Adaptive quantum state tomography with iterative particle filtering

Syed Muhammad Kazim · Ahmad Farooq · Junaid ur Rehman · Hyundong Shin

Abstract Several Bayesian estimation based heuristics have been developed to perform quantum state tomography (QST). Their ability to quantify uncertainties using region estimators and include a priori knowledge of the experimentalists makes this family of methods an attractive choice for QST. However, specialized techniques for pure states do not work well for mixed states and vice versa. In this paper, we present an adaptive particle filter (PF) based QST protocol which improves the scaling of fidelity compared to non-adaptive Bayesian schemes for arbitrary multi-qubit states. This is due to the protocol's unabating perseverance to find the states' diagonal bases and more systematic handling of enduring problems in popular PF methods relating to the subjectivity of informative priors and the invalidity of particles produced by resamplers. Numerical examples and implementation on IBM quantum devices demonstrate improved performance for arbitrary quantum states and the application readiness of our proposed scheme.

Keywords Quantum State Tomography · Multivariate Gaussian Distribution · Particle Filtering

Syed Muhammad Kazim
kazim@khu.ac.kr

Ahmad Farooq
ahmadfarooq@khu.ac.kr

Junaid ur Rehman
junaid@khu.ac.kr

Hyundong Shin (Corresponding Author)
hshin@khu.ac.kr

Department of Electronics and Information Convergence Engineering, Kyung Hee University,
Yongin-si, 17104 Korea

1 Introduction

With increasing research and development in the fields of quantum information and computing, accurate identification and estimation of system components has become vital in efficient operation of practical quantum devices. Quantum state tomography (QST) is one such technique, which is used to characterize an unknown given quantum state ρ. More formally, QST statistically reconstructs a density matrix ρ using data generated by measuring a large number N of identically prepared copies of ρ.

Generally, the reliability and accuracy of the estimated density matrix can be improved by increasing N. Due to the resource-intensive nature of QST, a key figure-of-merit of any QST technique is its scaling with respect to N, which can be improved by measuring ρ is some optimal measurement basis. Therefore, special interest in QST has been given to the development of optimal basis for measurement [1], [2]. Improvement in measurement precision has been experimentally demonstrated when projectors of the eigenstates of Pauli operators employed in standard measurement strategies are supplanted by mutually unbiased bases (MUB) [3]. Despite the success of MUBs, it has been argued that static approaches utilizing a fixed set of measurements do not take advantage of the information obtained from measurements during QST [4]. Therefore, adaptive realizations of QST are better positioned to reduce redundancy.

It has been shown analytically that the worst-case infidelity can be reduced to $O\left(1/N\right)$ if ρ is measured in its diagonal basis [5]. Since ρ is unknown, its diagonal basis is also a priori unknwon. Adaptive techniques for QST have been developed in the past decade, which attempt to approach the diagonal basis of ρ and employ it as a measurement basis. For example, popular maximum likelihood estimation (MLE) based adaptive schemes take a two-pronged approach: perform standard tomography on $N/2$ copies of ρ, estimate an intermediate state $\hat{\rho}$, and measure the remaining $N/2$ copies in the diagonal basis of $\hat{\rho}$ [6], [7]. This two-stage tomography using MLE provides improved estimates of ρ for any value of N compared to other standard procedures. However, MLE has inherent issues as a statistical estimator for QST. Although it is possible to reconstruct states using MLE, its incompatibility with error bars makes it difficult to explain the uncertainty and thus the reliability of the estimates. Moreover, as MLE is primarily a frequentist construct, it fits observed frequencies obtained from essentially probabilistic measurements to probabilities [8]. In case of small data sets, MLE can be very unreliable.

Bayesian QST is an alternative to the popular MLE-based QST, which does not suffer from the fundamental drawbacks of MLE. In Bayesian QST, we augment our a priori knowledge of ρ with a data driven likelihood function to produce a well-defined posterior distribution. The posterior distribution allows us to make statistical estimates about ρ and evaluate quantifiable error bars [8–10]. Moreover, unlike the two-stage MLE, we do not need to know N prior to our experiment.

Adaptive Bayesian schemes have been employed to good effect in QST. Adaptive Bayesian quantum tomography [4] employs a particle filter (PF) based approach, and uses an information-theoretic utility function which optimizes measurements in each iteration. Self-guided quantum tomography (SGQT) [11] is another technique that uses optimization to solve QST for pure states, and has reported considerable improvement in the estimation of pure states. More recently, practical adaptive quantum tomography (PAQT) [12], a hybrid of PF and SGQT, has demonstrated that SGQT can be applied to mixed states with good effect. PAQT also reports improvements in fidelity over contemporary PF based QST [4] for states regardless of purity. Yet, SGQT still outperforms PAQT for pure states. Therefore, despite significant advancements in Bayesian QST, there is still a need for a single technique that is equally adept at estimation of both pure and mixed states, and provides scaling better than or similar to other formulations. That is, given a random qubit, we should not have to choose between different techniques based on, a difficult to justify, a priori assumption of the state's purity.

In this paper, we present a unified and adaptive practical technique, and report improvements in estimation over existing heuristics for random states of one, two and three qubits. Moreover, we also develop a state specific pseudo-prior such that the performance of QST is not incumbent on a priori insight of the experiment. Furthermore, we provide a resampling algorithm that corrects the propensity of creating invalid particles of popular resampling techniques. Lastly, we demonstrate the practical nature of the proposed scheme in the estimation of pure and mixed states by providing a proof-of-concept implementation on IBM's quantum computers [13].

This paper is structured as follows. We delve into the analysis of our prior, explain and provide a pseudocode for resampling of particles, and expound on the functionality of our adaptive protocol in Section 2. We exemplify our protocol in Section 3, and discuss and conclude in Sections 4 and 5, respectively.

2 Methodology

An arbitrary quantum state can be represented by a positive matrix of unit trace, i.e., a density matrix, commonly denoted by ρ. A general system of n qubits, ρ can be represented in terms of its Bloch vector $\boldsymbol{r}$ [14]

$$\rho = \frac{1}{d} \left(\mathbb{I} + \boldsymbol{r} \cdot \boldsymbol{\sigma} \right), \tag{1}$$

where $d = 2^n$, $\mathbb{I}$ is the $d \times d$ identity matrix, $\boldsymbol{r} = (r_1, r_2, \cdots, r_{d^2-1}) \in \mathbb{R}^{d^2-1}$ is a Bloch vector, and $\boldsymbol{\sigma} = (\boldsymbol{\sigma}_1, \boldsymbol{\sigma}_2, \cdots \boldsymbol{\sigma}_{d^2-1})$ is a vector of Pauli words i.e. a tensor of two-dimensional Pauli operators. Normalization and positivity of ρ translate to the norm constraints on the Bloch vector $\boldsymbol{r}$, i.e., $\|\boldsymbol{r}\|_2 \leq 1$ where $\|\cdot\|_2$ is the Euclidean norm. Moreover, a sufficient description of valid quantum states requires ρ to be positive semidefinite and hermitian with unit trace. Throughout this paper, we denote true state and the estimated state by ρ and $\hat{\rho}$, respectively.

