





Improving application performance with biased distributions of quantum states

Sanjaya Lohani ^{1,2,*} Joseph M. Lukens ^{3,*} Daniel E. Jones ⁴ Thomas A. Searles ^{1,5}
Ryan T. Glasser^{2,†} and Brian T. Kirby^{2,4,‡}

¹IBM-HBCU Quantum Center, Howard University, Washington, D.C. 20059, USA

²Tulane University, New Orleans, Louisiana 70118, USA

³Quantum Information Science Group, Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831, USA

⁴United States Army Research Laboratory, Adelphi, Maryland 20783, USA

⁵Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA



(Received 8 August 2021; accepted 3 November 2021; published 29 November 2021)

We consider the properties of a specific distribution of mixed quantum states of arbitrary dimension that can be biased towards a specific mean purity. In particular, we analyze mixtures of Haar-random pure states with Dirichlet-distributed coefficients. We analytically derive the concentration parameters required to match the mean purity of the Bures and Hilbert–Schmidt distributions in any dimension. Numerical simulations suggest that this value recovers the Hilbert–Schmidt distribution exactly, offering an alternative and intuitive physical interpretation for ensembles of Hilbert–Schmidt-distributed random quantum states. We then demonstrate how substituting these Dirichlet-weighted Haar mixtures in place of the Bures and Hilbert–Schmidt distributions results in measurable performance advantages in machine-learning-based quantum state tomography systems and Bayesian quantum state reconstruction. Finally, we experimentally characterize the distribution of quantum states generated by both a cloud-accessed IBM quantum computer and an in-house source of polarization-entangled photons. In each case, our method can more closely match the underlying distribution than either Bures or Hilbert–Schmidt distributed states for various experimental conditions.

DOI: [10.1103/PhysRevResearch.3.043145](https://doi.org/10.1103/PhysRevResearch.3.043145)

I. INTRODUCTION

Ensembles of random density matrices have found broad applicability in quantum information science [1–7]. Of particular recent interest is their use in quantum state tomography systems, either as training sets for machine-learning-based techniques [8–12] or as prior distributions for Bayesian state reconstruction [13–19]. Significant effort has been devoted to developing such ensembles based on various underlying measures and determining their characteristics [20–24]. However, as quantum information science becomes increasingly reliant on classical computational resources for support, the opportunity has emerged to improve the performance of these systems through the creation of bespoke ensembles that more closely resemble the system under investigation.

In the case of pure states, ensembles are typically generated according to the Fubini–Study measure, induced by the Haar measure over the unitary group $U(D)$ [20,21,24]. A random pure state $|\psi_R\rangle$ can be found by first generating a Haar-random

unitary from $U(D)$ and applying it to any fixed quantum state $|\psi\rangle$ to obtain $|\psi_R\rangle = U|\psi\rangle$. Equivalently, one could use a column of the Haar-distributed unitary as $|\psi_R\rangle$, since $|\psi\rangle$ can be chosen arbitrarily.

Several methods for generating ensembles of mixed states exist. Two of the most widely used are based on the Hilbert–Schmidt (HS) and Bures measures. The HS measure has a simple physical interpretation as that which is induced when a Haar-random pure state $|\psi\rangle$ of dimension D^2 is traced down to dimension D . More generally, the HS measure is a special case of a family of induced measures related to tracing a DM -dimensional Haar-random pure state down to a D -dimensional mixed state. Although somewhat less intuitive physically, another metric commonly used for generating ensembles of random density matrices is the Bures measure. The Bures measure is unique as the sole monotone metric that is both Fisher and Fubini–Study adjusted, thereby aligning with standard metrics in both the classical and pure-state limits [22]. In this sense, the Bures distribution represents a canonical choice for situations of complete randomness, such as a prior for Bayesian state reconstruction when all possible input states can occur.

Recently, an additional ensemble of mixed quantum states based on sums of nonorthogonal Haar-random pure states has been explored as a prior for Bayesian quantum state reconstruction techniques, where the coefficients of these ensembles are distributed according to the symmetric Dirichlet distribution [17–19]. Motivated initially by both its computational simplicity and amenability to tuning effective rank [17],

*These authors contributed equally to this work.

[†]rglasser@tulane.edu

[‡]brian.t.kirby4.civ@army.mil

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.

this distribution allows for significant speed-ups in Bayesian quantum state estimation compared to alternative parametrizations [18]. Yet it is not clear from the existing literature how these states behave outside of limiting special cases, particularly in relation to other standard distributions, thus leaving a significant gap in the theoretical justification for this distribution.

In this paper, we determine various properties of the ensembles of density matrices generated by mixing Haar-random pure states with Dirichlet-distributed coefficients. We find that an immediate advantage of these ensembles compared to standard alternatives is that they can be biased toward a particular mean purity. We also determine the Dirichlet concentration parameter required to recover mean purities equal to those of the standard Bures or HS measures. Numerical simulation of up to ten-dimensional states indicates that for proper tuning of the concentration parameter, we can exactly recover the HS distribution, offering an alternative and intuitive physical interpretation for random quantum states distributed according to the HS measure. We then apply our results to generate training data for a machine-learning-based quantum state tomography system and as a prior distribution for Bayesian state reconstruction. In each case, we find measurable performance improvements using tailored ensembles of states as opposed to those based on Bures or HS. Finally, we perform experiments on both an IBM quantum computer (IBM Q) and a source of polarization-entangled photons to determine the typical ensemble of states generated by each throughout a characteristic experiment. In each case, we find that we can more closely match the distributions of these systems using Dirichlet-distributed mixtures than by standard methods.

II. BIASED DISTRIBUTIONS

A. Dirichlet distribution

The Dirichlet distribution is defined for vectors $\mathbf{x} = (x_1, \dots, x_K)$ where the elements of \mathbf{x} belong to the open $K - 1$ simplex: $x_j \geq 0$ and $\sum_{j=1}^K x_j = 1$. The probability density function of the Dirichlet distribution is defined by

$$\text{Dir}(\mathbf{x}|\boldsymbol{\alpha}) = \frac{\Gamma(\sum_{j=1}^K \alpha_j)}{\prod_{j=1}^K \Gamma(\alpha_j)} \prod_{j=1}^K x_j^{\alpha_j-1}, \quad (1)$$

where $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_K)$ with all $\alpha_j \geq 0$ defines the concentration parameters and $\Gamma(\cdot)$ is the standard gamma function.

The concentration parameters determine the properties of the distribution. To understand this intuitively, in Fig. 1 we plot 10^5 Dirichlet-distributed random points on a ternary plot when $K = 3$. The four examples shown in Fig. 1 are representative exceptional cases. In Fig. 1(a), the distribution is uniform, with an equal probability of sampling anywhere within the simplex. The uniform sampling of Fig. 1(a) contrasts with Figs. 1(b)–1(d), where the concentration parameters bias the distribution towards the center, away from the center, or into one corner of the simplex, respectively.

The behavior visualized in Fig. 1 is also reflected in the mean, variance, and covariance of the Dirichlet distribution. These statistical properties will be useful in what follows and

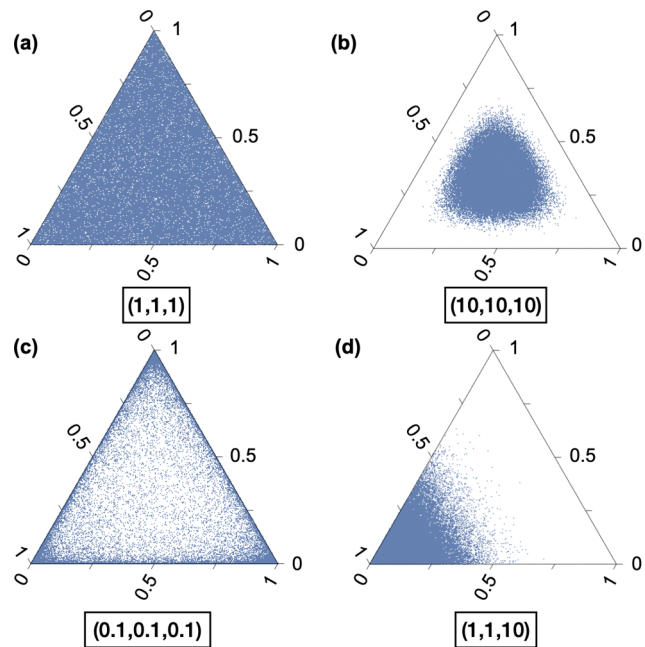


FIG. 1. Ternary plots of 10^5 Dirichlet vectors $\mathbf{x} = (x_1, x_2, x_3)$. Each plot is sampled from a distribution with a different set of concentration parameters, with (a)–(c) symmetric.

are given by

$$\begin{aligned} E[X_j] &= \tilde{\alpha}_j \\ \text{var}[X_j] &= \frac{\tilde{\alpha}_j(1 - \tilde{\alpha}_j)}{\alpha_0 + 1} \\ \text{cov}[X_j, X_k] &= \frac{\delta_{jk}\tilde{\alpha}_j - \tilde{\alpha}_j\tilde{\alpha}_k}{\alpha_0 + 1} \end{aligned} \quad (2)$$

for all $j, k \in \{1, \dots, K\}$, with δ_{jk} the Kronecker delta, $\alpha_0 = \sum_{j=1}^K \alpha_j$, and $\tilde{\alpha}_j = \alpha_j/\alpha_0$. In the equations below, we note that we do not always use capital letters to distinguish random variables from the specific values they assume when the result is more clear otherwise.

Throughout this paper, we will be interested in special cases of the Dirichlet distribution with all α_j equal, known as the symmetric Dirichlet distribution. Figures 1(a)–1(c) are examples of the symmetric Dirichlet distribution, which behave analogously for any value of K . In particular, the flat Dirichlet distribution occurs when all concentration parameters are equal to 1 for any K .

