\ln 2N + \ln 1/\delta), \quad (6)$$

where

$$v^2 = \left\| \mathbb{E}_T (\sigma_T^N)^2 \right\|_\infty = \left\| \mathbb{E}_\mu (\sigma^N(\mu))^2 \right\|_\infty, \quad (7a)$$

$$R \geq \left\| \sigma^N(\mu) - \mathbb{E}_\mu \sigma^N(\mu) \right\|_\infty, \quad (7b)$$

then the probability

$$\mathbb{P}(\|\sigma_T^N - \rho^N\|_\infty \leq \epsilon) \geq 1 - \delta.$$

Proof. Applying the matrix Bernstein inequality (see Ref. [26] and Appendix A in Ref. [30]) to $\sigma_T^N - \rho^N$ and using Eq. (4), we obtain

$$\mathbb{P}(\|\sigma_T^N - \rho^N\|_\infty \geq \tilde{\epsilon}) \leq 2N \exp\left(-\frac{T\tilde{\epsilon}^2/2}{v^2 + R\tilde{\epsilon}/3}\right).$$

Using the equivalence relation $\|X\|_\infty \leq \|X\|_1 \leq N\|X\|_\infty$ for an N -dimensional system, we have $\mathbb{P}(\|X\|_1 \geq \epsilon) \leq \mathbb{P}(N\|X\|_\infty \geq \epsilon)$. Substituting $\tilde{\epsilon}/N \rightarrow \epsilon$, we get

$$\mathbb{P}(\|\sigma_T^N - \rho^N\|_1 \geq \epsilon) \leq 2N \exp\left(-\frac{T\epsilon^2}{2N^2(v^2 + R\epsilon/3N)}\right).$$

Equating the error term on the right-hand side of the inequality to δ and rearranging the resulting expression, we obtain the result (6). ■

We replaced the original variance term $\|\mathbb{E}_\mu (\sigma^N - \rho^N)^2\|_\infty$ in Eq. (7a), which measures the fluctuations about the mean ρ^N , with the term $\|\mathbb{E}_\mu (\sigma^N)^2\|_\infty$ for ease of calculations. We noticed that this leads to the derived bounds being higher than they would otherwise have been in some cases.

We see that the scaling of the required sample size in Eq. (6) depends logarithmically on the inverse error probability δ and polynomially on the accuracy ϵ , recovering the corresponding portion of the qubit shadow result [4]. To determine the dependence on the occupation number cut-off N , we need to know the corresponding scaling of R and v [Eq. (7)], which are sometimes referred to as “shadow norms.” This scaling depends on the details of the protocol, and we proceed to apply this lemma to specific cases.

III. APPLICATIONS

We apply the general shadow bound from Sec. II to a large family of CV tomographic protocols that include homodyne and photon-number-resolving tomography.

A. Homodyne shadows

Homodyne tomography is a popular technique used in quantum optics to make quadrature measurements of CV quantum states [7]. This technique allows one to measure a quadrature operator such as oscillator position, momentum, or any of their superpositions.

Let $a, a^\dagger$ be the canonical mode lowering and raising operators satisfying $[a, a^\dagger] = 1$. A general quadrature operator,

$$X_\theta = \frac{1}{\sqrt{2}}(ae^{-i\theta} + a^\dagger e^{i\theta}), \quad (8)$$

is parameterized by $\theta \in [0, \pi)$, where $\theta = 0$ ($\theta = \pi/2$) corresponds to the pure position (momentum) quadrature. We consider all possible values of θ to be admissible and do not consider the problem of choosing an appropriate discretization. In practice, we use pseudorandom number generators to generate the values of θ_i . Each operator admits its own set of orthogonal but non-normalizable “eigenstates” $|x_\theta\rangle$ with $x_\theta \in \mathbb{R}$. Even though such “eigenstates” are not quantum states but are instead distributions, they do resolve the identity for each θ and thus yield valid positive operator-valued measures (POVMs) [36]. For a given θ , a homodyne protocol yields a measurement in the corresponding set of eigenstates.

To perform a homodyne measurement, the unknown system state interferes with an ancillary mode in a coherent state $|\alpha_{\text{LO}}\rangle$ [coherent state local oscillator (LO)] via a beam splitter, as shown in Fig. 1(a). The phase of the coherent state LO is precisely the phase θ that defines the quadrature that is measured, and is picked uniformly from $[0, \pi)$. Both output amplitudes of the beam splitter, N_1 and N_2 , respectively, are measured via homodyne detectors. Their difference N_- is proportional to the quadrature amplitude of the signal:

$$N_- = \sqrt{2} |\alpha_{\text{LO}}| x_\theta \quad \text{with probability} \quad \langle x_\theta | \rho | x_\theta \rangle. \quad (9)$$

To cast homodyne tomography into shadow terminology, we first specialize the expansion in Eq. (1) to this case. Substituting parameters θ and x_θ for μ , we get the expansion [see Eq. (23) in Ref. [7]]

$$\rho = \int_0^\pi d\theta \int_{\mathbb{R}} dx_\theta \frac{1}{\pi} \langle x_\theta | \rho | x_\theta \rangle F(x_\theta, \theta), \quad (10)$$

where $1/\pi$ corresponds to picking θ uniformly from $[0, \pi)$, and $\langle x_\theta | \rho | x_\theta \rangle$ corresponds to the conditional probability of the homodyne setup yielding x_θ as the outcome. The




FIG. 1. (a) Balanced homodyne detection. The unknown state signal interferes with a local oscillator (a coherent state $|\alpha\rangle$) and the intensities N_1 and N_2 at the output arms are measured. The difference in the intensities is proportional to the quadrature amplitude of the unknown state in the mode defined by the oscillator. (b) PNR measurement. The local oscillator is in a coherent state $|\alpha\rangle$, and detection of N photons at the detector occurs with probability $p(N|\alpha) = \langle N, \alpha | \rho | N, \alpha \rangle$. BS, beam splitter.

matrix elements of the basis operators $\{F(x_\theta, \theta)\}$ have standard expansions in the quantum optics literature. They are defined in terms of what are known as “pattern functions” $f_{mn}(x)$ [7,37]:

$$\langle m | F(x_\theta, \theta) | n \rangle = e^{i(m-n)\theta} f_{mn}(x_\theta). \quad (11)$$