2.1 Overview of PF-based QST

The first common step for all QST schemes is to perform measurements on an ensemble of identically prepared copies of ρ. In the case of single qubits, when ρ is measured in configuration $\alpha \in \mathcal{A}$ where $\mathcal{A}$ is a set of informationally complete projective measurement configurations, we observe one of the two possible outcomes $|\psi_\alpha^\ell\rangle$ for $\ell \in \{+1, -1\}$. Let N_0 qubits be measured in the configuration α, and n_α^ℓ be the number of times we observe the outcome ℓ. Then, the relative frequency $\hat{f}_\alpha^\ell = \frac{n_\alpha^\ell}{N_0}$ approximates the probability of outcome $|\psi_\alpha^\ell\rangle$ defined as $P\left(|\psi_\alpha^\ell\rangle\right) = \mathrm{tr}\left(|\psi_\alpha^\ell\rangle\langle\psi_\alpha^\ell|\rho\right)$. Then what remains of QST is to best estimate the state $\hat{\rho}$ based on $\hat{f}$ using a statistical method that also specifies the uncertainty of the estimate.

In the QST implementation of Bayesian PF [15] [16], we initialize particles $\{\gamma_k\}$ for $k \in \{1, \cdots, K\}$ in the Bloch sphere that represent prospective states based on a priori knowledge or some other bias. A particle γ_k is completely described at the tth iteration by its location $\boldsymbol{r}_k \in \mathbb{R}^{d^2-1}$ and its weight $w_k^t \in \mathbb{R}$. The initial distribution of particles is known as the prior [4, 17]

$$\Pr\left(\boldsymbol{r}\right) \approx \sum_k w_k^0 \delta\left(\boldsymbol{r} - \boldsymbol{r}_k\right), \tag{2}$$

where $w_k = \frac{1}{K}$. After performing N_0 measurements in α for the tth iteration, we calculate the likelihood [18],

$$\mathcal{L}\left(\gamma_k \mid \hat{f}, \alpha\right) = N_0! \prod_\ell \frac{\mathrm{tr}\left(\mathrm{p}_a^\ell \gamma_k\right)^{N_0 \hat{f}_a^\ell}}{\left(N_0 \hat{f}_a^\ell\right)!}, \tag{3}$$

where p_a^ℓ is the projector on the eigenspace corresponding to the eigenvalue ℓ of σ_a and update the weights of $\{\gamma_k\}$ for the $(t+1)$th iteration as follows

$$w_k^{t+1} \approx w_k^t \times \mathcal{L}\left(\gamma_k \mid \hat{f}, \alpha\right), \tag{4}$$

where w_k^{t+1} is normalized such that $\sum_k w_k^{t+1} = 1$. The resulting distribution $\{w_k^{t+1}\}$ is the posterior for the tth iteration. Then the Bloch vector of $\hat{\rho}$ is the Bayesian mean estimate (BME) of the distribution, which is simply the weighted aggregate of the posterior

$$\boldsymbol{r}_{\mathrm{BME}} = \sum_k w_k^{t+1} \boldsymbol{r}_k. \tag{5}$$

However, PF-based methods suffer from weight collapse where the whole posterior is concentrated on a single particle, giving it all the weight. This situation can be avoided by using an appropriate resampler that effectively reproduces the current distribution when the disparity in the weights of the particles exceeds a predetermined threshold.

In this section, we develop a prior $\Pr\left(\boldsymbol{r}\right)$, and identify the inefficiencies of contemporary resampling algorithms and introduce steps to resolve them. We

also demonstrate the need and advantages of iterative learning of the set of measurement configurations $\mathcal{A}$ in adaptive Bayesian QST.

2.2 Prior

In Bayesian QST, informative priors [9] can reduce the required number of copies N of ρ in the process. However, the quality of state estimation can suffer if the prior is based on incorrect insight. Although, attention has been afforded to increasing the robustness of protocols which rely on a priori knowledge, this robustness usually means the eventual convergence to the true state [17]. That is, the variance in the required number of samples N to achieve the same infidelity $\mathcal{I}$ for priors of varying authenticity of insight will be substantial. This is a significant problem since in real cases (state tomography of unknown states) where measurement metrics such as infidelity cannot be calculated to ascertain the accuracy of the process at any given stage, there is complete reliance on the statistical model to gauge infidelity for any value of N. Large variances in the output reduce our trust in the model, and hence reduces its practical applicability.

To counter these problems, we need a prior with statistically quantifiable errors. For this purpose, we utilize a small fraction of N to obtain an initial rough estimate of ρ and use a statistical model that specifies our region of interest in the Bloch sphere quantifying our uncertainty in the estimate.

In our protocol, we use a Multivariate Gaussian distribution $\mathcal{N}\left(\hat{\boldsymbol{\mu}}, \boldsymbol{\Sigma}\right)$ with mean $\hat{\boldsymbol{\mu}}$ and covariance $\boldsymbol{\Sigma}$ as a prior for QST, which henceforth we refer to as our preliminary guess (PG). Initially, we assume no knowledge of the true state ρ. We prepare $\left(d^2 - 1\right) N_0$ copies of ρ where $N_0 \ll N$ and measure each Pauli word $\boldsymbol{\sigma}_j$ for $j \in \left\{1, 2, \cdots, d^2 - 1\right\}$ on N_0 copies of the state so that n_j^+ $\left(n_j^-\right)$ is the number of times we observe the outcome corresponding to the +1 (-1) eigenstate of $\boldsymbol{\sigma}_j$. Then $\hat{\boldsymbol{\mu}} = (\hat{r}_1, \hat{r}_2, \cdots, \hat{r}_{d^2-1})$ is the Bloch vector of the most likely state $\hat{\rho}$ after $\left(d^2 - 1\right) N_0$ measurements, where $\hat{r}_j = \frac{n_j^+ - n_j^-}{n_j^+ + n_j^-}$. In case $\|\hat{\boldsymbol{\mu}}\|_2 > 1$, we project it onto the surface of a $\left(d^2 - 1\right)$ dimensional ball by normalizing it. Note that ensuring unit norm is sufficient for $d = 2$, but only necessary for $d > 2$. Then for $d > 2$, $\hat{\rho}$ may still not be a valid state, but we take no further actions at this stage.