B. Constructing mixtures of Dirichlet-weighted pure states

The properties of the Dirichlet distribution described above make it remarkably well suited for defining the coefficients of a mixture of pure quantum states. To begin with, the restriction of the components of \mathbf{x} to the $K - 1$ simplex is precisely the requirement for the coefficients of a convex sum, such as a density matrix, which is the convex sum of pure quantum states. Further, we can already see intuitively from Fig. 1 that the concentration parameters offer a powerful way of altering the overall properties of the coefficients.

An ensemble of D -dimensional mixed quantum states made from a convex sum of K Haar-random pure states $|\psi_j\rangle$

is given by

$$\rho = \sum_{j=1}^K x_j |\psi_j\rangle\langle\psi_j|, \tag{3}$$

where the vector \mathbf{x} is a random variable distributed according to $\text{Dir}(\mathbf{x}|\alpha)$, having been specialized to the symmetric case $\alpha = \{\alpha, \dots, \alpha\}$. To our knowledge, this ensemble of states was first considered by Mai and Alquier [17], inspired by work on low-rank matrix estimation via machine learning [25,26]. In light of this, and for convenience in the following analysis, we dub this construction the ‘‘Mai–Alquier’’ (MA) distribution. For our purposes here, we apply the MA label specifically to the case of symmetric Dirichlet weights, but admit any $K \in \mathbb{N}$; thus, for a given dimension D , the MA distribution requires specification of two parameters, α and K . Since first being defined in [17], the MA distribution has been used to analyze data from several quantum optical experiments via Bayesian quantum state tomography [18,19,27,28].

To characterize the properties of an ensemble of states generated according to Eq. (3), we begin by calculating the mean purity. We have chosen purity as our primary metric of interest for several reasons. The first is that, like the distributions we are analyzing, it is invariant under local rotations, and hence we can abstract away any considerations related to system alignment. Further, the purity of a quantum system can be estimated without full state tomography [29]. Finally, it applies to both individual and composite quantum systems equally well, as opposed to an entanglement metric, which is only applicable to the latter. The purity of ρ is given by

$$\text{Tr}(\rho^2) = \sum_{j=1}^K x_j^2 + 2 \sum_{j>k}^K x_j x_k F_{j,k}, \tag{4}$$

where $F_{j,k} = |\langle\psi_j|\psi_k\rangle|^2$. Since $2 \sum_{j>k}^K x_j x_k F_{j,k} \geq 0$, we see that for a given set of coefficients \mathbf{x} the minimum purity occurs when all $F_{j,k} = 0$, meaning that the states in the sum are orthogonal. This related case, where all states in Eq. (3) are orthogonal, was studied by Życzkowski in [30] in the context of estimating the volume of the separable states in the space of all mixed quantum states.

Random states of D dimension from the ensemble described in [30] can be generated by taking $K = D$ in Eq. (1), using \mathbf{x} as the entries in a diagonal matrix, and rotating this matrix by a Haar-random unitary from $U(D)$. Many studies have generated random density matrices in this fashion, but of particular interest to applications discussed in Sec. III A is [8], which uses this method to generate training data for a machine-learning-based separability-entanglement classifier. In the Appendix, we characterize the mean purity of the distribution described by Życzkowski as a function of the Dirichlet concentration coefficients and compare these results with the Bures, HS, and MA distributions. We find that the Życzkowski distribution is, as the MA distribution, capable of biasing based on mean purity. Also in the Appendix, we compare how well the Życzkowski and MA distributions fit the experimental scenarios of Sec. IV, and find in all three scenarios the MA distribution fits as well or better than the Życzkowski distribution, as measured by the Bhattacharyya coefficient [31]. As a final point of comparison, in Sec. II C

we will find that the MA distribution appears to simplify to the HS distribution given the correct choice of concentration parameter, which, to our knowledge, is not possible with the Życzkowski distribution. For these reasons, we focus our analysis on the MA distribution in our comparisons with Bures and HS here.

The expectation value of the purity can be found from

$$\begin{aligned} E_{MA}[\text{Tr}(\rho^2)] &= \sum_{j=1}^K E[x_j^2] + 2 \sum_{j>k}^K E[x_j x_k F_{j,k}] \\ &= \sum_{j=1}^K E[x_j^2] + 2 \sum_{j>k}^K E[x_j x_k] E[F_{j,k}], \end{aligned} \tag{5}$$

where we have used the linearity of the expectation operator and the fact that the states $|\psi_j\rangle$ are statistically independent of the weight coefficients x_j . (When needed for clarity in this and what follows, we apply a subscript $s \in \{MA, B, HS\}$ to expectations and variances in order to distinguish the density matrix distribution over which the quantity is computed.) For two Haar-randomly chosen pure states, the fidelity is distributed according to [23]

$$P(F) = (D - 1)(1 - F)^{D-2}. \tag{6}$$

From this we can calculate the first moment as

$$E[F] = \int_0^1 F P(F) dF = \frac{1}{D}, \tag{7}$$

which corresponds with the expression in [23]. The remaining expectation values can be found from Eq. (2) and the standard relationships

$$\begin{aligned} \text{var}[X] &= E[X^2] - E[X]^2 \\ \text{cov}[X_j, X_k] &= E[X_j X_k] - E[X_j] E[X_k]. \end{aligned} \tag{8}$$

The resulting expectation value of the purity is then given by

$$E_{MA}[\text{Tr}(\rho^2)] = \frac{D + \alpha(D + K - 1)}{D(1 + \alpha K)}, \tag{9}$$

where we have evaluated $\sum_{j>k}^K = \frac{K^2 - K}{2}$.

It is instructive to consider the behavior of Eq. (9) in several special cases. In the limit $\alpha \rightarrow \infty$, the variance of the symmetric Dirichlet distribution vanishes, and all coefficients approach $x_j = 1/K$. This limit corresponds to the case explored in both [32,33], and we find that for large α our Eq. (9) reduces to the expression found in [33]: $E[\text{Tr}(\rho^2)] = (D + K - 1)/DK$. (We will return to the results of [32] in Sec. II C in the context of comparisons with the HS ensemble.) In the opposite limit where all of the $\alpha \rightarrow 0$, the Dirichlet distribution favors a single entry of value 1 with all others 0 [34], and thus we find a mean purity of unity as all states are rank-1.

The ability to alter the properties of this ensemble of states is further exemplified in Fig. 2(a), where we plot the smoothed probability density function obtained from numerical samples of the purity for different α when $K = D = 4$ (corresponding to two qubits). In particular, we show ten values of α in steps of 0.1 from $\alpha = 0.1$ to $\alpha = 1$, with these extremal cases plotted in solid red and the intermediate cases in dashed blue. Note also that the α values of the solid-red lines correspond

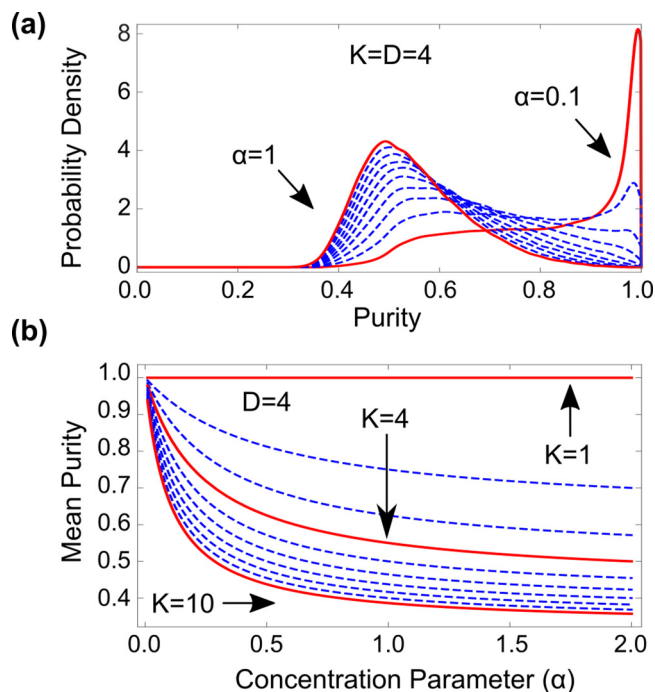


FIG. 2. Behavior of the probability density function and mean purity of the MA distribution. (a) Numerically sampled and smoothed probability density function of MA distributed random states for $K = D = 4$. Ten different curves are shown corresponding to $\alpha = 0.1$ to $\alpha = 1$ in steps of 0.1. The solid (red) lines are $\alpha = 0.1$ and $\alpha = 1$, with the other values shown as dashed (blue) lines. (b) Mean purity of MA distributed states for $D = 4$ and $K = 1$ to $K = 10$ in steps of 1. The solid (red) lines indicate $K = 1, 4, 10$, with the other values shown with dashed (blue) lines. The $K = 4$ curve corresponds to the curves in part (a).

to those visualized for the $K = 3$ distribution in Figs. 1(a) and 1(c).

In Fig. 2(b) we consider the impact of increasing K for a fixed D , in other words, making each mixed state in the ensemble a sum of more pure states. Specifically plotted is Eq. (9) with $D = 4$ and K ranging from 1 to 10 in single increments, as a function of α , with the end cases in solid red. Further, for ease of comparison with Fig. 2(a), we also make the $K = D = 4$ case in solid red, where the mean of each distribution corresponds to the y value along the $K = 4$ line. As expected, the more terms in the sum of each ρ , the more rapidly the states in the ensemble become mixed as α increases.