These pattern functions are symmetric in m and n , and are defined with use of Fock state wave functions $\psi_m(x)$ (m th energy eigenstate of a harmonic oscillator) and $\phi_n(x)$ (n th non-normalizable solution of the Schrödinger equation of a harmonic oscillator) as

$$f_{mn}(x) = \frac{\partial}{\partial x} (\psi_m(x) \phi_n(x)) \quad \text{for } n \geq m. \quad (12)$$

Because the quadrature probability distributions are normalized for each θ , the above equation can automatically be interpreted in the statistical fashion of Eq. (2). In this interpretation, basis operators F are sampled according to the probability distribution defined by θ and $\langle x_\theta | \rho | x_\theta \rangle$. A regularized *homodyne CV shadow* of size T , constructed with use of the samples $\{x_{\theta_i}, \theta_i\}_{i=1}^T$, is therefore

$$\sigma_T^N = \frac{1}{T} \sum_{i=1}^T P_N F(x_{\theta_i}, \theta_i) P_N. \quad (13)$$

The randomized choice of θ_i independent being the state is what allows us to reinterpret this reconstruction method as an instance of classical shadows. We also remark that a homodyne shadow is distinct from the Wigner function projection of the state. A shadow estimates the density operator in the Fock basis from a collection of homodyne samples with randomly chosen θ . A Wigner function projection $\int dp_\theta W(x_\theta, p_\theta) = \langle x_\theta | \rho | x_\theta \rangle$, on the other hand, is the probability of obtaining specific measurement outcomes at a fixed value of θ .

We now adapt Lemma 1 to homodyne shadows, relying on the bounds on the pattern-function matrix from Ref. [20]. The following theorem gives an upper bound on the sample complexity of constructing shadows through homodyne samples, for an error $\|\sigma_T^N - \rho^N\|_1 \leq \epsilon$ with high probability (greater than $1 - \delta$).

Theorem 1. With use of homodyne tomography, for $\epsilon, \delta \in (0, 1)$, some positive constant C_1 [20], and an N -dimensional homodyne shadow σ_T^N (13) of a single-mode CV state ρ ,

$$\mathbb{P}(\|\sigma_T^N - \rho^N\|_1 \leq \epsilon) \geq 1 - \delta$$

when the size of the shadow is at least

$$T = \frac{2N^{13/3}C_1}{\epsilon^2} (\ln 1/\delta + \ln 2N).$$

Proof. From Theorem 3.2 in Ref. [20], we use $\|\sigma_T^N\|_\infty^2 \leq C_1 N^{7/3}$ and $\|\sigma_T^N\|_2^2 \leq C_2 N^3$, which is valid for any value of homodyne angle θ and gives

$$\|\sigma_T^N - \rho^N\|_\infty \leq \|\sigma_T^N\|_\infty + 1 \leq \sqrt{C_1} N^{7/6} + 1 = R \quad (14)$$

and

$$v^2 = \|\mathbb{E}_T \sigma_T^2\|_\infty = \left\| \int dx d\theta \frac{1}{\pi} \langle x_\theta | \rho | x_\theta \rangle \sigma_T^2 \right\|_\infty \quad (15a)$$

$$\leq \int dx d\theta \frac{1}{\pi} \langle x_\theta | \rho | x_\theta \rangle \|\sigma_T^2\|_\infty \quad (15b)$$

$$\leq \int dx d\theta \frac{1}{\pi} \langle x_\theta | \rho | x_\theta \rangle C_1 N^{7/3} \quad (15c)$$

$$\leq C_1 N^{7/3}. \quad (15d)$$

Therefore, with use of Lemma 1, the size of the shadow needs to be at least

$$\begin{aligned} T &= \frac{2N^2 (\sigma^2 + R\epsilon/2N)}{\epsilon^2} (\ln 2N - \ln \delta) \\ &\leq \frac{2N^2 (C_1 N^{7/3} + 1 + \sqrt{C_1} N^{1/6} \epsilon/2)}{\epsilon^2} (\ln 2N + \ln 1/\delta). \end{aligned} \quad (16)$$

■

We see that the upper bound on the minimum number of samples scales as $N^{13/3} \ln N$ for a photon-cutoff N .

Regarding the utility of homodyne shadows in calculating expectation values of observables, the pattern-function operators F are unequivocally more complex than the original qubit shadows [4]. Nevertheless, efficient methods are known to compute them using recursive relations [38].

B. PNR shadows

Our second application is based on the T -operator formalism developed in Ref. [29]. The relevant basis expansion of a single-mode CV state ρ can be written in terms of what are called “ T operators,”

$$\rho = \frac{1}{\pi} \int_{\mathbb{C}} d^2\alpha \operatorname{Tr}[\rho T(\alpha, \alpha^*)] \bar{T}(\alpha, \alpha^*), \quad (17)$$

where $\operatorname{Tr}[\rho T(\alpha, \alpha^*)]$ plays the role of a quasiprobability distribution over the complex plane $\mathbb{C}$ and the dual operators $\bar{T}(\alpha, \alpha^*)$ are used to reconstruct the state given this distribution.

The T operators are defined in terms of displaced Fock states, which are the result of applying a displacement operator $D(\alpha) = \exp(\alpha a^\dagger - \alpha^* a)$ on a Fock state: $|\alpha, n\rangle \equiv D(\alpha)|n\rangle$. There exists a family of $T(\alpha, \alpha^*)$, and their corresponding duals $\bar{T}(\alpha, \alpha^*)$, parameterized by $-1 < r < 1$, such that the dual operator is given by

$$\bar{T}_r(\alpha, \alpha^*) = T_{-r}(\alpha, \alpha^*). \quad (18)$$

$T(\alpha, \alpha^*)$ belonging to this family can be diagonalized as

$$\begin{aligned} T_r(\alpha, \alpha^*) &= \sum_{n=0}^{\infty} \lambda_r^{(n)} D(\alpha)|n\rangle\langle n|D(\alpha)^\dagger \\ &= \sum_n \lambda_r^{(n)} |\alpha, n\rangle\langle\alpha, n|, \end{aligned} \quad (19)$$

where the eigenvalues $\lambda_r^{(n)}$ are

$$\lambda_r^{(n)} = \frac{2(-1)^n(1-r)^n}{\pi(1+r)(1+r)^n}. \quad (20)$$

With use of the above expansions, Eq. (17) becomes

$$\rho = \int_{\mathbb{C}} \frac{d^2\alpha}{\pi} \sum_{n=0}^{\infty} \langle n, \alpha | \rho | n, \alpha \rangle \lambda_r^{(n)} \bar{T}_r(\alpha, \alpha^*), \quad (21)$$

which we interpret as an expectation value of samples $\lambda_r^{(n)} \bar{T}_r(\alpha, \alpha^*)$ distributed according to the probabilities $1/\pi \langle n, \alpha | \rho | n, \alpha \rangle$. However, in a physical interferometric measurement procedure corresponding to this expansion [see Fig. 1(b)], the displacement parameter α comes from a noncompact set. This means that we need to first truncate the sample region to make the above expression a valid (i.e., normalizable) probability distribution.