To find the covariance matrix of PG, we must quantify our uncertainty in $\hat{\boldsymbol{\mu}}$. That is, we must find the variance of the sample distribution of our estimate $\hat{r}_j$, or simply the square of the standard error of r_j for the diagonal elements of $\boldsymbol{\Sigma}$. First, note that the off-diagonal elements of $\boldsymbol{\Sigma}$ are negligible because we perform projective measurements on Pauli words. Furthermore, since we perform N_0 measurements only, we can use an unbiased estimate of the variance to evaluate the standard error. Then, the jth diagonal entry

$$\boldsymbol{\Sigma}_{jj} = \frac{\sum_{i=1}^{N_0} \left(l_i - \hat{r}_j\right)^2}{N_0 \left(N_0 - 1\right)} = \frac{n_j^+ \left(1 - \hat{r}_j\right)^2 + n_j^- \left(-1 - \hat{r}_j\right)^2}{N_0 \left(N_0 - 1\right)} + \epsilon,$$

where ϵ is a small constant defined to accommodate all ρ that are eigenstates of our measurement operators. The $\left(d^2 - 1\right) \times \left(d^2 - 1\right)$ covariance matrix $\boldsymbol{\Sigma}$ is a measure of the variances of the distances of $\hat{\rho}$ from ρ after N_0 measurements in each element of $\boldsymbol{\sigma}$. Furthermore, the diagonal elements of $\boldsymbol{\Sigma}$ are state specific. For $N_0 = 50$, when $\hat{r}_j \to \{+1, -1\}$, $\Sigma_{jj} \to 0$ and $\hat{r}_j \to 0$, $\Sigma_{jj} \to 0.02$. This state specificity is especially useful for states of high purity since it reduces our volume of interest for the same confidence level. Note that since we perform measurements to approximate parameters of $\mathcal{N}\left(\hat{\boldsymbol{\mu}}, \boldsymbol{\Sigma}\right)$ for PG, we can think of PG as a pseudo-prior rather than a conventional prior since the latter is constructed solely from a priori knowledge.

By utilizing only a fraction of copies, we have accomplished (i) an approximate state that serves as the first step for an adaptive QST technique explained later in the section, and (ii) reduced the region of interest substantially while retaining a statistical description of the uncertainty of PG. To inspect (ii), let $\hat{\boldsymbol{r}} = (\hat{r}_x, \hat{r}_y, \hat{r}_z)$ be the Bloch vector of a single qubit, and based on our previous discussion, let each element of $\hat{\boldsymbol{r}}$ be an independent and identically distributed (i.i.d) Gaussian random variable with $\mathcal{N}(r_i, \Sigma_{ii})$. An ellipsoid can be drawn to represent specific confidence levels using

$$\sum_{i \in \{x, y, z\}} \frac{(\hat{r}_i - r_i)^2}{\Sigma_{ii}} = s, \tag{6}$$

where s represents the confidence interval of a χ^2 distribution with three degrees of freedom and $s = 11.345$ for a 99% confidence interval. In the case of a maximally mixed state, worst case for our strategy, the diagonal elements of $\boldsymbol{\Sigma}$ are all equal to 0.02, and our ellipsoid corresponds to a sphere with diameter $2\sqrt{cs} = 0.135$. Therefore, compared to an uninformative prior which suggests all valid states are equiprobable, we have reduced the volume of interest by at least 99.97%. Therefore, as demonstrated in Fig. 9, by using a 'just enough' informative prior which is characterized by its uncertainty works just as well as a good informative prior, and better than one characterized by inaccurate a priori knowledge.

2.3 Resampling

In the PF-based Bayesian estimation, particles have to be resampled when $\sum_{k=1}^{K-1} w_k \approx 0$ and $w_K \approx 1$ [19]. Depending on the resampling algorithm and proximity of the previous state to the surface of the Bloch sphere, resampled particles can be invalid states. Invalid two-dimensional states can be visualized as particles not circumscribed by the Bloch sphere, or more generally, as states represented by density matrices with eigenvalues less than zero or greater than one. A common workaround is to reduce the negative eigenvalues of such particles to 0, and normalize the remaining eigenvalues to produce valid states [17]. Although this approach is simple, it deforms the resampling distribution, and increases computational expense since correction of thousands of particles can
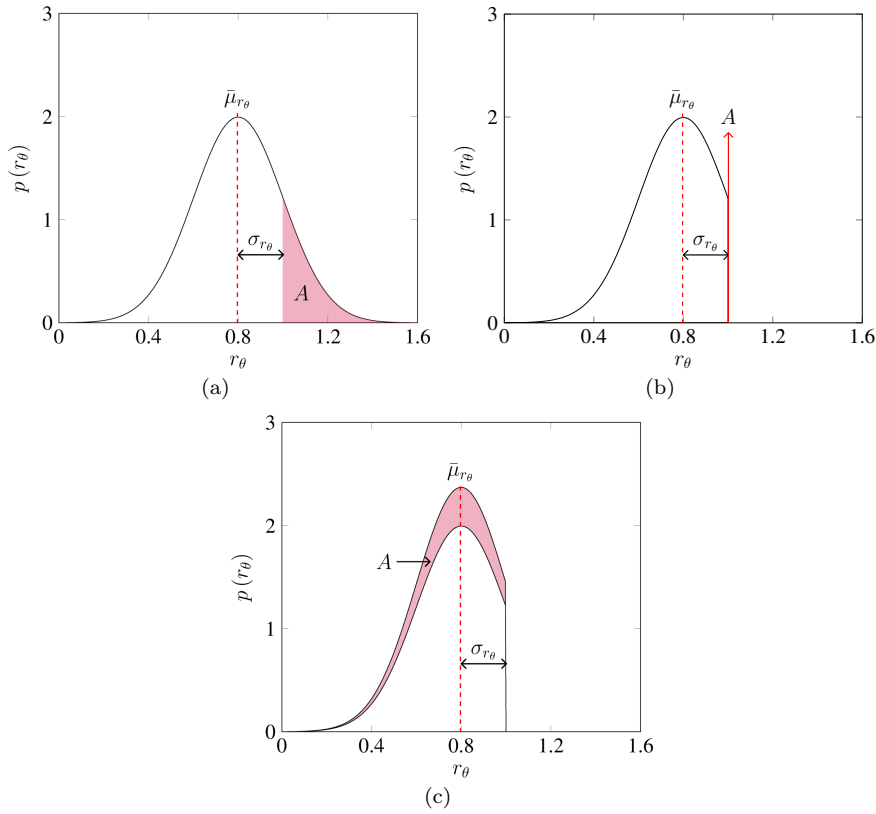
Fig. 1 1-D representation of Bloch sphere where $\theta \in \{x, y, z\}$ and $\|r_\theta\| \leq 1$. (a) Resampler with PDF $\mathcal{N}(\bar{\mu}_{r_\theta}, \sigma_{r_\theta})$. (b) Output of resampler with concurrent negative eigenvalue truncation schemes morph the PDF such that there is an impulse at $r_\theta = 1$ of size A. (c) Proposed truncated Gaussian distribution such that the resampler outputs valid states without changing spatial probability ratios of samples.

be required every time the particles are resampled. That is, when a high density of particles are close to the surface of the Bloch sphere (when estimating a pure state, or a state close by), a high number of samples produced after resampling will be invalid states, which are essentially projected onto the Bloch sphere as pure states.