Finally, we can also place an upper bound on the variance of the purity of ρ using the Bhatia–Davis inequality as [35]

$$\text{var}_{MA}[\text{Tr}(\rho^2)] \leq \frac{\alpha(1 + \alpha)(D - 1)^2(K - 1)}{(D + \alpha DK)^2}. \quad (10)$$

We see for a fixed D and K that in the limit $\alpha \rightarrow 0$ we find $\text{var}_{MA}[\text{Tr}(\rho^2)] \rightarrow 0$, which is consistent with the intuition that all but one element in \mathbf{x} approaches zero in this limit, meaning that ρ becomes a pure state. In the opposite limit of $\alpha \rightarrow \infty$, we find for fixed D and K that

$$\lim_{\alpha \rightarrow \infty} \text{var}_{MA}[\text{Tr}(\rho^2)] \leq \frac{(D - 1)^2(K - 1)}{D^2 K^2}. \quad (11)$$

This limit includes but does not equal 0, which is consistent with the interpretation that when $\alpha \rightarrow \infty$, the state will not always approach the completely mixed state. This is obvious if we consider $K < D$, as it is impossible for the state to become full rank, and thus no set of coefficients could result in an identity matrix. However, we note that a completely mixed state can be reached consistently for any α and D if we instead take the limit

$$\lim_{K \rightarrow \infty} \text{var}_{MA}[\text{Tr}(\rho^2)] = 0. \quad (12)$$

Physically, this can be interpreted to mean that if ρ is composed of a sufficiently large sum of Haar-random pure states, it will always approach the maximally mixed state, regardless of coefficients.

Finally, we note that although extensions of our results beyond the symmetric distribution are straightforward, we have focused on the symmetric distribution in this section because the equivalence under exchange of index simplified the presentation. It is possible that asymmetric distributions may offer additional flexibility in specific cases, and generalizations in this direction represent an interesting area for further research.

C. Comparisons with standard methods

We begin by reviewing standard methods for generating random quantum states according to the Bures and HS distributions. While each can in principle be induced through the partial trace on higher dimensional pure states [21], it is simpler in practice to derive each from the complex Ginibre ensemble [24], which consists of $D \times D$ complex matrices with independently chosen entries from a standard normal distribution [36]. A random quantum state ρ from the Bures ensemble can be generated according to

$$\rho = \frac{(\mathbb{1} + U)GG^\dagger(\mathbb{1} + U^\dagger)}{\text{Tr}[(\mathbb{1} + U)GG^\dagger(\mathbb{1} + U^\dagger)]}, \quad (13)$$

and from the Hilbert–Schmidt ensemble by

$$\rho = \frac{GG^\dagger}{\text{Tr}(GG^\dagger)}, \quad (14)$$

where G is a random matrix from the Ginibre ensemble and U is a Haar-distributed random unitary from $U(D)$ [24].

The average purity of Bures-distributed states was found in [24] and is given by

$$E_B[\text{Tr}(\rho^2)] = \frac{5D^2 + 1}{2D(D^2 + 2)}. \quad (15)$$

Similarly, for the HS distribution, the average purity was found to be [21]

$$E_{HS}[\text{Tr}(\rho^2)] = \frac{2D}{D^2 + 1}. \quad (16)$$

We immediately see that [37]

$$E_B[\text{Tr}(\rho^2)] > E_{HS}[\text{Tr}(\rho^2)] \quad (17)$$

for any dimension $D \geq 2$.

For ease of comparison, we fix $K = D$ in Eq. (9), which becomes

$$E_{MA}[\text{Tr}(\rho^2)]_{K=D} = \frac{D + \alpha(2D - 1)}{D + \alpha D^2}. \quad (18)$$




FIG. 3. Comparison of the numerically sampled probability distribution of purity for MA, Bures, and HS distributed random quantum states. Both plots show $D = K = \{2, \dots, 10\}$, and each curve was created using 10^6 samples. (a) The Bures distribution is plotted with solid (red) lines, and the corresponding MA distribution with $\alpha = \alpha_B$ is plotted with dashed (black) lines. (b) The HS distribution is plotted with solid (red) line, with the corresponding MA distribution with $\alpha = \alpha_{HS}$ plotted with dashed (black) lines. We see complete overlap between the HS and MA distributions when $\alpha = \alpha_{HS}$.

We note that $K = D$ is a natural choice as it is the minimum K capable of producing a full rank state. From these expressions, we find that $E_{MA}[\text{Tr}(\rho^2)]_{K=D}$ can be made equal to either the Bures or HS averages by setting α equal to the following:

$$\alpha_B = \frac{2D - 1}{4 + D}, \quad \alpha_{HS} = D. \tag{19}$$

The probability density for the purity of the Dirichlet-distributed mixed states and the Bures and Hilbert–Schmidt ensembles are plotted in Fig. 3. The Bures and Hilbert–Schmidt probability densities are in solid (red) with the $\alpha = \alpha_B$ and $\alpha = \alpha_{HS}$ MA ensembles plotted with dashed (black) lines. The probability densities are shown for $K = D$ and $D \in \{2, \dots, 10\}$: the $D \in \{2, 10\}$ cases are labeled, with the others in sequential order. For all dimensions considered, the tailored MA distribution appears to reproduce the HS ensemble exactly.

For all $D \geq 2$, $\alpha_{HS} > \alpha_B$, and hence we arrive at the following inequalities:

$$\mu_{MA}^{\alpha < \alpha_B} > \mu_B > \mu_{HS} > \mu_{MA}^{\alpha > \alpha_{HS}}, \tag{20}$$

where we have used the shorthand $\mu = E[\text{Tr}(\rho^2)]$ and the MA distributions assume $K = D$. The findings summarized in Eqs. (19), (20) represent major contributions of our present investigation, revealing quantitatively how the parameter α of the MA distribution can be tuned to obtain an equal, higher, or lower mean purity compared to well-known fiducial density matrix measures. As we will see below, for many applications

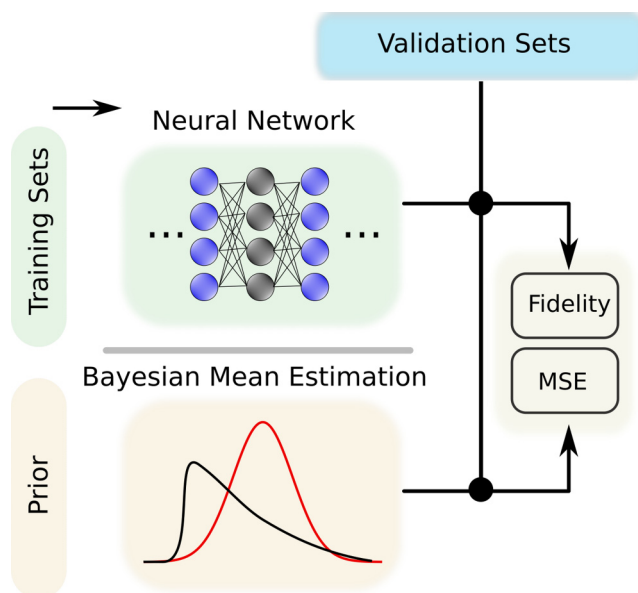


FIG. 4. Two advanced methods for reconstructing quantum states: neural network (top) and Bayesian mean estimation (bottom). Both methods require a defined distribution as input: the training set for the neural network, and prior for BME. In both methods, fidelity and mean-squared error (MSE) between the target and reconstructed quantum states are evaluated for common validation sets.

the system under investigation may produce states that are nearly pure on average, and hence randomly sampling states with $\alpha < \alpha_B$ can bias the distribution to better reflect the states of interest. Such tunability in rank is unavailable from Bures (by definition); for HS, changing to draws of complex-normal $D \times K$ matrices G (with $K < D$) can be used to reduce the rank—and increase purity—of ρ formed using Eq. (14) [15], but this comes at the cost of eliminating full-rank states from consideration altogether. In contrast, adjusting α in the MA ensemble is able to favor pure or mixed states while still providing support over the full D -dimensional Hilbert space.

Finally, as a side note, the α values found to match the mean purities of Bures and HS distributions in Eq. (19) are almost as interesting for what they are *not* as for what they are. For example, neither case corresponds to $\alpha = 1$, despite the fact that $\alpha = 1$ gives a uniform distribution of coefficients over the simplex [Fig. 1(a)] and therefore ostensibly may seem the default choice for a uniform density matrix distribution. At the other extreme, nor does the limit $\alpha \rightarrow \infty$ give mean purities commensurate with Bures or HS. Indeed this latter limit corresponds to the sum of D evenly weighted Haar-random pure states, which was noted in [32] to differ from the HS ensemble in general; thus our conclusion $\alpha_{HS} \neq \infty$ is consistent with these findings as well.

III. APPLICATIONS

We now explore applications that show measurable performance improvements when substituting the MA distribution for either the Bures or HS distributions. In particular, we consider two different methods for reconstructing quantum states from measured tomography data as shown in Fig. 4. The first

method reconstructs states using pretrained neural networks. In Sec. III A, we study the impact of the chosen training set on the fidelity of the reconstruction. The second approach we consider for state reconstruction is Bayesian mean estimation (BME). In Sec. III B, we examine the significance of using a carefully selected prior distribution for the rate of convergence of a BME-based reconstruction method.

So that we can study the efficacy of these techniques in a range of situations, we have opted to perform state reconstruction as a function of the number of measurements performed. In particular, we perform reconstruction of ($D = 4$)-dimensional states using simulated data sets from random Pauli measurements (one measurement per each random basis selection) with the total number of measurements ranging from 1 to 1000. This approach allows us to view the disparity in performance both in the high-statistical noise regime where very little data is available and in the more ideal asymptotic case where the observed frequencies approach the actual probability distributions. The low-count regime is of particular interest for tomography of high-dimensional systems where the size and complexity of the system may limit the amount of measurements that can be feasibly performed.

In each section, we perform numerical simulations using sets of randomly sampled quantum states. For training, we utilize full 1000-measurement data sets from 10^5 randomly chosen states from a given metric: Fubini–Study (pure states), Bures, HS, or MA of a particular α . Pure states of dimension $D = 4$ were sampled by taking the first column of a Haar-distributed unitary from $U(D)$. We sample the Bures and HS distributions using Eqs. (13) and (14). Validation is then performed using 2000 randomly chosen states from the metric under consideration (which need not match the training set), where we consider total measurement numbers of $M \in \{1, 15, 25, 50, 75, 100, 200, 400, 600, 800, 1000\}$. The Bayesian approach does not utilize a training set (the analog functionality is realized by the mathematically specified prior distribution), but performance is tested using the same validation sets as the neural network cases.