We restrict the integration of Eq. (21) to be over a phase-space region $\mathcal{A}$ of finite area A . We can set this region to be, for example, a disk of radius $\alpha_{\max}$. We will then be able to account for the truncated Fock space by setting $\alpha_{\max}^2 = N$ since that is the average occupation number of a coherent state with $|\alpha| = \alpha_{\max}$. In that case, $A = 4\pi N$ scales linearly with the maximum photon number cutoff.

Projection into an N -dimensional subspace implies that we consider only states with some maximum photon number $N - 1$ and regularized T operators $\bar{T}_r^N(\alpha, \alpha^*)$. Combining this with the truncation of the phase-space integration yields the following approximate expression for the truncated density matrix:

$$\rho^N \approx \int_{\mathcal{A}} \frac{d^2\alpha}{\pi A} \sum_{n=0}^{N-1} \langle n, \alpha | \rho | n, \alpha \rangle A \lambda_r^{(n)} \bar{T}_r^N(\alpha, \alpha^*), \quad (22)$$

in which we interpret $A \lambda_r^{(n)} \bar{T}_r^N(\alpha, \alpha^*)$ as a *PNR shadow* sampled from the probability distribution

$$p(n, \alpha) = \frac{1}{A} \langle n, \alpha | \rho | n, \alpha \rangle \quad (23)$$

with parameters $0 \leq n \leq N - 1$ and $\alpha \in \mathcal{A}$.

The eigenvalues $\lambda_r^{(n)}$ exhibit different behaviors for different values of r , and in the Appendix we analyze how many samples are required for density matrix estimation in each case. We obtain quartic scaling with N at $r = 0$, up to logarithmic corrections, and unfavorable exponential scaling for $0 < |r| < 1$. In the favorable case, the upper bound on the minimum number of samples is given by

$$T = \frac{32N^2 A^2}{\pi^4 \epsilon^2} (\ln 2N + \ln 1/\delta) = \mathcal{O}(N^4 \ln N), \quad (24)$$

beating our theoretical estimates for homodyne shadows. Therefore, we see that T operators at $r = 0$ have a favorable sample complexity scaling and can be used as a tomographic method to estimate an unknown state. The matrix elements of the T operator in the Fock basis are equal to the matrix elements of the displacement operator up to a phase

$$\langle m | T_0(\alpha, \alpha^*) | n \rangle = (1)^n \langle m | D(2\alpha) | n \rangle, \quad (25)$$

which can be used to efficiently compute the classical shadows. Moreover, the case of $r = 0$ can also be realized via photon-parity measurement tomography [39–42], providing another experimental technique to realize these shadows.

It is worth noting that in the limit $r = -1$, $\bar{T}_r(\alpha, \alpha^*) = T_{-r}(\alpha, \alpha^*)$ becomes a projector onto the coherent state $|\alpha\rangle$, making the measurement protocol equivalent to what is known as heterodyne measurement, achieved with a setup similar to the setup used in homodyne measurement [43]. The basis expansion is given in terms of the Q function $Q(\alpha) = 1/\pi \langle \alpha | \rho | \alpha \rangle$. However, the sample complexity upper bound diverges in this case. It is possible to avoid this divergence by taking a limit $r \rightarrow -1$ as the number of samples increases. Unfortunately, this method only achieves an upper bound that is exponential in N because, as noted above, for any $0 < |r| < 1$, our sample complexity upper bound diverges exponentially with N (see the Appendix).

IV. MULTIMODE SHADOWS

In this section, we show that it is possible to construct efficient representations of reduced density matrices on constant numbers of modes from local CV shadow data.

Multimode CV states are states that are made up of multiple *modes* of the underlying system, such as different frequencies of light in an optical system. This is analogous to having a multiqubit state in DVs. Each mode is equipped with a generally continuous and noncompact space $\mathcal{X}_i$ and M -mode, possibly entangled, quantum states reside in $\mathcal{X} = \bigotimes_{i=1}^M \mathcal{X}_i$. Let the space $\mathcal{X}$ be parameterized by a variable $\mu = (\mu^1, \dots, \mu^M)$ and spanned by the set of basis operators $\{\sigma(\mu)\}$. Extending Eq. (1), we can then write the density matrix of a general multimode state belonging to this space as

$$\rho = \int_{\mathcal{X}} d\mu p(\mu) \sigma(\mu), \quad (26)$$

where $p(\mu)$ are the corresponding coefficients. Boldface symbols are used to represent multimode variables and states. All the density operators and shadows considered in this section are the regularized versions, projected into the N -photon subspace of each mode. The superscripts indicating this are omitted for brevity.

Analogous to the single-mode case in Sec. II, we require a basis expansion of ρ such that $p(\mu) \geq 0$ for all ρ , the non-negativity of $p(\mu)$ allowing it to be interpreted as a probability distribution over μ . We proceed by using an expansion of the form $p(\mu) = \langle \mu | \rho | \mu \rangle$, where the quantum states or distributions $|\mu\rangle$ parameterize the underlying space. For example, we can take $|\mu\rangle$ to be local position state eigenstates in the case of homodyne tomography.