More specifically, suppose that we utilize a univariate Gaussian distribution $\mathcal{N}(\bar{\mu}_{r_\theta}, \sigma_{r_\theta})$ to resample particles as shown in Fig. 1(a) where the shaded region A represents the probability of sampling an invalid state. By projecting the invalid states onto the surface of the Bloch sphere, we have essentially sampled pure states with probability A and used a distribution represented in Fig. 1(b). To accommodate the constraints of a two-dimensional valid density matrix without deforming the distribution of the particles and involving further computation required to correct invalid states, we resample using a truncated Gaussian (TG) distribution as shown in Fig. 1(c). To extend this

---

Algorithm 1: Practical resampling for valid states

Input: Particle weights $\{w_k\}$, locations $\{\boldsymbol{r}_k\}$ for $k \in \{1, \cdots, K\}$
Input: $a \in [0, 1]$
Output: Updated weights $\{w'_k\}$, locations $\{\boldsymbol{r}'_k\}$

1   $\boldsymbol{\mu} \leftarrow \text{Mean}(\{w_k\}, \{\boldsymbol{r}_k\})$
2   $h \leftarrow \sqrt{1 - a^2}$
3   $\boldsymbol{\Sigma} \leftarrow h^2 \, \text{Cov}(\{w_k\}, \{\boldsymbol{r}_k\})$
4   $[\mathbf{V}, \boldsymbol{\lambda}] \leftarrow \text{eig}\,(\boldsymbol{\Sigma})$ ▶ $\mathbf{V}$ is a matrix of eigenvectors, $\boldsymbol{\lambda}$ is a vector of eigenvalues
5   $\tau_1 = \tau_2 = \cdots = \tau_{d^2-1} = 0$
6   **for** $k \in 1 \to K$ **do**
7   draw k^{th} particle $\boldsymbol{r}_k$ with probability w_k
8      $\boldsymbol{\mu}' \leftarrow \mathbf{V}^{\mathrm{T}}\left(a\boldsymbol{r}_k + (1-a)\boldsymbol{\mu}\right)$ ▶ $\boldsymbol{\mu}' = \left(\mu'_0, \mu'_1, \cdots, \mu'_{d^2-1}\right)$
9   **for** $j \in \left\{1, 2, \cdots, d^2 - 1\right\}$ **do**
10          $C_1 \leftarrow -\sqrt{1 - \sum_{l=0}^{j} \tau_l^2}$
11          $C_2 \to -C_1$
12          $\tau_j \leftarrow \text{TG}_{[C_1, C_2]}\left(\mu'_j, \sqrt{\lambda_j}\right)$
13      $\boldsymbol{r}'_k \leftarrow \mathbf{V} \cdot (\tau_0, \tau_1, \tau_2)^{\mathrm{T}}$
14      $w'_i \leftarrow 1/K$
15   **return** $\{w'_k\}\{\boldsymbol{r}'_k\}$

---

approach to higher dimensional quantum states, we sample from a $\left(d^2 - 1\right)$ dimensional unit ball instead of the more complicated space of valid quantum states. Although this simplification precludes the ability of this algorithm to *always* sample valid states, it ensures smooth distributions for resampling. Procedure for sampling from the probability density function (PDF) of Fig. 1(c) is given in Algorithm 1. For each particle, we sample r_k for $k \in \left\{1, 2, \cdots, d^2 - 1\right\}$ from marginal Gaussian distributions corresponding to orthogonal axes, which are principal components of the current particle distribution, sequentially. We update the domain $[C_1, C_2]$ of the next PDF based on the sample such that the constraint $\|\boldsymbol{r}\| \leq 1$ is not violated. Since we sample from marginal distributions of each axis independently, we must ensure that we can extract marginal PDFs from the available joint PDF. Therefore, we use $\mathbf{V}$ to change the system's bases to the principal components of the existing particle distribution. In this way, we remove existing correlations between bases. Lastly, for each orthogonal axis, we calculate C_1 and C_2 of Gaussian distribution, and sample τ. $\boldsymbol{r}'_k$ and w'_k are the updated particle's location and weight, respectively.

2.4 Adaptive Bayesian QST

Infidelity scales as $O\left(1/N\right)$ when ρ is measured in its own diagonal basis, which is the best possible scaling for QST. However, when this is not the case, it scales as $O\left(1/\sqrt{N}\right)$ for qubits close to the surface of the Bloch sphere [5], [6]. Since ρ is unknown, the purpose of introducing adaptive protocols is to approach the
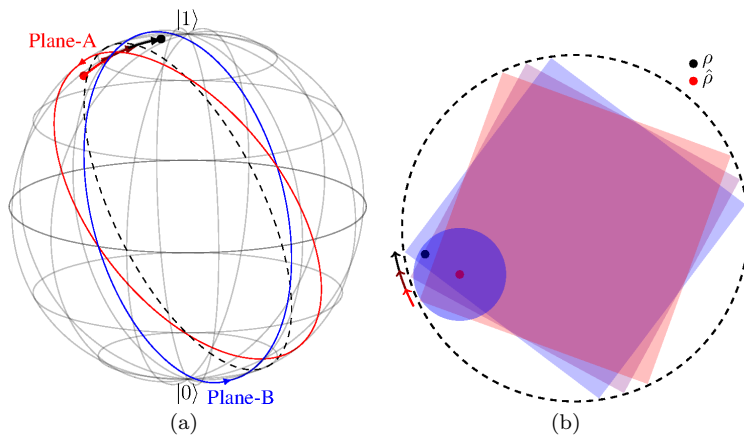
Fig. 2 (a) Rotation of planes inside Bloch sphere based on estimated $\hat{\rho}$ where on average the plane move towards ρ with increasing number of iterations. (b) Opposite vertices of inscribed squares represent eigenstates of measurement operators where the adaptive scheme rotates the square such that it moves closer and eventually encompasses the true state to improve scaling of QST.

diagonal basis without losing information gained from intermediary measurement bases. Particle filtering naturally incorporates adaptivity in QST with minimal computational overhead.

The formation of PG is conceptually the first step of our adaptive QST protocol, in that we initially perform measurements on ρ using Pauli operators, and estimate a state, $\hat{\rho}$. We then proceed to change the measurement configuration for the next set of measurements. We use the eigenbases of $\hat{\rho}$ to rotate Pauli words such that the updated operators $\Omega_1, \Omega_2, \cdots, \Omega_{d^2-1}$ maintain $\mathrm{Tr}\left(\Omega_i^\dagger \Omega_j\right) = 2\delta_{ij}$ and one of the operators diagonalizes $\hat{\rho}$. The particle filter updates the existing distribution based on the outcomes of the measurements in this configuration. We estimate an updated $\hat{\rho}$ and repeat the process iteratively until we have exhausted N or some other criterion is met.