A. Neural network quantum state tomography

Machine learning has found broad applicability in quantum information science in topic areas as diverse as experiment design [38], state classification [8,39,40], and even studies on quantum foundations [41]. Recently, several studies have applied machine-learning methods to quantum state reconstruction [8–12]. Here we consider how well a machine-learning-based quantum state tomography system can reconstruct states of one distribution when trained on another. Specifically, we will first consider reconstructing pure two-qubit quantum states by training a network exclusively on either (i) Haar-random pure states, (ii) randomly generated states from the Bures distribution, or (iii) randomly generated states from the HS distribution. Further, we will demonstrate the versatility of the MA distribution by separately training a network on randomly generated states from ten different MA distributions (α from 0.1 to 1 in steps of 0.1), and we will show how well each of these networks reconstructs pure-, HS-, and Bures-distributed validation states. We note that a

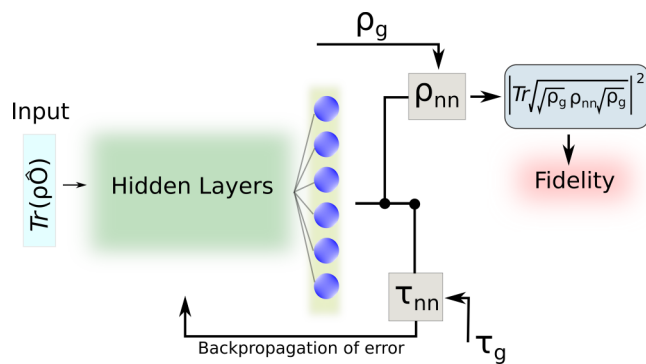


FIG. 5. Architecture of the neural network. ρ_g and ρ_{nn} represent the ground truth and predicted density matrices. Similarly, τ_g , and τ_{nn} , respectively are the ground truth and predicted tau vectors.

previous study has considered the issue of matching training distributions to validation distributions and dealt with it using a pipeline of neural networks trained on different distributions [10]. Instead, in this study we consider the generalizability of a single neural network for reconstructing states from the entire Hilbert space.

The architecture of the neural network [42] we consider is shown in Fig. 5. The network consists of an input layer that receives the measured tomography values, $\text{Tr}(\rho \hat{O})$, which connects to a convolutional layer with a kernel size of (2, 2), stride lengths of 1, ReLU as an activation function, and filters of size 64. Then we attach a max-pooling unit with a pool-size of (2, 2), stride of length 2, and the “valid” padding, which is, again, followed by another convolutional layer with the same configuration as mentioned before. Next, we implement a flattening layer, which is then fully connected to a dense layer with 3000 neurons using the ReLU activation function. We then apply a dropout unit with a 50% dropout rate, followed by another dense layer with 1200 neurons using a ReLU activation function, followed by a dropout unit with the same rate. After this, we fully connect another dense layer with a linear activation as shown by blue circles in Fig. 5. The prediction made at this layer (blue circles in Fig. 5) is branched into two pipelines. The first pipeline evaluates the τ_{nn} vectors and compares them with the ground truth values τ_g . Note that the target τ_g vectors are obtained from the Cholesky decomposition of the target density matrices (ρ_g). In the two-qubit scenario, for any random quantum state, $\rho_g = \xi \xi^\dagger$, where

$$\xi = \begin{bmatrix} \tau_0 & 0 & 0 & 0 \\ \tau_4 + i\tau_5 & \tau_1 & 0 & 0 \\ \tau_{10} + i\tau_{11} & \tau_6 + i\tau_7 & \tau_2 & 0 \\ \tau_{14} + i\tau_{15} & \tau_{12} + i\tau_{13} & \tau_8 + i\tau_9 & \tau_3 \end{bmatrix}, \quad (21)$$

which can be rearranged as the vector

$$\xi \rightarrow \tau = (\tau_0, \tau_1, \tau_2, \tau_3, \dots, \tau_{15}). \quad (22)$$

Consequently, the mean-squared loss between the predicted and the target τ vectors is obtained and back-propagated for the learning, while the second pipeline makes predictions for the corresponding density matrices from the predicted τ_{nn} vectors. At the end, the network compares the predicted density matrices (ρ_{nn}) with the ground truth states, and evaluates the

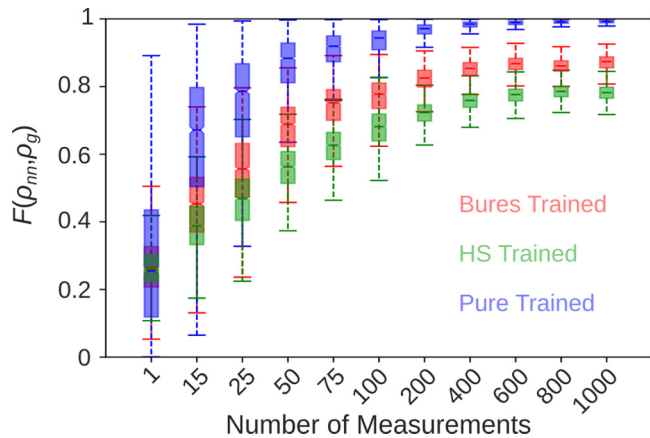


FIG. 6. Reconstructing pure states from networks separately trained with either Haar-random pure states, Bures-drawn mixed states, or HS-drawn mixed states.

fidelity

$$F(\rho_{nm}, \rho_g) = \left[\text{Tr} \sqrt{\sqrt{\rho_g} \rho_{nm} \sqrt{\rho_g}} \right]^2. \quad (23)$$

Ultimately, the final network settings are chosen to maximize the average fidelity of the reconstructed states to the ground truth.

To illustrate the proof of concept, we reconstruct two-qubit quantum states using networks separately trained using each of the ensembles introduced above. Each network is trained for 75 epochs at a learning rate of 0.008. Once the networks are trained, we make predictions for a validation set of pure quantum states that are unknown to the networks. The reconstruction fidelities for the predicted quantum states for various measurement scenarios are shown in Fig. 6. We find that the network trained with pure states quickly reconstructs the validation set with near unit fidelity as shown by blue box plots. However, the networks separately trained with Bures (red box) and HS (green box) states struggle to reconstruct the validation set with reconstruction fidelities leveling off around 0.87 and 0.79, respectively. For this and all subsequent box plots below, the notch gives the median and each box encloses $[Q_1, Q_3]$, while the whiskers extend from $Q_1 - 1.5(Q_3 - Q_1)$ to $Q_3 + 1.5(Q_3 - Q_1)$, where Q_1 and Q_3 are the first and third quartiles. At first, this is somewhat surprising, as the pure states are a subset of the mixed states. Hence, a naive assumption is that a system capable of reconstructing mixed states with high fidelity would also reconstruct pure states with similar fidelity. Therefore, we are faced with a problem of network generalization: the reconstruction fidelity of a neural network depends heavily on the distribution used to train it. One possible way to overcome this generalization issue is to train multiple neural networks for different situations, as we performed in [9], or even to combine them together into a pipeline, as was done in [10]. Instead, we attempt to solve the same problem by biasing the training distribution to better cover the Hilbert space through proper tuning of the MA distribution. In Fig. 7, we vary α of the MA distribution from 0.1 to 1.0 in steps of 0.1 and sample training sets for 1000 random Pauli measurements. For each α , we separately train a network on states randomly sampled from the corresponding MA




FIG. 7. Reconstruction fidelity versus various concentration parameters (α). The fidelity at various measurements with the network trained at $\alpha = 0.4$ is shown in the inset, which corresponds to Fig. 6.

distribution. Once the network is trained, we reconstruct the validation sets containing Bures-, HS-, and Haar-distributed pure states individually. The reconstruction fidelities for Bures and HS states gradually increase with an increase in α (red and green box plots), whereas the reconstruction fidelities for Haar-random pure states slowly decrease with an increase in α (blue box plots). However, reconstruction fidelities for the three cases approximately coalesce when training states are sampled from a distribution with $\alpha \approx 0.4$. The difference in reconstruction fidelity seen in Fig. 7 at $\alpha = 0.1$ and $\alpha = 1$ for different validation sets stresses the importance of carefully selecting training sets for machine-learning-based reconstruction techniques. For comparison, we also include an inset in Fig. 7 in the same form as Fig. 6. We see in the inset that reconstruction of all three validation sets converge, unlike the situation in Fig. 6.

This investigation highlights that by tuning the degree of sparsity of the Dirichlet distribution, we are able to hedge our bets on accurate state estimation. A neural network trained with a parameter setting of $\alpha = 0.4$ attains significantly higher reconstruction fidelities for validation on pure states compared to a neural network trained on Bures. Further, as we see in Fig. 7, this gain in pure state reconstruction fidelity (blue boxes), comes with a relatively minor reduction in Bures and HS reconstruction fidelity (red and green boxes, respectively). Hence, through proper tuning of the training set, we are able to create a neural network reconstruction system with significantly more generality than one trained on Bures, HS, or pure states alone.

B. Bayesian quantum state tomography

Here, we continue to explore inference for the same data sets, but now employ Bayesian mean estimation (BME) [13,43], ultimately finding similar advantages for the custom states as in the neural network cases. In contrast to alternative approaches in quantum state tomography, Bayesian inference defines a complete probability distribution for ρ , utilizing Bayes' theorem to combine prior knowledge and experimental results within a single consistent framework. Bayesian tomography enjoys several appealing features, including automatic

uncertainty quantification and the return of reliable estimates under any measurement conditions [13]. Moreover, the mean of the posterior distribution is optimal in that it minimizes the mean-squared error with respect to the ground truth for any number of measurements [44]—hence the emphasis on “mean” in BME. Although BME remains fairly uncommon in quantum state tomography, due in large part to the computational difficulties associated with high-dimensional integrals, several practical Monte Carlo approaches for quantum BME [14–18] have appeared following the initial proposal [13], making it an increasingly attractive prospect in quantum state estimation.