We construct shadows of multimode states using local, joint single-mode measurement. For an estimator of a state ρ to qualify as a shadow, as introduced in Sec. II, it must be equal to the state ρ in expectation. Below we show that it can indeed be satisfied for shadows constructed through local measurements. For this calculation, we use another expansion of a general quantum state ρ in terms of some set of separable operators $\{\rho_\alpha\}$ and corresponding complex coefficients $\{c_\alpha\}$,

$$\rho = \sum_{\alpha} c_{\alpha} \rho_{\alpha} = \sum_{\alpha} c_{\alpha} \bigotimes_{i=1}^M \rho_{\alpha}^i. \quad (27)$$

Such an expansion is possible irrespective of whether the state is separable or entangled. The estimator becomes

$$\mathbb{E}_{\mu} \sigma(\mu) = \int_{\mathcal{X}} d\mu \times \langle \mu | \rho | \mu \rangle \times \bigotimes_{i=1}^M \sigma(\mu^i) \quad (28a)$$

$$= \sum_{\alpha} c_{\alpha} \int_{\mathcal{X}} d\mu \times \bigotimes_{i=1}^M \langle \mu^i | \rho_{\alpha}^i | \mu^i \rangle \sigma(\mu^i) \quad (28b)$$

$$= \sum_{\alpha} c_{\alpha} \bigotimes_{i=1}^M \int_{\mathcal{X}_i} d\mu^i \langle \mu^i | \rho_{\alpha}^i | \mu^i \rangle \sigma(\mu^i) \quad (28c)$$

$$= \sum_{\alpha} c_{\alpha} \rho_{\alpha} = \rho. \quad (28d)$$

In the following theorem, we calculate the variance and thereby the sample complexity of estimating a multimode state through local CV shadows. This result reduces to the single-mode result from Lemma 1 when $M = 1$.

Theorem 2. Fix $\epsilon, \delta \in (0, 1)$, and let σ_T be an N^M -dimensional classical shadow obtained through local measurements $\sigma(\mu) = \bigotimes_{i=1}^M \sigma(\mu^i)$ of an M -mode CV state ρ [Eq. (27)]. If the size of the classical shadow is at least

$$T = \frac{2N^{2M} (v_1^{2M} \sum_{\alpha} |c_{\alpha}| + R\epsilon/3N^M)}{\epsilon^2} (M \ln 2N + \ln 1/\delta), \quad (29)$$

then

$$\mathbb{P}(\|\sigma_T - \rho\|_1 \leq \epsilon) \geq 1 - \delta,$$

written in terms of the single-mode variance v_1^2 [Eq. (7a)] and $R \geq \|\sigma(\mu) - \mathbb{E}_{\mu} \sigma(\mu)\|_{\infty}$.

Proof. The measurement variance for an M -mode state proceeds as in the single-mode case with the help of Eq. (27):

$$v_M^2 = \left\| \mathbb{E}_{\mu} (\sigma(\mu))^2 \right\|_{\infty} \quad (30a)$$

$$= \left\| \int_{\mathcal{X}} \prod_{j=1}^M d\mu^j \times \langle \mu | \rho | \mu \rangle \times \bigotimes_{i=1}^M (\sigma(\mu^i))^2 \right\|_{\infty} \quad (30b)$$

$$= \left\| \sum_{\alpha} c_{\alpha} \bigotimes_i \int_{\mathcal{X}_i} d\mu^i \langle \mu^i | \rho_{\alpha}^i | \mu^i \rangle (\sigma(\mu^i))^2 \right\|_{\infty} \quad (30c)$$

$$\leq \sum_{\alpha} |c_{\alpha}| \prod_i \left\| \int_{\mathcal{X}_i} d\mu^i \langle \mu^i | \rho_{\alpha}^i | \mu^i \rangle (\sigma(\mu^i))^2 \right\|_{\infty} \quad (30d)$$

$$\leq \sum_{\alpha} |c_{\alpha}| \prod_i v_1^2 \quad (30e)$$

$$\leq v_1^{2M} \sum_{\alpha} |c_{\alpha}|. \quad (30f)$$

Using the matrix Bernstein inequality and following a calculation similar to that in Lemma 1, replacing N with N^M , we get the desired result in Eq. (29). ■

As a consequence, we obtain a sample complexity bound on estimating k -local reduced density matrices, which are the outcome of tracing over all but k modes of a multimode state. Since we are tracing out some of the modes, we are left with a “reduced density matrix.”

Corollary 1. If ρ_k denotes the reduced density matrix of ρ in $k \leq M$ modes and its estimator constructed through local measurements is $\sigma_{T,k}$ such that the size of the shadow is at least

$$T = \frac{2N^{2k}(v_1^{2k} \sum_{\alpha} |c_{\alpha}| + R\epsilon/3N^k)}{\epsilon^2} (k \ln 2N + \ln 1/\delta),$$

then

$$\mathbb{P}(\|\sigma_{T,k} - \rho_k\|_1 \leq \epsilon) \geq 1 - \delta.$$

Proof. Without loss of generality, let us assume ρ_k is the reduced density matrix over the modes $1, \dots, k$, with $k \leq M$, i.e., $\rho_k = \text{Tr}_{k+1, \dots, M}(\rho)$. Define $\mu_k = (\mu^1, \dots, \mu^k)$ to parameterize this k -mode subspace, and the local operators

$$\sigma_k(\mu_k) = \text{Tr}_{k+1, \dots, M}(\sigma(\mu)) = \bigotimes_{i=1}^k \sigma(\mu^i).$$

The classical shadow of the reduced matrix $\sigma_{T,k}$ is then the average of T such operators. The variance of estimating this reduced matrix, following earlier calculations, is given by

$$v_k^2 = \left\| \mathbb{E}_{\mu_k} (\sigma_k(\mu_k))^2 \right\|_{\infty} \leq v_1^{2k} \sum_{\alpha} |c_{\alpha}|. \quad (31)$$

■

One can further generalize this result to derive sample complexity bounds for determining the reduced density matrix for any choice of k -local modes. The proof follows from an application of the union bound and increases the sample complexity by a factor of k [4,30].

The sum of the absolute values of the coefficients in Eq. (27) is a factor that affects the sampling complexity of multimode states. These coefficients can, in general, be complex. When the state is separable (i.e., can be written as a product state) this factor becomes 1, thereby resulting in a complexity that is equivalent to sampling all the modes individually, as expected of a product state. When the coefficients are all positive, they again sum up to 1 (since the trace of the state must equal 1), resulting in no increase in complexity as compared with a separable state. In the general case, we can still upper-bound the sampling complexity, but the overall scaling with the number of modes has a larger prefactor in the exponent. These results are summarized in Table I.

We add that the proposed local-joint measurements are not the most general measurements for multimode CV states. On trying many different bases of entangled/global states, we have not found a suitable one that fulfills our criteria of a positive weight function for all states [$p(\mu)$ in Eq. (26)], but have not ruled out the possibility of their existence. Studies of whether global measurements can facilitate such a statistical interpretation and of the sample complexity when such bases are used are interesting avenues for future work.