This iterative process is analogous to rotating the plane of measurement (in a two-dimensional space) and cube (in a three-dimensional space) such that the true state eventually lies within it, as shown in Fig. 2, to achieve the best possible scaling of infidelity. The circular plane represented by a dashed line on the left in Fig. 2 is observed in the same figure on the right. The opposite corners of the squares circumscribed by the circle represent the eigenstates of the measurement operators. At a certain step in QST, state $\hat{\rho}$ is approximated. The plane of measurement is rotated such that eigenbases of one of the operators diagonalizes $\hat{\rho}$. The next set of measurements are performed in this configuration, and based on the outcomes a new plane of measurement is selected. The iterative rotation is indicated by the color of the planes in the direction corresponding to the arrows. States close to the surface of the Bloch sphere require a greater number of iterations. In the proposed scheme,
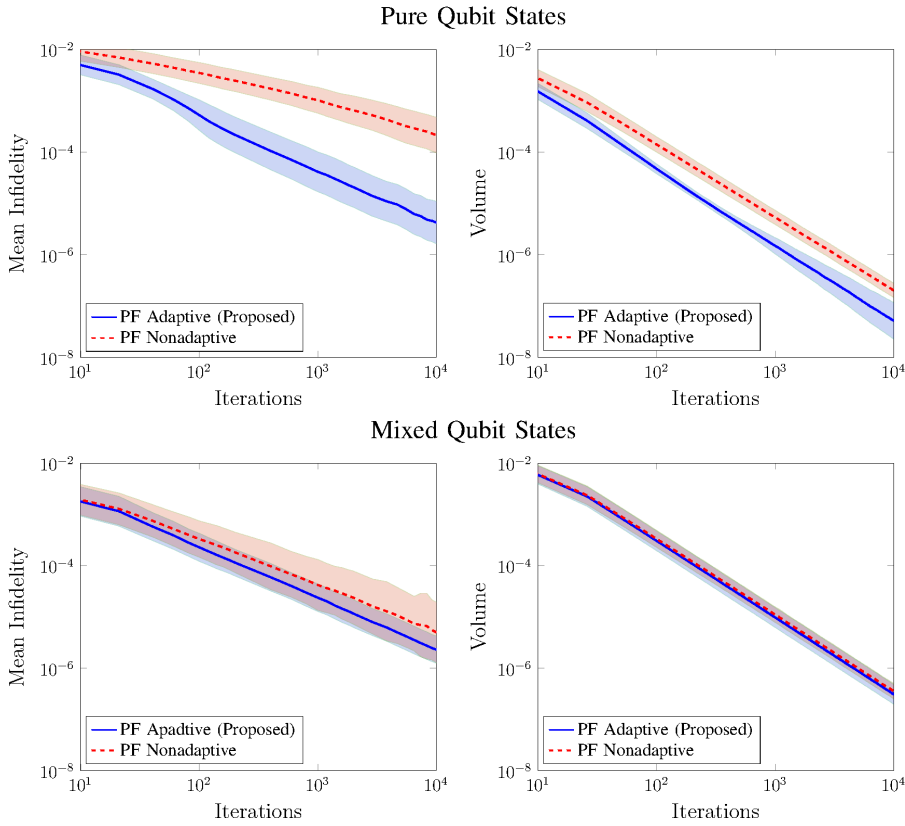
Fig. 3 Infidelity and minimum volume enclosing ellipsoid (MVEE) for 99% credible region for the proposed adaptive protocol and its non-adaptive version for pure and mixed states using the method detailed in Section 2 averaged over 1000 states. The shaded region indicates the 16% and 84% quantiles over all measurements.

we perform N_0 measurements on ρ using a single operator Ω_j per iteration, where Ω_j is chosen randomly.

3 Numerical and Experimental Results

In this section, we demonstrate the performance of our protocol detailed in Section 2. We perform QST using open-source packages Qinfer [20] and Qutip [21] on pure and mixed states using our adaptive protocol, and compare the results to non-adaptive PF algorithm. Moreover, we also perform QST on IBM quantum experience [13] to show the readily applicable nature of our work. QST is performed on ρ which is randomly sampled from Hilbert-Schmidt uniform and Haar uniform distributions for mixed and pure states, respectively, with the resampling parameter $\alpha = 0.1$. We report infidelity $\mathcal{I}$ between the estimated
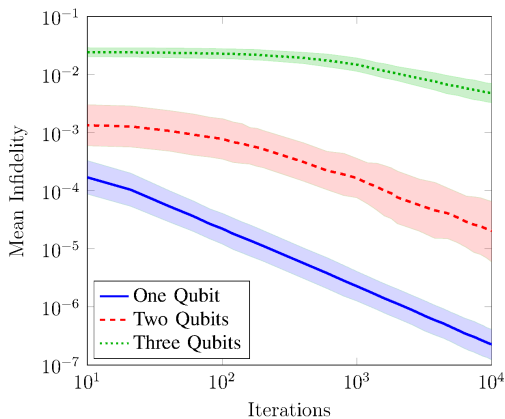
Fig. 4 Infidelity of one, two and three qubit mixed states. One qubit numerical results used 500 shots per operator for PG, whereas both two and three qubit cases used 2500 shots per operator. Each iteration consists of 500 shots. For one, two and three qubits, we used 2000, 4000 and 8000 particles, respectively. One and two qubit cases are averaged over 500 randomly sampled states from Hilbert-Schmidt uniform distribution whereas three qubit case is averaged over 50 states from the same distribution. The shaded area corresponds to ± 1 standard deviation.

states $\hat{\rho}$ and true states ρ defined as [22]

$$\mathcal{I}(\rho, \hat{\rho}) = 1 - \mathrm{tr}\left(\sqrt{\sqrt{\rho}\hat{\rho}\sqrt{\rho}}\right), \tag{7}$$

where fidelity $\mathcal{F} = \mathrm{tr}\left(\sqrt{\sqrt{\rho}\hat{\rho}\sqrt{\rho}}\right)$ so that $\mathcal{I} = 1 - \mathcal{F}$. The infidelity captures the idea of closeness between ρ and $\hat{\rho}$ such that $\mathcal{I}(\rho, \hat{\rho}) = 0$ if and only if $\rho = \hat{\rho}$. Note that $\hat{\rho}$ can be very close to ρ and still be an invalid state due to our simplification of the valid space in Section 2.3. Therefore, as a final step, we check for the validity of $\hat{\rho}$, and ensure semi-definiteness by reducing negative eigenvalues to zero, and normalizing the remaining to ensure unit trace, when required. We also report volumes of covariance $\boldsymbol{\Sigma}$-based ellipsoids enclosing 99% credible regions defined as [10], [23]

$$\mathrm{Vol}(\boldsymbol{\Sigma}) = \frac{\pi^{(d-1)/2}}{\Gamma\left(\frac{d}{2}+1\right)} \det(\boldsymbol{\Sigma})^{-1/2}, \tag{8}$$

where d is the dimension of $\boldsymbol{\Sigma}$ and Γ is the Gamma function for our simulations and experiments. This measure specifies the concentration of particles by calculating the volume enclosed by a fixed percentage of particles. Thereby helping us gauge the quality of convergence in successive iterations.