Formally, in the Bayesian viewpoint we assign a probability distribution to \mathbf{w} , a vector of all parameters necessary to define a given density matrix $\rho(\mathbf{w})$. The length and content of \mathbf{w} can vary depending on the chosen parametrization, but it is assumed that all allowed values produce a physical $\rho(\mathbf{w})$, i.e., unit trace, Hermitian, and positive semidefinite. According to Bayes’ rule, the posterior probability density for \mathbf{w} following an experiment can be written as

$$\pi(\mathbf{w}) = \frac{1}{\mathcal{Z}} L_{\mathcal{D}}(\mathbf{w}) \pi_0(\mathbf{w}), \quad (24)$$

where the likelihood $L_{\mathcal{D}}(\mathbf{w}) \propto P(\mathcal{D}|\mathbf{w})$ (the probability of observing data set \mathcal{D} given parameters \mathbf{w}), $\pi_0(\mathbf{w})$ is the prior on \mathbf{w} , and \mathcal{Z} is a normalizing constant (which does not need to be determined in the following examples). Considering a collection of M measurements on a repeatedly prepared input state, with each result described by an operator Λ_m , the likelihood can be written as $L_{\mathcal{D}}(\mathbf{w}) = \prod_{m=1}^M \text{Tr} \rho(\mathbf{w}) \Lambda_m$; in the special case of projective measurements (such as the Pauli measurements in our example), this becomes

$$L_{\mathcal{D}}(\mathbf{w}) = \prod_{m=1}^M \langle \psi_m | \rho(\mathbf{w}) | \psi_m \rangle, \quad (25)$$

with $|\psi_m\rangle$ the eigenstate observed in measurement m .

Of primary interest for our present purposes, however, is the prior $\pi_0(\mathbf{w})$. Just as the training set in the neural network case specifies the realm of states expected in subsequent experiments, $\pi_0(\mathbf{w})$ expresses all *a priori* assumptions about the system under test in the form of an explicit probability distribution over states $\rho(\mathbf{w})$. As frequently emphasized in BME, the exact functional form of any well-chosen prior—e.g., one which assigns reasonable probability to all states in the Hilbert space—will have minimal impact on the posterior $\pi(\mathbf{w})$ in the limit $M \rightarrow \infty$, since the likelihood $L_{\mathcal{D}}(\mathbf{w})$ then dominates. However, $\pi_0(\mathbf{w})$ can have a profound impact in the case of few measurements, such as the conditions explored here.

Given the neural network findings in Fig. 7, we focus our comparison specifically on priors corresponding to either the Bures, HS, or MA ($K = 4$, $\alpha = 0.4$) distribution. For each option, we select a parametrization amenable to numerical sampling. Bures states can be represented by the vector $\mathbf{w} = (w_1, \dots, w_{2D^2})$, with each element independently distributed according to a complex standard normal distribution $w_k \sim \mathcal{CN}(0, 1)$. We assign D^2 of the components to populate the $D \times D$ Ginibre matrix G in Eq. (13); the remaining D^2 elements comprise a second, independent Ginibre matrix, which is then fed into the algorithm of [45] to produce

the Haar-random unitary U also required in Eq. (13). Thus the Bures prior can be written as $\pi_0(\mathbf{w}) \propto \prod_{k=1}^{2D^2} e^{-\frac{1}{2}|w_k|^2}$. Since HS states can be represented using a single Ginibre draw [Eq. (14)], we can shorten \mathbf{w} to a length- D^2 vector and define the HS prior as $\pi_0(\mathbf{w}) \propto \prod_{k=1}^{D^2} e^{-\frac{1}{2}|w_k|^2}$.

For the MA prior, we take $\mathbf{w} = (\mathbf{u}, \mathbf{v}_1, \dots, \mathbf{v}_D)$, where $\mathbf{u} = (u_1, \dots, u_D)$ is a real D -dimensional vector of positive scalars, and \mathbf{v}_k is a D -dimensional complex vector. The prior follows $\pi_0(\mathbf{w}) \propto \prod_{k=1}^D u_k^{\alpha-1} e^{-u_k} e^{-\frac{1}{2} \mathbf{v}_k^\dagger \mathbf{v}_k}$, and we define $\rho(\mathbf{w})$ according to

$$\rho(\mathbf{w}) = \sum_{k=1}^D \left(\frac{u_k}{\sum_l u_l} \right) \frac{\mathbf{v}_k \mathbf{v}_k^\dagger}{|\mathbf{v}_k|^2}. \quad (26)$$

This equation is equivalent to Eq. (3) expressed in the computational basis, though for convenience we now utilize unnormalized parameters; our combined gamma and complex-normal prior ensures that the normalized entities $\mathbf{u}/\sum_l u_l$ and $\mathbf{v}_k/|\mathbf{v}_k|$ are Dirichlet- and Haar-distributed, respectively, as required [17].

For either prior, the Bayesian mean estimate for ρ is

$$\rho_{BME} = \int d\mathbf{w} \pi(\mathbf{w}) \rho(\mathbf{w}), \quad (27)$$

which we compare to the ground truth ρ_g by fidelity $F(\rho_{BME}, \rho_g)$ [Eq. (23)]. Computation of Eq. (27) is effected utilizing preconditioned Crank-Nicolson (pCN) Markov chain Monte Carlo (MCMC) techniques [46] recently introduced to quantum state tomography [18], with the only modification being an improved, reversible proposal distribution for the gamma random variables u_k based on [47]. A total of 2^{10} samples of \mathbf{w} are kept in estimating Eq. (27), and a thinning value of 2^8 was found sufficient for convergence in fidelity, thus corresponding to total MCMC chain lengths of 2^{18} .

The distribution of fidelities $F(\rho_{BME}, \rho_g)$ obtained for several sets of quantum states appear in Fig. 8, plotted as a function of M : Figure 8(a) corresponds to ground truth draws (i.e., a validation set) from the Bures distribution, and Fig. 8(b) to pure states. The fidelities increase rapidly with measurements M , and all three priors produce nearly identical distributions for the Bures ground truth states, whereas the MA prior displays a noticeable edge over both Bures and HS when the ground truth draws are pure.

To further quantify this comparison, we can take advantage of theoretical optimality guarantees offered by the Bayesian mean estimate ρ_{BME} , namely, that ρ_{BME} minimizes the expected value of every operational divergence, a type of reward-scheme-motivated metric quantifying the closeness of an estimator to the ground truth ρ_g [13]. While fidelity does not correspond to such an operational divergence, the squared Frobenius distance $D_F^2(A, B) = \|A - B\|_F^2 = \text{Tr}[(A - B)^\dagger (A - B)]$ does. Accordingly, if we define a mean squared error (MSE) according to $\langle D_F^2(\tilde{\rho}, \rho_g) \rangle$ —the average squared distance over all validation states for a particular type of estimator $\tilde{\rho} - \tilde{\rho} = \rho_{BME}$ is guaranteed to minimize MSE, for any number of measurements, provided that the prior $\pi_0(\mathbf{x})$ used matches the actual distribution from which ρ_g is drawn.

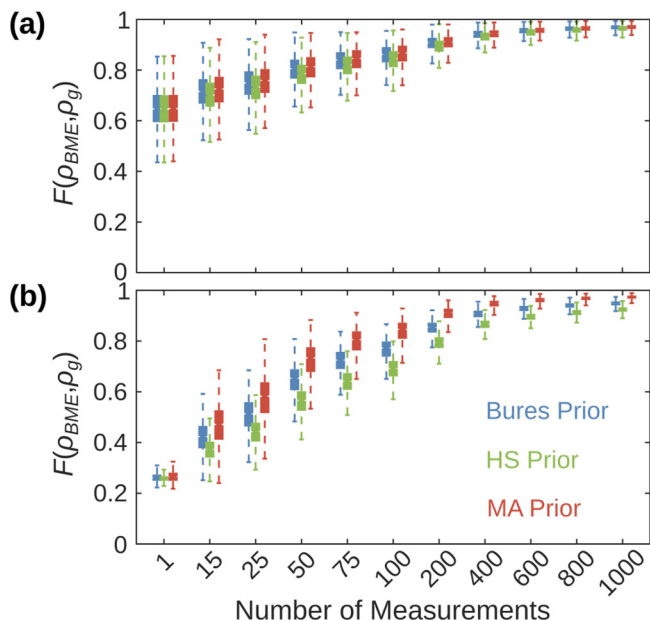


FIG. 8. Bayesian inference fidelity for simulated two-qubit experiments. (a) Validation set drawn from Bures metric. (b) Validation set drawn from Fubini–Study metric. Inference is performed using either the Bures, HS, or MA prior (with $K = D$ and $\alpha = 0.4$), for a specified number of random Pauli measurements.

Calculating MSE for all examples in Fig. 8, as well as for the outputs of the Bures-, HS-, and MA-trained neural networks of Sec. III A with the same input data, we obtain the results in Fig. 9. As expected, the BME results based on a Bures prior attain the smallest error of all inference procedures when the validation states are themselves Bures-

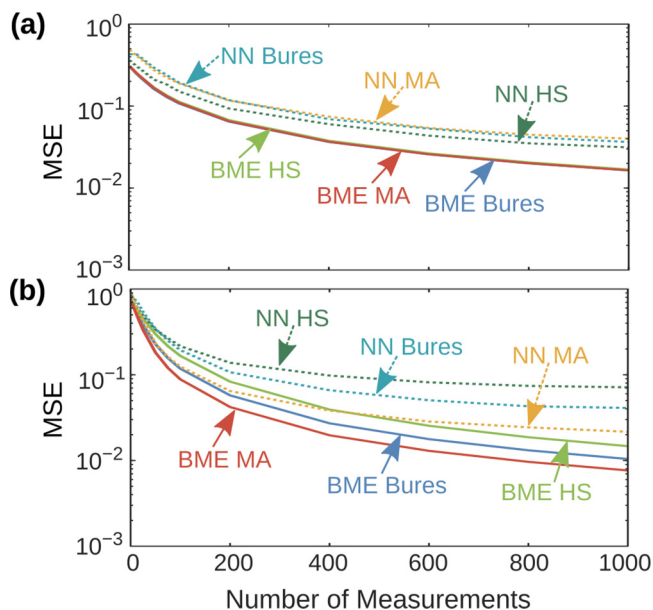


FIG. 9. Mean-squared Frobenius error $\langle D_F^2(\tilde{\rho}, \rho_g) \rangle$ for each estimation procedure. Simulated data sets from (a) Bures-drawn states and (b) Haar-random pure states. BME (neural network) results corresponding to either the Bures or MA priors (training sets) are plotted.

distributed [Fig. 9(a)]. Yet the MA prior with $\alpha = 0.4$ results prove extremely close and are difficult to distinguish on this scale. However, when the validation set is restricted to pure states, the MA prior and MA-trained network improve on their Bures and HS counterparts in realizing more accurate estimates [Fig. 9(b)]. Importantly, in this latter case, none of the BME priors matches the actual validation distribution, so that MSE optimality does not apply to any curve here, in contrast to the case “BME Bures” in (a). As a surprising side note, the neural network trained on HS realizes lower MSE on a Bures validation set than the network actually trained on Bures directly [Fig. 9(a)]; the cause of this is unknown, but we suspect it results from the fact that the neural network was trained to maximize fidelity rather than minimize the Frobenius error, for which performance is not guaranteed.