Corollary 2. Similarly, we can estimate the expectation values for a list of multimode observables $\{O_1, \dots, O_M\}$, local to $\{k_1, \dots, k_M\}$ modes, respectively, and each mode within an N -photon subspace, from a shadow of size upper-bounded by $T = \mathcal{O}\left(\max_i \frac{\|O_i\|}{\epsilon^2} \ln(M/\delta)\right)$ such that

$$\mathbb{P}(|\text{Tr}(O_i \rho) - \text{Tr}(O_i \sigma_T)| \leq \epsilon) \geq 1 - \delta \text{ for all } 1 \leq i \leq M. \quad (32)$$

Proof. Since the 1-norm of a matrix is the sum of its singular values, with use of Hölder’s inequality,

$$|\text{Tr}(O \rho)| \leq \|O\|_{\infty} \|\rho\|_1. \quad (33)$$

Therefore,

$$\begin{aligned} \mathbb{P}(|\text{Tr}(O_i(\rho - \sigma_T))| \geq \epsilon) &\leq \mathbb{P}(\|O\|_{\infty} \|\rho\|_1 \geq \epsilon) \\ &= \mathbb{P}\left(\|\rho\|_{\infty} \geq \frac{\epsilon}{\|O\|_1}\right). \end{aligned}$$

With use of the matrix Bernstein inequality, this implies

$$\begin{aligned} \mathbb{P}(|\text{Tr}(O_i(\rho - \sigma_T))| \geq \epsilon) \\ \leq 2N^{k_i} \exp\left(-\frac{T(\epsilon/\|O\|_{\infty})^2}{2N^{2k_i}(v^{2k_i} + R\epsilon/3N^{k_i}\|O\|_{\infty})}\right), \end{aligned}$$

where v and R are as defined in Theorem 2. Equating the upper bound on probability to δ/M and using the union bound [4], we find that a shadow size upper-bounded by

$$T = \max_i \frac{2\|O_i\|_{\infty} N^{2k_i} v^{2k_i}}{\epsilon^2} (k_i \ln 2N + \ln M/\delta) \quad (34)$$

suffices to estimate the list of multimode observables $\{O_i\}_{i=1}^M$ with the desired accuracy. ■

V. NUMERICAL AND EXPERIMENTAL TESTS

In this section, we verify our shadow construction methods and present numerical results on sample complexity scaling, with the maximum occupation number N , of single-mode shadows. We also analyze experimental homodyne data and demonstrate the reconstruction of multimode states.




FIG. 2. Single-mode shadows. (a) Density matrix elements and (b) Wigner function representation of the original state, a homodyne shadow, and a PNR shadow of an even CSS state $|\psi\rangle \propto |\alpha\rangle + |-\alpha\rangle$ with $\alpha = \sqrt{10}$. (c),(d) Numerical sample complexity of single-mode shadows; the precision and confidence parameters are $\epsilon = 0.1$ and $\delta = 0.1$. (c) Homodyne shadows: vacuum state ($T \propto N^{1.389}$), CSS state with amplitude $\alpha = \sqrt{10}$ ($T \propto N^{1.820}$), and random pure states ($T \propto N^{1.212}$). (d) PNR shadows: vacuum state ($T \propto N^{3.51}$) and pure random states.

A. Single-mode simulations

To analyze the applicability of our theoretical guarantees, we construct estimators of density matrices of certain target states using homodyne and PNR shadows. The measurement bases in all cases were generated using uniform pseudorandom-number generators. We then analyzed how close they are to the target states in spectral norm, for ease of simulation. Figure 2(a) shows how a reconstructed CSS state with mean photon number of 10 compares with the original density matrix elements for both homodyne and PNR tomography. A comparison in terms of the Wigner function is shown in Fig. 2(b). To achieve equally accurate shadows such that $\|\sigma_T^N - \rho^N\|_\infty = 0.1$, we had to use 5×10^4 samples for homodyne tomography and 10^6 samples for PNR tomography.

Of primary interest in the classical-shadow formalism is the number of samples required to achieve a given target precision with high probability. This minimum shadow size for a given state T was numerically obtained by our repeating the estimation process for several shadow sizes and identifying the one that results in precision ϵ with a probability of at least $1 - \delta$. We observed that if we are looking for a fixed error in the norm (i.e., fixing ϵ), increasing the shadow size decreases the probability with which this error occurs (i.e., δ) as the estimation becomes more accurate. We observed that T for fixed ϵ is proportional to $\ln 1/\delta$ and this helped us narrow down our search for minimum shadow sizes corresponding to our chosen pair $(\epsilon, \delta) = (0.1, 0.1)$. The scaling with the occupation number cutoff is then obtained by our finding the minimum shadow size for different values of N , and the error bars indicate the statistical errors in estimating T via this process. We again note that in our numerics we analyze wrt spectral norm distance. The scaling we desire, wrt trace distance, can however be easily extrapolated from these simulations using Lemma 1 and $\|X\|_\infty \leq \|X\|_1$. The desired scaling is greater than the obtained scaling utmost by a factor of N^2 . In Figs. 2(c) and 2(d), we plot this observed minimum

size for different states for homodyne and PNR shadows, respectively.

We ran our simulations for a vacuum state $|\psi\rangle = |0\rangle$, a CSS state $|\psi\rangle \propto |\alpha\rangle + |-\alpha\rangle$ with mean photon number $|\alpha|^2 = 10$, and random pure states. The random pure states of size N had nonzero support up to only the $(N - 1)$ th Fock state. Homodyne tomography results are slightly better than, but close to, the theoretical upper bound. PNR tomography for vacuum states also agrees with our theoretical predictions. We also tested the scaling of PNR tomography of random pure states, but were computationally limited in the accessible sizes. Over the range we observed, the results are consistent with the scaling of the vacuum state.