Fig. 3 demonstrates the advantage of proposed adaptive scheme as compared to non-adaptive PF-based scheme, both in terms of infidelity and the volume for 99% credible regions. The difference between the two schemes is more pronounced for pure states because of the presence of only a single eigenvalue, which reduces uncertainty in measurement in an adaptive setting when

(a)   (b)
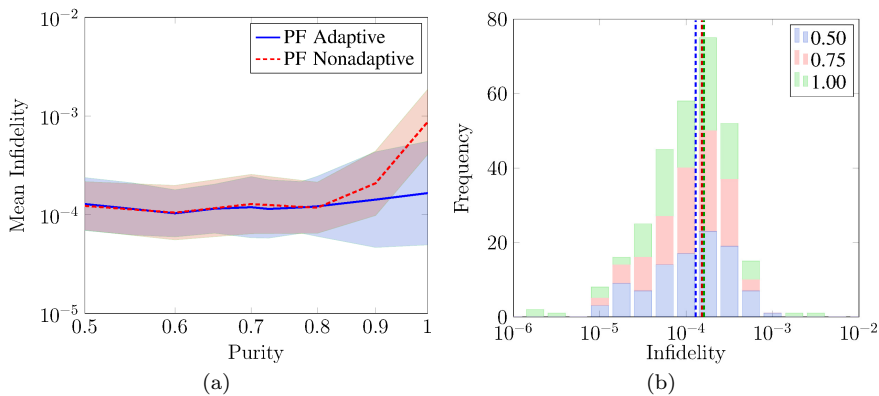
Fig. 5 Invariance under purity of target state. (a) Mean infidelity of the proposed adaptive protocol and its nonadaptive version after 10^4 measurements shown as a function of purity of the two-dimensional state to be estimated, averaged over 100 states. The shaded region corresponds to ± 1 standard deviation. (b) Stacked histogram for infidelity using adaptive protocol of 100 states with purity 0.5 (blue), 0.75 (red) and 1.00 (green) each with their corresponding means given by dashed lines of the same colours, respectively.

the measurement operator diagonalizes a $\hat{\rho}$ that is close to ρ. Fig. 4 demonstrates the infidelities of one, two and three qubit mixed states.

Fig. 5 demonstrates the performance of the protocol as a function of the purity, given by $\mathrm{tr}\left(\rho^2\right)$. For each value of purity, 100 states were randomly sampled from the Hilbert-Schmidt uniform distribution. Then given a specific purity and keeping the eigenvectors of the sampled state unchanged, eigenvalues were updated accordingly. Fig. 5(a) shows that performance of the adaptive protocol remains mostly invariant with respect to purity. Contrarily for the nonadaptive protocol, the performance declines as states become more pure. Fig. 5(b) shows the distribution of infidelity of the states after 10^4 measurements using the adaptive protocol. The distributions are very similar showing that the proposed protocol is unhindered by the state's purity.

To perform QST on a mixed state on IBM's quantum computer [24], we first prepare a two-qubit pure state $|\psi\rangle = [0.73, -0.25, 0.528, 0.348] \in H_B \otimes H_A$ so that by taking the partial trace of $|\psi\rangle\langle\psi|$ with respect to Hilbert space H_B, we attain $\rho \in H_A$ defined by the density matrix $\rho = \frac{1}{2}\left(\mathbb{I} + 0.6\sigma_x + 0.2\sigma_z\right)$, where σ_x and σ_z are Pauli X and Z respectively. At any specific iteration, we calculate the weighted aggregate $\boldsymbol{r}_{\mathrm{BME}}$ of the particle distribution to estimate $\hat{\rho}$, and rotate our measurement operators using the unitary operator $\mathcal{U}$ as detailed in Section 2 and demonstrated by the quantum circuit flow diagram in Fig. 6. We measure the first qubit and update our particle filter accordingly. We also execute this process for a pure state, and report infidelity $\mathcal{I}$ for both states in Fig. 7 for 15 iterations of 1000 shots each. In this paper we used *ibmq_athens*, which is one of the IBM Quantum Falcon Processors with average readout error and average CNOT error of $1.950e^{-2}$ and $1.064e^{-2}$, respectively.
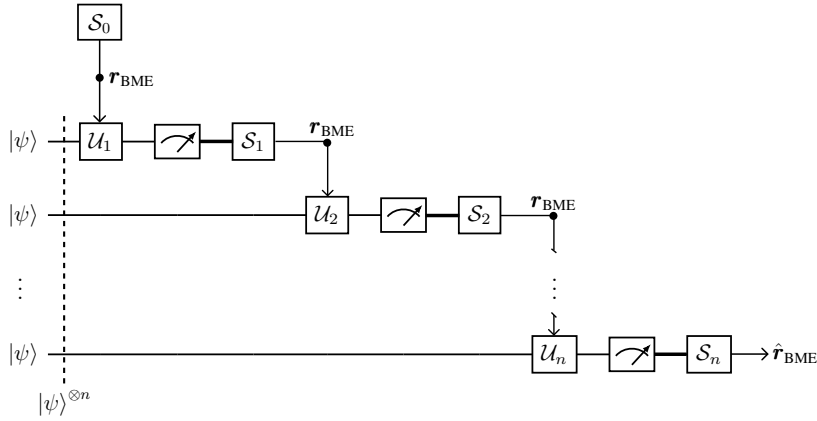
Fig. 6 The PF Adaptive QST for single-qubit state $|\psi\rangle$. The $\mathcal{U}_{i+1}$ gate is configured based on the $\boldsymbol{r}_{\mathrm{BME}}$ of PF distribution $\mathcal{S}_i$ at iteration i. $\mathcal{U}_{i+1}$ changes the basis of measurement, and $\mathcal{S}_{i+1}$ is updated based on the measurements counts. The process is initialized with PG, $\mathcal{S}_0 = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, and $\hat{\boldsymbol{r}}_{\mathrm{BME}}$ is the Bloch vector of our estimate.
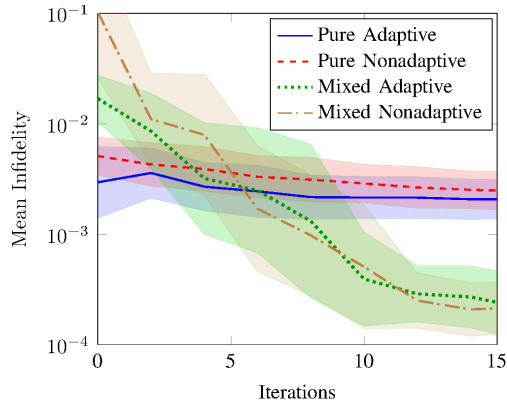
Fig. 7 Mean infidelity of pure and mixed qubit states from measurements on IBM quantum computers for the proposed scheme and its nonadaptive version, averaged over 50 states. Each iteration consists of 10^3 shots where fidelity at 0th iteration is due to PG which performs three iterations of 10^3 shots along each Pauli axis. The shaded area corresponds to ± 1 standard deviation.