IV. COMPARISONS WITH EXPERIMENTAL SCENARIOS

In the previous section, our use of simulated data sets allowed us to test the performance of the MA distribution directly against fiducial density matrix distributions. Yet in practice, a quantum system may produce states whose distribution deviates significantly from standard measures such as Fubini–Study, Bures, and HS. Therefore in this section, we will apply the tunability of the MA distribution toward matching the purity distributions of actual experimental systems, specifically a cloud-accessed quantum computer (IBM Q) and a commercial polarization-entangled photon source.

Our primary metric of comparison will be the Bhattacharyya coefficient [31]. We will calculate the Bhattacharyya coefficient from numerically sampled histograms of 100 bins normalized such that the total sums to unity. The number of sampled states used to generate these histograms will vary by scenario, and is explicitly mentioned in each section. To denote the vector of histogram bin heights, we adopt the notation of h_s , where $s \in \{MA, B, HS, IBM1, IBM2, EPS\}$; the first three are for the MA, Bures, and HS distributions, respectively, and the IBM1, IBM2, and EPS subscripts refer to the measured data in Figs. 11, 12, and 14, respectively. We calculate the Bhattacharyya coefficient from

$$\mathcal{B}(h_j, h_k) = \sum_{l=1}^{100} \sqrt{h_j(l)h_k(l)}. \quad (28)$$

The Bhattacharyya coefficient between identical histograms is 1, and two completely nonoverlapping (orthogonal) histograms has a coefficient of 0. While the Bhattacharyya coefficient in this discrete form can depend on how many bins we separate the data into, we find that 100 is sufficient to capture the essential features we want to exemplify.

A. IBM Quantum Computer

To illustrate the proof of concept, we implement two different scenarios on an IBM quantum computer that result in distributions of states with significantly different mean purity. First, we assume a simple approach and initialize a four-qubit quantum state with all qubits in the $|0\rangle$ state and perform tomography without intermediary gates. The intention of this




FIG. 10. IBM Q tomography setup. An arbitrary four-qubit quantum circuit generates random quantum states on a seven-qubit computer (*ibmq_jakarta*), followed by Pauli measurements. MLE is used to reconstruct the quantum states.

scenario is to obtain a distribution of states with very high purity in order to mimic the operation of an “ideal” system.

In the second scenario, we operate on the four initial qubits with a quantum circuit consisting of nine U gates and three $CNOT$ gates, as shown in Fig. 10. The U gates can be represented in matrix form as

$$U(\theta, \phi, \lambda) = \begin{bmatrix} \cos(\frac{\theta}{2}) & -e^{i\lambda} \sin(\frac{\theta}{2}) \\ e^{i\phi} \sin(\frac{\theta}{2}) & e^{i(\phi+\lambda)} \cos(\frac{\theta}{2}) \end{bmatrix}. \quad (29)$$

To generate a distribution of random states, we randomly select the values of (θ, ϕ, λ) for the U gates such that $(\theta, \phi, \lambda) \in 2\pi\mathcal{N}(0, 1)$, where $\mathcal{N}(0, 1)$ is the standard normal distribution.

We perform quantum state tomography using 81 quantum circuits that project on all combinations of the four local Pauli bases $\{X, Y, Z\}_1 \otimes \{X, Y, Z\}_2 \otimes \{X, Y, Z\}_3 \otimes \{X, Y, Z\}_4$. All measurements were executed over the cloud with a seven qubit NISQ-era quantum computer, *ibmq_jakarta* [48]. To reduce statistical noise in the IBM Q measurement results, we execute circuits for 5000 shots. Further, we implement a measurement correction fitter for a full calibration with the method “least_squares,” and perform maximum likelihood estimation (MLE) with the method “lstsq” from the Qiskit Ignis API to reconstruct quantum states. The “lstsq” method first computes the least-squares estimate of the density matrix, then applies the technique of [49] to impose positive semidefiniteness.

Overall we perform complete state tomography 240 times (120 with no intermediate circuits, and 120 random circuits), 5000 shots each. We then use these measurement results to reconstruct the corresponding density matrices and evaluate their purity. Histograms of the measured purity distributions found from the reconstructed states with no gates and with random instances of the circuit shown in Fig. 10 are shown, respectively, in Figs. 11 and 12. The mean purity of IBM Q measured data in Figs. 11 and 12 is 0.98 and 0.75, respectively. For comparison, we also show the ($D = 16$)-dimensional Bures and HS distributions, and the $D = K = 16$ MA distribution with mean purity tuned to match the data sets obtained from IBM Q. Each simulated histogram was created from 10^5 random samples.

To quantitatively compare the distributions in Figs. 11 and 12, we also determine the overlap of the histograms in terms of the Bhattacharyya coefficients. In particular, we find

$$\mathcal{B}(h_{IBM1}, h_{HS}) = \mathcal{B}(h_{IBM1}, h_B) = 0,$$




FIG. 11. Histograms comparing the Bures, HS, and MA distributions to the measured distribution of purity from the IBM Q for four-qubit quantum circuits initialized at $|0\rangle$ with no gates before tomography. The MA distribution has $D = K = 16$ and α is tuned so that the mean of the distribution matches that of the measured IBM Q distribution.




FIG. 12. Histograms comparing the Bures, HS, and MA distributions to the measured distribution of purity from the IBM Q for states resulting from random instances of the circuit shown in Fig. 10. The MA distribution has $D = K = 16$ and α is tuned so that the mean of the distribution matches that of the measured IBM Q distribution.




FIG. 13. (a) Schematic of the setup used to characterize the distribution of quantum states generated by a polarization-entangled photon source (EPS). DSF: dispersion-shifted fiber; DS: detector station; PA: polarization analyzer consisting of several waveplates (red) and a polarizer (blue); SPD: single-photon detector. (b) Mean density matrix resulting from performing quantum state tomography on the entangled state output by the EPS 1000 consecutive times.

$$\begin{aligned} \mathcal{B}(h_{IBM2}, h_{HS}) &= \mathcal{B}(h_{IBM2}, h_B) = 0, \\ \mathcal{B}(h_{IBM1}, h_{MA}) &= 0.92, \\ \mathcal{B}(h_{IBM2}, h_{MA}) &= 0.55. \end{aligned} \quad (30)$$

In other words, after generating 10^5 random states from both the Bures and HS distributions, neither resulted in a single state in the same bin as any of the measured states in either Figs. 11 or 12. We also see a very high overlap between the MA distribution and the measured data in Fig. 11, where the average state has very high purity.

B. Polarization-entangled photons

A schematic diagram of the experimental setup used to characterize the distribution of quantum states generated by a polarization-entangled photon source is shown in Fig. 13(a). The setup consists of an entangled-photon source (EPS) [50] connected to two separate detector stations (DS) with telecom optical fibers. The EPS creates signal and idler photons via four-wave mixing [51] by pumping a dispersion-shifted fiber (DSF) with a 50 MHz pulsed fiber laser that operates at 1552.52 nm. The DSF is stored in a laboratory-grade freezer at -86°C in order to reduce the generation of Raman-scattered noise photons. The average number of generated photon pairs per pulse can be tuned in the 0.001 – 0.1 range but is set at ~ 0.1 here. By arranging the DSF in a Sagnac loop with a polarizing beam splitter (PBS), the signal and idler photons are entangled in polarization [52]. The signal and idler photons are then spectrally demultiplexed into 100 GHz-spaced ITU outputs, resulting in photons with a temporal duration of about 15 ps. For this experiment, we use channels 28 (1554.94 nm) and 34 (1550.12 nm).

The detector stations include polarization analyzers (PA) and gated single-photon detectors (SPDs) with a detection efficiency of $\eta \sim 20\%$ and a dark count probability of $\sim 4 \times 10^{-5}$ per gate. Automated FPGA-based software controls the detectors and analyzers in order to perform full polarization state tomography from measurements in 36 different bases. Each of the 36 measurements is performed over 10 million detector gates, resulting in up to several thousand detected coincidences per measurement. The density matrix is then reconstructed using MLE [53]. The time required to perform all 36 measurements and MLE is about 12 s. Tomography is

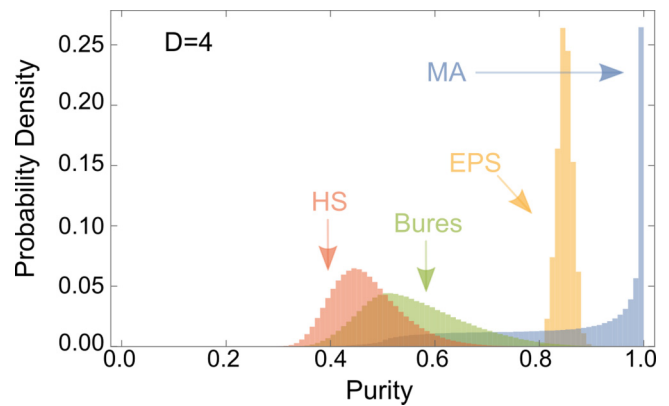


FIG. 14. Histograms comparing the Bures, HS, and MA distributions to the measured distribution of purity from a two-qubit entangled photon source (EPS). The MA distribution has $D = K = 4$ and α is tuned so that the mean of the distribution matches that of the measured EPS distribution.

performed 1000 consecutive times, resulting in a total experiment duration of ~ 4 hr. The mean density matrix is shown in Fig. 13(b). The purity is calculated for all 1000 density matrices. Analogously to Figs. 11 and 12, we plot in Fig. 14 the histograms of the measured data, the $D = 4$ HS and Bures distributions, and the $D = K = 4$ MA distribution with mean purity tuned to match the EPS distribution.