B. Experimental data

Here we analyze the sample complexity of experimental homodyne data from Ref. [9]. Once again, our analysis uses spectral norm distance, the outcomes of which are extrapolated to the case of trace norm distance. The experiment consisted of creating CSSs, where the created states were approximations to small-amplitude, even CSS states $|\psi\rangle \propto |\alpha\rangle + |-\alpha\rangle$. A squeezed vacuum was generated through spontaneous parametric down-conversion in a KNbO₃ nonlinear crystal: an up-converted laser pulse at 430 nm created the squeezed vacuum at 860 nm. After spectral filtering, the squeezed vacuum was incident on a weakly reflecting beam splitter with reflectivity R . Photons detected in the reflected path herald a CSS emerging from the transmission port of the beam splitter. The resulting CSS was then directed to a homodyne detection setup. The homodyne setup consisted of a 50:50 beam splitter, two high-efficiency photodiodes, and a low-noise amplification circuit. The local oscillator was created with a strong laser field (approximately 10^9 photons per pulse) that allows access to the quadratures through balanced detection. The




FIG. 3. Experimental test of single-mode shadows. (a) Density matrix elements and (b) Wigner function representation of the ideal CSS state from theory (left) and the reconstructed state from experimental homodyne measurements (right). Ideal CSS state: $|\psi\rangle \propto |\alpha\rangle + |-\alpha\rangle$, $\alpha = 1.32$, $T = 3 \times 10^5$. (c) Sample complexity scaling with maximum photon number N , from the homodyne data from the experiment. The two methods “TES 2” and “APD 1” correspond to photon detection through a TES and a SPAD, respectively. The scaling shown is for $\epsilon = 0.4$ and $\delta = 0.1$. (d),(e) Homodyne two-mode shadows. Density matrix elements of the original states and numerically reconstructed states. The number of samples $T = 5 \times 10^4$, $N = 2$, and $\alpha = \sqrt{1.5}$ in both cases. (d) Two-mode separable state: $|\Psi\rangle = |\psi_\alpha\rangle \otimes |\psi_\alpha\rangle$, where $|\psi_\alpha\rangle \propto |\alpha\rangle + |-\alpha\rangle$, and the infinity norm distance of the reconstructed state from the ideal state is $\epsilon = 0.03$. (e) Two-mode entangled state: $|\Psi\rangle \propto |\alpha, \alpha\rangle + |-\alpha, -\alpha\rangle$ and $\epsilon = 0.05$.

heralded photon detection was done with either two single-photon avalanche diodes (SPADs) [44] or one transition-edge sensor (TES) [45]. The experiment analyzed here used both heralded-photon detection schemes. The CSSs were heralded on the detection of two heralding photons, generating the approximation of a small-amplitude even cat state.

Using the homodyne data from this experiment, we constructed an estimate of the state, an example of which is shown in Figs. 3(a) and 3(b), and studied how close the reconstructed states are to the target states. Before analyzing the sample complexity, we preprocessed the experimental data to account for propagation losses and detection inefficiencies [9].

We then performed the same numerical tests as in Sec. V A to find the minimum shadow size (T) required to achieve a fixed (ϵ, δ) for different occupational number (N) cutoffs. Working with a given dataset means our shadow size is limited to the size of the dataset and that accuracy cannot be increased arbitrarily. Figure 3(c) shows how T changed with N , and we observe a scaling of $T \propto N^{0.659 \pm 0.001}$ for a CSS with $\alpha = 1.16$ using the TES, and as $T \propto N^{0.685 \pm 0.001}$ for a CSS with $\alpha = 1.32$ using the SPADs. Adjusted to trace-norm distance, the scaling of $T \propto N^{2.6}$ agrees with the upper bounds in our theoretical predictions and numerical simulations for single-mode homodyne shadows.

C. Multimode simulations

We now provide an example reconstruction of both separable and entangled two-mode states.

In Fig. 3(d) we can see original and reconstructed density matrices of a separable two-mode state $|\Psi\rangle = |\psi_\alpha\rangle \otimes |\psi_\alpha\rangle$, where $|\psi_\alpha\rangle$ is an even CSS state of amplitude $\alpha = \sqrt{1.5}$. In Fig. 3(e), we show the reconstruction of an entangled state $|\Psi\rangle = |\alpha\rangle \otimes |\alpha\rangle + |-\alpha\rangle \otimes |-\alpha\rangle$, with the same number of measurements. The infinity-norm distance ϵ between the original state and the reconstructed state is 0.03 and 0.05, respectively. The increased sample complexity needed to reach a given precision for the entangled state tomography is consistent with our general expectations from the theoretical analysis in Sec. IV.

VI. OUTLOOK AND CONCLUSIONS

We developed a formalism to derive worst-case sample complexity of a large class of state-reconstruction procedures for single-mode CV states (Lemma 1 and Theorem 1). In an experimental context, our work provides a tool to compare the sample complexity of different measurement methods. For photon-number-resolving measurements, we show through numerical simulations that our analytical derived upper bound is close to the observed complexity. In the case of homodyne tomography, we notice that the simulation results surpass our derived bound, rendering homodyne tomography more efficient than PNR tomography and photon-parity tomography in practice.

We obtain the sample complexity of estimating multimode CV states by reconstructing shadows of the true states from local measurements on each mode, in the spirit of the qubit-based shadow tomography proposal [3–5]. We also consider a CV manifestation of a global version of the

original shadow protocol, where one uses global measurements over various linear combinations of the constituent modes for state reconstruction. Of particular note is that in contrast to the discrete-variable systems, our analysis does not rely on the construction of state designs (approximate or otherwise) such as in Sec. VI.A in Ref. [8]. Instead, we focus on adapting existing techniques for CV tomography to fit within the framework of randomized measurements.

In summary, we recast existing homodyne and photon-number-resolving protocols as shadow-tomography protocols and show that such protocols yield good local estimates of a multimode state using a number of samples that is polynomial in the number of modes 1. Our CV shadow framework can be extended to analyze other metrics such as the total variation distance between two outcome probability distributions, which is valuable for verifying sampling experiments such as boson sampling [46]. Another remaining open question is to determine the robustness of CV protocols to noise, e.g., within the framework of robust shadow estimation [47]. Deriving rigorous lower bounds that can match the numerically and experimentally observed scaling for homodyne tomography is also an important outstanding challenge.

ACKNOWLEDGMENTS

We thank Yi-Kai Liu and Dominik Hangleiter for helpful discussions. M.J.G. and V.V.A. thank Joe Iosue and Kunal Sharma for discussions, as well as collaborations on Ref. [8]. We thank S. Becker, N. Datta, L. Lami, and C. Rouzé for helpful discussions about the relations of our results to the results in Ref. [27] and for pointing out a change in our proof of homodyne bounds that allows an improvement from $\mathcal{O}(N^5)$ scaling to $\mathcal{O}(N^{4+1/3})$ scaling. This work was supported in part by NIST Grant No. 70NANB21H055_0 and NSF Quantum Leap Challenge Institutes program Grant No. OMA-2120757.