In Fig. 7, the infidelity in the experimental implementation, scales better for mixed states than pure states. This difference is explained by state preparation errors of real quantum devices. Due to the presence of noise, instead of preparing the state ρ, the quantum device prepares a different random state ρ_1 in each iteration, where $\rho_1 \approx \rho$. Let's assume that $\boldsymbol{\sigma}$ diagonalizes ρ and ρ_1. If ρ is pure, then $1 - \epsilon \leq \operatorname{tr}(\rho_1 \boldsymbol{\sigma}) \leq 1$ for $\epsilon > 0$. This is why we also observe the early saturation of pure curves in Fig. 7. On the other hand, assuming state preparation errors to be inherently random, if ρ is mixed, then $\operatorname{tr}(\rho \boldsymbol{\sigma}) - \epsilon \leq \operatorname{tr}(\rho_1 \boldsymbol{\sigma}) \leq \operatorname{tr}(\rho \boldsymbol{\sigma}) + \epsilon$. Since the expectation of the prepared states can be higher or lower than that of ρ, when ρ_1 is prepared independently in

(a) Pure States
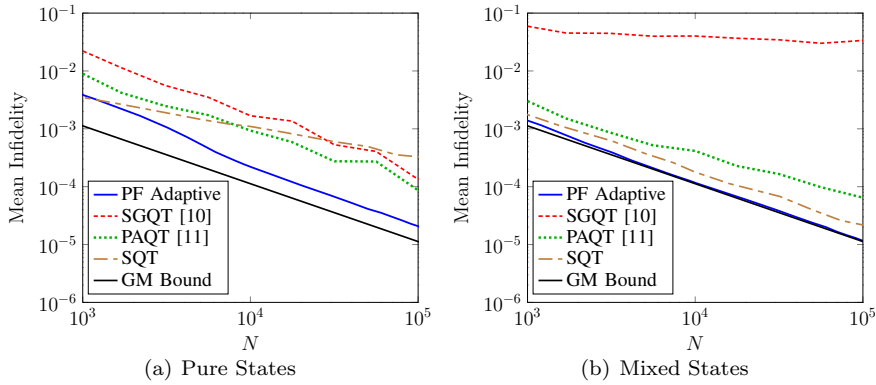
(b) Mixed States

Fig. 8 Comparison of mean infidelity of pure and mixed state estimations based on proposed protocol, PF Adaptive and its nonadaptive version, PF NA (averaged over 1000 states with 2000 particles) along with SGQT, PAQT (4000 particles) [26] and the Gill-Massar (GM) bound. Each iteration consists of 50 shots along a measurement axis defined by the protocols.

each iteration, the state preparation errors introduced in one iteration can be offset by other. Moreover, improvement in later iterations is small, which is a typical behavior of tomography experiments on noisy systems [25]. To obtain these results, we also utilized the noise mitigation offered by QISKIT module to reduce noise from quantum circuits and measurements that increased the number of measurements by a polynomial factor.

4 Discussion

Given that we set out to demonstrate that our method is adept at QST regardless of the purity of the state, we now show (in addition to Fig. 5) its advantage over other contemporary Bayesian methods. Self-guided quantum tomography (SGQT) is a method that learns pure states through 'Simultaneous Perturbation Stochastic Approximation' (SPSA) [11]. However, it does not report error bars and works poorly for mixed states. Practical adaptive quantum tomography (PAQT) [12] is a more rounded technique that builds on SGQT by applying measurements learned by SGQT on a Bayesian particle filter. PAQT can therefore report error bars, and estimate both pure and mixed states. In the process of making a unified technique, PAQT compromises on the infidelities reported by SGQT for pure states. In Fig. 8, we demonstrate the mean infidelities of SGQT, PAQT and our method for both one qubit pure and mixed states against the number of measurements N. Data used for SGQT and PAQT plots is provided by the authors in [26]. Our advantage over both techniques is visible especially when we compare there performance for states of arbitrary purity.

The accuracy of estimated quantum state can be studied by means of relevant inequalities which are Cramér-Rao inequality $C \leq F^{-1}$, quantum Cramér-Rao inequality $C \leq J^{-1}$ and the Gill-Massar (GM) ineqaulity $\mathrm{tr}\left(FJ^{-1}\right)$
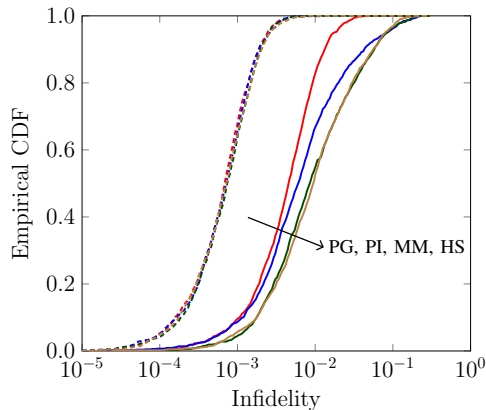
Fig. 9 CDF of infidelity of 1000 randomly sampled mixed qubits. The initial guesses of partial information (PI), maximally mixed (MM) and Hilbert-Schmidt (HS) based priors for our QST protocol are $\hat{\rho}_0 = (1 - p)\rho + \pi p$, $\frac{\mathbb{I}}{2}$, and randomly chosen state from HS distribution, respectively, where $p = 0.1$, and the initial PF distribution is uniform. PG utilized $3N_0$ samples prior to the first iteration, after which each iteration measures N_0 samples for all priors. For a fair comparison, 1st (solid) and 23rd (dashed) iterations of PG are plotted with 4th (solid) and 26th (dashed) iterations of PI, MM and HS respectively.

$\leq d - 1$, where C, F and J are covarinace, classical Fisher information and quantum Fisher information matrices, respectively [27]. These fundamental inequalities established the lower bounds for several accuracy metrics, such as mean square error and infidelity, considering the impact of the finite ensemble size on the estimation uncertainty. The GM bound for the mean squared Bures distance in d dimensional quantum system for finite number of copies N is given as [28]

$$D_B^2 = \frac{1}{4N}(d + 1)^2(d - 1),\tag{9}$$

which can be defined through fidelity

$$1 - \mathcal{F}^2(\rho, \hat{\rho}) = D_B^2 - D_B^4/4.\tag{10}$$

Mean infidelity of the proposed protocol in Fig. 8 saturates the GM bound only for mixed single qubits. Moreover, for two or more qubits, infidelities obtained do not reach the GM bound. The diminished advantage for $d > 2$ is explained by the increasing number of redundant iterations for larger systems. In the case of three qubits, after a sufficient number of iterations, only a single operator will produce outcomes which will be useful for state characterization, while the remaining 62 will produce approximately $\frac{N_0}{2}$ counts where N_0 measurements are performed in each iteration. Note that the GM bound in Fig. 8 is adjusted according to the definition of $\mathcal{I}$ provided in equation 7.