To make a quantitative comparison between the various distributions in Fig. 14, we have also calculated the Bhattacharyya coefficients to be

$$\begin{aligned} \mathcal{B}(h_{EPS}, h_{HS}) &= 0.006, & \mathcal{B}(h_{EPS}, h_B) &= 0.1, \\ \mathcal{B}(h_{EPS}, h_{MA}) &= 0.3. \end{aligned} \quad (31)$$

These results indicate that the MA distribution is closer to the measured EPS distribution than either the Bures or HS distances in terms of the Bhattacharyya coefficients.

V. DISCUSSION

From a purely theoretical standpoint, the idea of performing inference with a prior (or training set) that does not correspond to the actual distribution of validation states seems unwarranted; after all, why should one intentionally select a prior that does not match the quantum states under investigation? And as we explored in Sec. IV, the flexibility of the MA distribution permits construction of priors or training sets that are well-tailored to an experimental system, so that such a situation of mismatched priors can be at least partially mitigated. Yet in some cases, such detailed prior specification may be undesirable. For example, if one has strong beliefs that a quantum system produces pure maximally entangled states—after all, that was its design—it would be unwise to impose these beliefs on the prior in the context of state tomography: The goal is to *show* that this belief is true from subsequent measurements, not *assume* it *a priori*. Accordingly, selection of an appropriate prior or training set might be motivated less by experimenter beliefs and more by the goal of providing a generic benchmark designed to let experiments guide the posterior distribution to the ground truth.

In this practical sense, our examples highlight the value of exploring nontraditional priors and training sets for quantum state tomography. Even if lacking some of the theoretical properties of archetypal quantum state measures (like Bures), custom distributions like the MA distribution can attain performance comparable to these measures over the general Hilbert space—thus ensuring good fiducial uniformity as a test distribution—while enabling improvements for specific subspaces (e.g., pure states). Such performance hedging for desired outcomes bears resemblance to the recent estimation approach of classical shadows [54], which, when compared to BME, accepts much higher estimation error on average in exchange for remarkably low error for specific cases of interest [55].

Interestingly, we have demonstrated measurable performance improvements by matching the MA distribution to only a single feature of the underlying distribution, the mean purity. A topic of future research would be to determine how much further these results can be improved by matching the underlying distribution in more sophisticated ways, such as aiming to maximize overlap rather than merely matching means.

The apparent recovery of the HS distribution by a properly tuned MA distribution offers, to our knowledge, a new and relatively simple physical interpretation of the HS distribution. Traditionally, the HS distribution is motivated as naturally induced by tracing out Haar-random states of a higher dimension. Our results indicate that alternatively, one can think of a D -dimensional HS distribution resulting from a sum of D Haar-random pure states with appropriate Dirichlet weights. It would be interesting to determine if the distributions found via other induced measures can also be interpreted similarly by finding an appropriate K and Dirichlet weights.

ACKNOWLEDGMENTS

A portion of this work was performed at Oak Ridge National Laboratory, operated by UT-Battelle for the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. J.M.L. acknowledges funding by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, through the Quantum Algorithm Teams and Early Career Research Programs. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. Additionally, we acknowledge use of the IBM Quantum for this work. The views expressed are those of the authors and do not reflect the official policy or position of IBM Quantum. This material is based upon work supported by, or in part by, the Army Research Laboratory and the Army Research Office under Contract/Grants No. W911NF-19-2-0087 and No. W911NF-20-2-0168. T.A.S. acknowledges support from the IBM-HBCU Quantum Center and the Martin Luther King Visiting Scholars Program at MIT.

APPENDIX

We can create random density matrices according to [30] by taking \mathbf{x} of length D , creating a diagonal matrix from it, and rotating it with a Haar-random unitary from $U(D)$. The purity of such a state is given by

$$\text{Tr}(\rho^2) = \sum_{j=1}^D x_j^2, \quad (\text{A1})$$

and the expectation value of the purity by

$$\mathbb{E}_Z[\text{Tr}(\rho^2)] = \sum_{j=1}^D \mathbb{E}[x_j^2] = \frac{1 + \beta}{1 + D\beta}, \quad (\text{A2})$$

where we have used β as the concentration parameter of the Dirichlet distribution so as not to be confused with the MA expressions (where α was used).

The β values, which tune the mean purity of this distribution to match that of the Bures, HS, and MA distributions [analogous to Eq. (19) for the MA distribution], are given by

$$\beta_{MA} = \frac{\alpha(D-1)}{(\alpha+1)D}, \quad \beta_B = \frac{(D-1)(2D-1)}{3D(D+1)},$$

$$\beta_{HS} = \frac{D-1}{D+1}. \quad (\text{A3})$$

In Fig. 15, we overlay the numerically sampled and smoothed probability density functions of the Bures, HS, MA, and

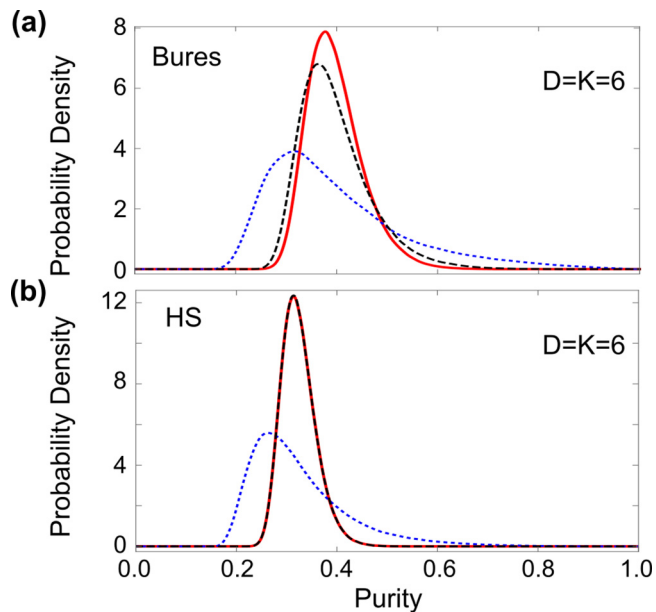


FIG. 15. Comparison of the Bures, HS, MA, and Życzkowski distributions all with the same mean purity for $D = K = 6$. Smoothed probability density functions are created from sampling each distribution 10^6 times. An analogous plot over several D values is shown in Fig. 3. (a) The solid (red), dashed (black), and dotted (blue) lines correspond to the Bures, MA, and Życzkowski distributions, respectively. (b) The solid (red), dashed (black), and dotted (blue) lines correspond to the HS, MA, and Życzkowski distributions, respectively.

Życzkowski distributions for $D = 6$ ($K = D$ for MA). These plots are generated using 10^6 random samples from each distribution, and the HS (solid red), Bures (solid red), and MA (dashed black) lines are equivalent to the $D = 6$ curves in Fig. 3. In Fig. 15(a), we compare the distribution of Życzkowski (blue dotted) and MA (dashed black) distributions against the Bures distribution (solid red) when all have the same mean purity. Similarly, in Fig. 15(b) we compare the distribution of Życzkowski (blue dotted) and MA (dashed black) distributions against the HS distribution (solid red) when all have the same mean purity. We see that, unlike the MA distribution, the distribution of Życzkowski does not reproduce either the HS or Bures distributions. Further, the mode of Życzkowski's distribution appears to be more mixed than any of the other three distributions for the same mean purity for the parameters plotted in Fig. 15.

Finally, as was performed in Sec. IV, we fit Życzkowski's distribution by way of the mean purity to the data obtained in our experimental scenarios and compare with the

MA distribution. Adopting the same notation as in Sec. IV we write the vector of histogram bins as \mathbf{h}_s , where $s \in \{MA, B, HS, Z, IBM1, IBM2, EPS\}$ labels the distribution, with the new addition of Z for Życzkowski. Calculating the Bhattacharyya coefficients we find

$$\frac{\mathcal{B}(\mathbf{h}_{IBM1}, \mathbf{h}_{MA})}{\mathcal{B}(\mathbf{h}_{IBM1}, \mathbf{h}_Z)} \approx \frac{\mathcal{B}(\mathbf{h}_{IBM2}, \mathbf{h}_{MA})}{\mathcal{B}(\mathbf{h}_{IBM2}, \mathbf{h}_Z)} \approx 1$$

$$\frac{\mathcal{B}(\mathbf{h}_{EPS}, \mathbf{h}_{MA})}{\mathcal{B}(\mathbf{h}_{EPS}, \mathbf{h}_Z)} = 1.2 \quad (\text{A4})$$

In these expressions a value above one means that the MA distribution has a higher overlap with the experimental data, and a value of one means the two distributions are similar. Hence, in all three cases we find that the MA distribution fits as well or slightly better than the distribution of Życzkowski as measured by the Bhattacharyya coefficient when tuning to match mean purity.