APPENDIX: SCALING WITH PNR TOMOGRAPHY

Beginning with the approximate expression for a truncated density matrix in the displaced Fock basis (A1), we want to estimate the sample complexity of PNR shadows:

$$\rho^N \approx \int_{\mathcal{A}} \frac{d^2\alpha}{\pi A} \sum_{n=0}^{N-1} \langle n, \alpha | \rho | n, \alpha \rangle A \lambda_r^{(n)} \tilde{T}_r^N(\alpha, \alpha^*). \quad (\text{A1})$$

The above expansion in terms of $T_r(\alpha, \alpha^*)$ is valid for $-1 < r < 1$ as mentioned in Sec. III B. Depending on whether r is positive or negative, one of $\lambda_r^{(n)}$ or $\lambda_{-r}^{(n)}$ (the latter being an eigenvalue of $T_{-r}(\alpha, \alpha^*)$) increases with n and the other decreases with n . For $r = 0$, the eigenvalues $\lambda_0^{(n)}$ are independent of n . With use of the triangle inequality

and Eq. (20), the shadow norm bound (7b) becomes

$$\|A \lambda_r^{(n)} T_{-r}^N - \rho^N\|_{\infty} \leq A \|\lambda_r^{(n)} T_{-r}^N\|_{\infty} + A \quad (\text{A2a})$$

$$\leq A |\lambda_{-|r|}^{(N)} \lambda_{|r|}^{(0)}| + A \quad (\text{A2b})$$

$$\leq \frac{2A |\lambda_{-|r|}^{(N)}|}{\pi(1 + |r|)} + A \equiv R, \quad (\text{A2c})$$

with the equality holding at $r = 0$. As for the variance from Eq. (7a), we have

$$v^2 = \left\| \mathbb{E}_{n,\alpha} (A \lambda_r^{(n)} T_{-r}^N)^2 \right\|_{\infty} \quad (\text{A3a})$$

$$= \left\| \int_{\mathcal{A}} \frac{d^2\alpha}{\pi A} \sum_{n=0}^{N-1} p(n, \alpha) (A \lambda_r^{(n)} T_{-r}^N)^2 \right\|_{\infty} \quad (\text{A3b})$$

$$\leq \int_{\mathcal{A}} \frac{d^2\alpha}{\pi A} \sum_{n=0}^{N-1} p(n, \alpha) A^2 \left\| (\lambda_r^{(n)} T_{-r}^N)^2 \right\|_{\infty} \quad (\text{A3c})$$

$$\leq \int_{\mathcal{A}} \frac{d^2\alpha}{\pi A} \sum_{n=0}^{N-1} p(n, \alpha) A^2 |\lambda_{|r|}^{(0)} \lambda_{-|r|}^{(N)}|^2 \quad (\text{A3d})$$

$$\leq A^2 |\lambda_{|r|}^{(0)} \lambda_{-|r|}^{(N)}|^2. \quad (\text{A3e})$$

The remaining norm can be bounded with use of Eq. (20) as was done in Eq. (A2), and the remaining sum and integral can be upper-bounded by a factor of A , yielding

$$v^2 \leq \frac{16 A^2}{\pi^4 (1 - |r|^2)^2} \left(\frac{1 + |r|}{1 - |r|} \right)^{2N}. \quad (\text{A4})$$

Applying Lemma 1 and using $A = aN$, we find that the upper bound on the minimum number of samples is given by

$$T = \frac{32 a^2 N^4}{\pi^4 \epsilon^2 (1 - |r|^2)^2} \left(\frac{1 + |r|}{1 - |r|} \right)^{2N} \ln \left(\frac{2N}{\delta} \right), \quad (\text{A5})$$

growing exponentially with N for $r \neq 0$.

-
- [1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, 2011).
 - [2] S. L. Braunstein and P. van Loock, Quantum information with continuous variables, *Rev. Mod. Phys.* **77**, 513 (2005).
 - [3] S. Aaronson, *Shadow tomography of quantum states*, [arXiv:1711.01053](https://arxiv.org/abs/1711.01053) (2018).
 - [4] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
 - [5] A. Elben, S. T. Flammia, H.-Y. Huang, R. Kueng, J. Preskill, B. Vermersch, and P. Zoller, *The randomized measurement toolbox*, [arXiv:2203.11374](https://arxiv.org/abs/2203.11374) (2022).