-
- [1] A. Montanaro, On the distinguishability of random quantum states, *Commun. Math. Phys.* **273**, 619 (2007).
- [2] A. Hama, S. Santra, and P. Zanardi, Quantum Entanglement in Random Physical States, *Phys. Rev. Lett.* **109**, 040502 (2012).
- [3] F. M. Miatto, K. Piché, T. Brougham, and R. W. Boyd, Recovering full coherence in a qubit by measuring half of its environment, *Phys. Rev. A* **92**, 062331 (2015).
- [4] D. Girolami and G. Adesso, Quantum discord for general two-qubit states: Analytical progress, *Phys. Rev. A* **83**, 052108 (2011).
- [5] B. T. Kirby, S. Santra, V. S. Malinovsky, and M. Brodsky, Entanglement swapping of two arbitrarily degraded entangled states, *Phys. Rev. A* **94**, 012336 (2016).
- [6] X.-M. Lu, J. Ma, Z. Xi, and X. Wang, Optimal measurements to access classical correlations of two-qubit states, *Phys. Rev. A* **83**, 012327 (2011).
- [7] M. Roncaglia, A. Montorsi, and M. Genovese, Bipartite entanglement of quantum states in a pair basis, *Phys. Rev. A* **90**, 062303 (2014).
- [8] S. Lu, S. Huang, K. Li, J. Li, J. Chen, D. Lu, Z. Ji, Y. Shen, D. Zhou, and B. Zeng, Separability-entanglement classifier via machine learning, *Phys. Rev. A* **98**, 012315 (2018).
- [9] S. Lohani, B. T. Kirby, M. Brodsky, O. Danaci, and R. T. Glasser, Machine learning assisted quantum state estimation, *Mach. Learn.: Sci. Technol.* **1**, 035007 (2020).
- [10] O. Danaci, S. Lohani, B. Kirby, and R. T. Glasser, Machine learning pipeline for quantum state estimation with incomplete measurements, *Mach. Learn.: Sci. Technol.* **2**, 035014 (2021).
- [11] S. Lohani, T. A. Searles, B. T. Kirby, and R. T. Glasser, On the experimental feasibility of quantum state reconstruction via machine learning, *IEEE Trans. Quantum Eng.* **2**, 2103410 (2021).
- [12] S. Ahmed, C. S. Muñoz, F. Nori, and A. F. Kockum, Classification and reconstruction of optical quantum states with deep neural networks, *Phys. Rev. Research* **3**, 033278 (2021).
- [13] R. Blume-Kohout, Optimal, reliable estimation of quantum states, *New J. Phys.* **12**, 043034 (2010).
- [14] Y.-L. Seah, J. Shang, H. K. Ng, D. J. Nott, and B.-G. Englert, Monte Carlo sampling from the quantum state space. II, *New J. Phys.* **17**, 043018 (2015).
- [15] C. Granade, J. Combes, and D. G. Cory, Practical Bayesian tomography, *New J. Phys.* **18**, 033024 (2016).
- [16] B. P. Williams and P. Lougovski, Quantum state estimation when qubits are lost: a no-data-left-behind approach, *New J. Phys.* **19**, 043003 (2017).
- [17] T. T. Mai and P. Alquier, Pseudo-Bayesian quantum tomography with rank-adaptation, *J. Stat. Plan. Inference* **184**, 62 (2017).
- [18] J. M. Lukens, K. J. H. Law, A. Jasra, and P. Lougovski, A practical and efficient approach for Bayesian quantum state estimation, *New J. Phys.* **22**, 063038 (2020).
- [19] H.-H. Lu, E. M. Simmerman, P. Lougovski, A. M. Weiner, and J. M. Lukens, Fully Arbitrary Control of Frequency-Bin Qubits, *Phys. Rev. Lett.* **125**, 120503 (2020).
- [20] W. K. Wootters, Random quantum states, *Found. Phys.* **20**, 1365 (1990).
- [21] K. Życzkowski and H.-J. Sommers, Induced measures in the space of mixed quantum states, *J. Phys. A: Math. Gen.* **34**, 7111 (2001).
- [22] H.-J. Sommers and K. Życzkowski, Bures volume of the set of mixed quantum states, *J. Phys. A: Math. Gen.* **36**, 10083 (2003).
- [23] K. Życzkowski and H.-J. Sommers, Average fidelity between random quantum states, *Phys. Rev. A* **71**, 032313 (2005).
- [24] V. Al Osipov, H.-J. Sommers, and K. Życzkowski, Random Bures mixed states and the distribution of their purity, *J. Phys. A: Math. Theor.* **43**, 055302 (2010).
- [25] T. T. Mai and P. Alquier, A Bayesian approach for noisy matrix completion: Optimal rate under general sampling distribution, *Electron. J. Stat.* **9**, 823 (2015).

- [26] V. Cottet and P. Alquier, 1-bit matrix completion: PAC-Bayesian analysis of a variational approximation, *Mach. Learn.* **107**, 579 (2018).
- [27] N. B. Lingaraju, H.-H. Lu, S. Seshadri, D. E. Leaird, A. M. Weiner, and J. M. Lukens, Adaptive bandwidth management for entanglement distribution in quantum networks, *Optica* **8**, 329 (2021).
- [28] M. Alshowkan, B. P. Williams, P. G. Evans, N. S. Rao, E. M. Simmerman, H.-H. Lu, N. B. Lingaraju, A. M. Weiner, C. E. Marvinney, Y.-Y. Pai, B. J. Lawrie, N. A. Peters, and J. M. Lukens, A reconfigurable quantum local area network over deployed fiber, *PRX Quantum* **2**, 040304 (2021).
- [29] A. K. Ekert, C. M. Alves, D. K. L. Oi, M. Horodecki, P. Horodecki, and L. C. Kwek, Direct Estimations of Linear and Nonlinear Functionals of a Quantum State, *Phys. Rev. Lett.* **88**, 217901 (2002).
- [30] K. Życzkowski, Volume of the set of separable states. II, *Phys. Rev. A* **60**, 3496 (1999).
- [31] C. A. Fuchs and J. Van De Graaf, Cryptographic distinguishability measures for quantum-mechanical states, *IEEE Trans. Inf. Theory* **45**, 1216 (1999).
- [32] I. Bengtsson and K. Życzkowski, *Geometry of Quantum States: An Introduction to Quantum Entanglement* (Cambridge University Press, Cambridge, 2017).
- [33] J. L. Alonso, J. Clemente-Gallardo, J. C. Cuchí, P. Echenique, and F. Falceto, Ehrenfest dynamics is purity non-preserving: A necessary ingredient for decoherence, *J. Chem. Phys.* **137**, 054106 (2012).
- [34] M. Telgarsky, Dirichlet draws are sparse with high probability, [arXiv:1301.4917](https://arxiv.org/abs/1301.4917).
- [35] R. Bhatia and C. Davis, A better bound on the variance, *Am. Math Mon.* **107**, 353 (2000).
- [36] J. Ginibre, Statistical ensembles of complex, quaternion, and real matrices, *J. Math. Phys.* **6**, 440 (1965).
- [37] H.-J. Sommers and K. Życzkowski, Statistical properties of random density matrices, *J. Phys. A: Math. Gen.* **37**, 8457 (2004).
- [38] A. A. Melnikov, H. P. Nautrup, M. Krenn, V. Dunjko, M. Tiersch, A. Zeilinger, and H. J. Briegel, Active learning machine learns to create new quantum experiments, *Proc. Natl. Acad. Sci. USA* **115**, 1221 (2018).
- [39] G. Sentís, M. Guţă, and G. Adesso, Quantum learning of coherent states, *EPJ Quantum Technol.* **2**, 1 (2015).
- [40] C. Harney, S. Pirandola, A. Ferraro, and M. Paternostro, Entanglement classification via neural network quantum states, *New J. Phys.* **22**, 045001 (2020).
- [41] K. Bharti, T. Haug, V. Vedral, and L.-C. Kwek, Machine learning meets quantum foundations: A brief survey, *AVS Quantum Sci.* **2**, 034101 (2020).
- [42] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin *et al.*, TensorFlow: Large-scale machine learning on heterogeneous systems, 2015, Software available from [tensorflow.org](https://www.tensorflow.org).
- [43] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press, Cambridge, 2003).
- [44] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods* (Springer, New York, 1999).
- [45] F. Mezzadri, How to generate random matrices from the classical compact groups, *Not. Am. Math. Soc.* **54**, 592 (2007).
- [46] S. L. Cotter, G. O. Roberts, A. M. Stuart, and D. White, MCMC methods for functions: Modifying old algorithms to make them faster, *Stat. Sci.* **28**, 424 (2013).
- [47] P. A. Lewis, E. McKenzie, and D. K. Hugus, Gamma processes, Technical Report, Naval Postgraduate School, 1986.
- [48] M. S. Anis, H. Abraham, A. Offei, R. Agarwal, G. Agliardi, M. Aharoni, I. Y. Akhalwaya, G. Aleksandrowicz, T. Alexander, M. Amy *et al.*, Qiskit: An open-source framework for quantum computing, 2021, qiskit.org.
- [49] J. A. Smolin, J. M. Gambetta, and G. Smith, Efficient Method for Computing the Maximum-Likelihood Quantum State from Measurements with Additive Gaussian Noise, *Phys. Rev. Lett.* **108**, 070502 (2012).
- [50] NuCrypt, “Quantum optical instrumentation,” available at <http://nucrypt.net/quantum-optical-instrumentation.html>
- [51] M. Fiorentino, P. L. Voss, J. E. Sharping, and P. Kumar, All-fiber photon-pair source for quantum communications, *IEEE Photonics Technol. Lett.* **14**, 983 (2002).
- [52] S. X. Wang and G. S. Kanter, Robust multiwavelength all-fiber source of polarization-entangled photons with built-in analyzer alignment signal, *IEEE J. Sel. Top. Quantum Electron.* **15**, 1733 (2009).
- [53] J. B. Altepeter, E. R. Jeffrey, and P. G. Kwiat, Photonic state tomography, *Adv. At. Mol. Opt. Phys.* **52**, 105 (2005).
- [54] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [55] J. M. Lukens, K. J. H. Law, and R. S. Bennink, A Bayesian analysis of classical shadows, *npj Quantum Info.* **7**, 113 (2021).