- [6] R. Stricker, M. Meth, L. Postler, C. Edmunds, C. Ferrie, R. Blatt, P. Schindler, T. Monz, R. Kueng, and M. Ringbauer, *Experimental single-setting quantum state tomography*, [arXiv:2206.00019](#) (2022).
- [7] A. I. Lvovsky and M. G. Raymer, Continuous-variable optical quantum-state tomography, *Rev. Mod. Phys.* **81**, 299 (2009).
- [8] J. T. Iosue, K. Sharma, M. J. Gullans, and V. V. Albert, *Continuous-variable quantum state designs: Theory and applications*, [arXiv:2211.05127](#) (2022).
- [9] T. Gerrits, S. Glancy, T. S. Clement, B. Calkins, A. E. Lita, A. J. Miller, A. L. Migdall, S. W. Nam, R. P. Mirin, and E. Knill, Generation of optical coherent-state superpositions by number-resolved photon subtraction from the squeezed vacuum, *Phys. Rev. A* **82**, 031802(R) (2010).
- [10] S. M. Girvin, Circuit QED: Superconducting qubits coupled to microwave photons, *Quantum Machines: Measurement and Control of Engineered Quantum Systems* (2014), p. 113.
- [11] B. de Neeve, T.-L. Nguyen, T. Behrle, and J. P. Home, Error correction of a logical grid state qubit by dissipative pumping, *Nat. Phys.* **18**, 296 (2022).
- [12] M. Aspelmeyer, T. J. Kippenberg, and F. Marquardt, Cavity optomechanics, *Rev. Mod. Phys.* **86**, 1391 (2014).
- [13] G. S. MacCabe, H. Ren, J. Luo, J. D. Cohen, H. Zhou, A. Sipahigil, M. Mirhosseini, and O. Painter, Nano-acoustic resonator with ultralong phonon lifetime, *Science* **370**, 840 (2020).
- [14] K. Vogel and H. Risken, Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase, *Phys. Rev. A* **40**, 2847 (1989).
- [15] D. T. Smithey, M. Beck, M. G. Raymer, and A. Faridani, Measurement of the Wigner distribution and the density matrix of a light mode using optical homodyne tomography: Application to squeezed states and the vacuum, *Phys. Rev. Lett.* **70**, 1244 (1993).
- [16] U. Leonhardt and H. Paul, Measuring the quantum state of light, *Prog. Quantum Electron.* **19**, 89 (1995).
- [17] A. Bisio, G. Chiribella, G. M. D'Ariano, S. Facchini, and P. Perinotti, Optimal quantum tomography, *IEEE J. Sel. Top. Quantum Electron.* **15**, 1646 (2009).
- [18] G. M. D'Ariano, Tomographic measurement of the density matrix of the radiation field, *Quantum Semiclass. Opt.: J. Eur. Opt. Soc. Part B* **7**, 693 (1995).
- [19] M. Rosati, *A learning theory for quantum photonic processes and beyond*, [arXiv:2209.03075](#) (2022).
- [20] L. M. Artiles, R. D. Gill, and M. I. Guta, *An invitation to quantum tomography*, [arXiv:math/0405595](#) **67**, 109 (2005).
- [21] K. Banaszek, Maximum-likelihood estimation of photon-number distribution from homodyne statistics, *Phys. Rev. A* **57**, 5013 (1998).
- [22] R. Gill and M. Guta, *An invitation to quantum tomography*, [arXiv:quant-ph/0303020](#) (2003).
- [23] P. Albin, E. De Vito, and A. Toigo, Quantum homodyne tomography as an informationally complete positive-operator-valued measure, *J. Phys. A: Math. Theor.* **42**, 295302 (2009).
- [24] P. Alquier, K. Meziani, and G. Peyré, Adaptive estimation of the density matrix in quantum homodyne tomography with noisy data, *Inverse Probl.* **29**, 075017 (2013).
- [25] J. A. Tropp, The expected norm of a sum of independent random matrices: An Elementary Approach. In: Houdré, C., Mason, D., Reynaud-Bouret, P., Rosiński, J. (eds) *High Dimensional Probability VII. Progress in Probability*, vol 71. Birkhäuser, (2016), pp. 173–202.
- [26] J. A. Tropp, An introduction to matrix concentration inequalities, *Found. Trends Mach. Learn.* **8**, 1 (2015).
- [27] S. Becker, N. Datta, L. Lami, and C. Rouzé, *Classical shadow tomography for continuous variables quantum systems*, [arXiv:2211.07578](#) (2022).
- [28] U. Leonhardt and M. Munroe, Number of phases required to determine a quantum state in optical homodyne tomography, *Phys. Rev. A* **54**, 3682 (1996).
- [29] A. Wunsche, Displaced Fock states and their connection to quasiprobabilities, *Quantum Opt.: J. Eur. Opt. Soc. Part B* **3**, 359 (1991).
- [30] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, Provably efficient machine learning for quantum many-body problems, *Science* **377**, 1397 (2022).
- [31] M. Hillery, R. O'Connell, M. Scully, and E. Wigner, Distribution functions in physics: Fundamentals, *Phys. Rep.* **106**, 121 (1984).
- [32] C. Ferrie, Quasi-probability representations of quantum theory with applications to quantum information science, *Rep. Prog. Phys.* **74**, 116001 (2011).
- [33] A. Wünsche, Tomographic reconstruction of the density operator from its normally ordered moments, *Phys. Rev. A* **54**, 5291 (1996).
- [34] Note, the words shadow, estimator and reconstructed state are used interchangeably throughout the paper.
- [35] For example, for a regulator R one can find a sufficiently large N such that $\|P_N R \rho R^\dagger P_N - R \rho R^\dagger\| < \epsilon$ for all density matrices ρ and apply our results directly on the basis of the triangle inequality. More generally, the bounds we prove on the projected shadows could be adapted to regulators with smooth cutoffs, but we leave this approach for future work.
- [36] A. Holevo, *Probabilistic and Statistical Aspects of Quantum Theory* (Edizioni della Normale, Pisa, 2011).
- [37] G. M. D'Ariano, C. Macchiavello, and M. G. A. Paris, Detection of the density matrix through optical homodyne tomography without filtered back projection, *Phys. Rev. A* **50**, 4298 (1994).
- [38] T. Richter, Realistic pattern functions for optical homodyne tomography and determination of specific expectation values, *Phys. Rev. A* **61**, 063819 (2000).
- [39] L. G. Lutterbach and L. Davidovich, Method for direct measurement of the Wigner function in cavity QED and ion traps, *Phys. Rev. Lett.* **78**, 2547 (1997).
- [40] P. Bertet, A. Auffeves, P. Maioli, S. Osnaghi, T. Meunier, M. Brune, J. M. Raimond, and S. Haroche, Direct measurement of the Wigner function of a one-photon Fock state in a cavity, *Phys. Rev. Lett.* **89**, 200402 (2002).
- [41] B. Vlastakis, G. Kirchmair, Z. Leghtas, S. E. Nigg, L. Frunzio, S. M. Girvin, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, Deterministically encoding quantum information using 100-photon Schrödinger cat states, *Science* **342**, 607 (2013).
- [42] C. Wang, Y. Y. Gao, P. Reinhold, R. W. Heeres, N. Ofek, K. Chou, C. Axline, M. Reagor, J. Blumoff, K. M. Sliwa,

- L. Frunzio, S. M. Girvin, Liang Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, A Schrödinger cat living in two boxes, *Science* **352**, 1087 (2016).
- [43] A. Cives-Esclop, A. Luis, and L. L. Sánchez-Soto, An eight-port detector with a local oscillator of finite intensity, *J. Opt. B: Quantum Semiclass. Opt.* **2**, 526 (2000).
- [44] S. Cova, G. Ripamonti, and A. Lacaita, Avalanche semiconductor detector for single optical photons with a time resolution of 60 ps, *Nucl. Instrum. Methods Phys. Res. Sect. A: Accelerators, Spectrometers, Detectors and Associated Equipment* **253**, 482 (1987).
- [45] A. E. Lita, A. J. Miller, and S. W. Nam, Counting near-infrared single-photons with 95% efficiency, *Opt. Express* **16**, 3032 (2008).
- [46] D. Hangleiter and J. Eisert, *Computational advantage of quantum random sampling*, [arXiv:2206.04079](https://arxiv.org/abs/2206.04079) (2022).
- [47] S. Chen, W. Yu, P. Zeng, and S. T. Flammia, Robust shadow estimation, *PRX Quantum* **2**, 030348 (2021).