Fig. 9 compares the cumulative density function (CDF) of infidelity of our complete proposed scheme in which we utilize PG detailed in Section 2 and three modified versions of our scheme which utilize priors of varying degrees
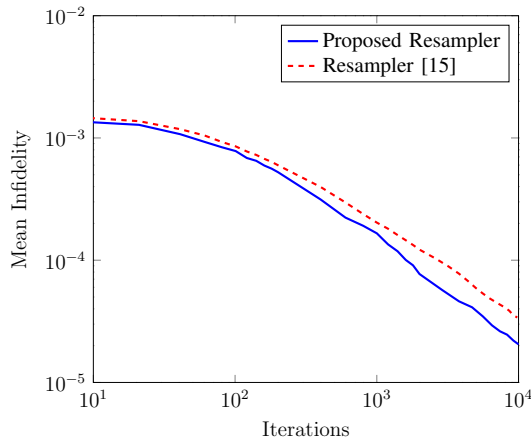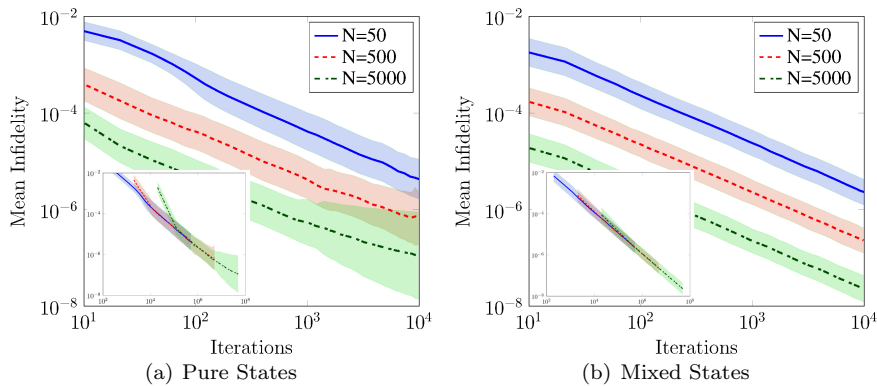
Fig. 10 Mean infidelity of the proposed protocol with truncated Gaussian resampler and the conventional Liu-West resampler, averaged over 250 random states. Resampling parameter $\alpha = 0.1$ and 4000 particles are used for state tomography of four-dimensional mixed states.

of information. Partial information (PI), maximally mixed (MM) and Hilbert-Schmidt (HS) based priors are named to allude to the process with which we have estimated the initial $\hat{\rho}_0$ required to kickstart the adaptive QST protocol. So for PI, $\hat{\rho}_0 = (1-p)\,\rho + \pi p$ where $p = 0.1$. For MM, $\hat{\rho}_0 = \mathbb{I}/d$ and for HS, $\hat{\rho}_0$ is a random sample from the Hiblert-Schmidt uniform distribution. However, since we are unsure of our uncertainty in $\hat{\rho}_0$, we have utilized a uniform particle filter distribution for each.

While strictly speaking, PG is not a prior, it provides improved characterization of the quantum states initially. This improvement can be attributed to the "free flowing" rotation of measurement operators and the concentration of particles in the proximity of the true state. In the case of PI, note that $\hat{\rho}$ shares the same eigenvectors as ρ. Therefore, in the first iteration, one measurement operator diagnoalizes ρ. This operator is chosen with probability $\frac{1}{d^2-1}$. If this operator is chosen, then the probabilities of particles close to ρ will increase in the posterior. However, since N in small ($N = 50$ is used for $d = 2$ in this paper) and the particles are uniformly distributed, the cumulative probability of particles farther away decrease the proximity of the BME from ρ. This BME will be used to rotate the set of measurement operators in the next iteration, and therefore no operator will perfectly diagnolize ρ. In the case of MM, the operators will not rotate in the first iteration. After the first measurement the likelihood of the particles near $\mathrm{tr}\,(\Omega\rho)$ will increase where $\Omega \in \{X, Y, Z\}$. Due to the spread of the particles, the operators will rotate in the next iteration but this rotation will be restricted. Lastly in the case of HS, $\hat{\rho}$ is just a randomly sampled state from the Hilbert-Schmidt uniform distribution, and such prior information cannot be expected to offer any improvement. In all three cases, the advantage reduces with a higher number of iterations (once the resampler is called and rotations become more "free") as shown by the overlapping CDFs (dashed) in Fig. 9.

(a) Pure States   (b) Mixed States

Fig. 11 Comparison of mean infidelities of the proposed adaptive protocol for pure and mixed qubit states for different N, where N is the number of shots along a measurement axis in each iteration. Each curve is averaged over 1000 randomly sampled particles and the shaded region indicates the 16% and 84% quantiles over all measurements.

Fig. 10 demonstrates the advantage of the proposed resampler over conventional resamplers [16] used in PF implementations in QST. For this figure, only the resampler is changed without making any other amendments to the adaptive scheme. Fig. 1 demonstrated that conventional resampling strategies inadvertently allocate substantial probability to the edges of the Bloch space. Therefore, a larger number of particles are required to ensure they cover the required Bloch space. The problem exacerbates for higher dimensions.

For our protocol, we used multiple values of the resampling parameter and found $\alpha = 0.1$ to be optimal. The intuition behind this is that on average the BME of the distribution at any iteration is closer to ρ than any individual particle. Therefore, when resampling, we allow the mean of the resampler's distribution to be more heavily influenced by the BME than by the individual particles.

The conventional resampler in Fig. 10 uses a Gaussian distribution to resample particles and then performs eigenvalue correction on each particle. In case we do not perform any correction to the sampled particles and considering that we sample a large number of particles (4000 for Fig. 10) multiple times for QST of a single state, it is highly likely that we will sample a significant percentage of invalid particles. Note that such a resampler samples invalid particles even for the simplest case when $d = 2$. Therefore, when using an unconstrained Gaussian distribution for the proposed protocol, some correction of states is generally required. The eigenvalue truncation and normalization used in conventional resamplers, and truncating Gaussian distributions as done in the proposed resampler are two ways to proceed. Compared to the unconstrained Gaussian resampler, TG i) always samples valid states for $d = 2$ and ii) reduces the probability of sampling invalid states by removing the possibility of sampling states outside the unit-ball. In comparison to the conventional resamplers, the advantage of TG in terms of infidelities is slight. However, TG

requires fewer computational resources. Since TG resampler is inherently sequential for each particle (although resampling of each particle is independent and therefore parallelizable), it scales as $O\left(d^2 P\right)$ while conventional resampler scales as $O\left(d^3 P\right)$ due to the requirement of eigendecomposition, where $P \gg d$ is the total number of particles. Moreover, it is interesting to note, that despite the unit-ball simplification, there are no evident drawbacks, and even states with high purity seldom require eigenvalue correction as the final step.

Fig. 11 shows the robustness of our scheme in the presence of statistical noise which can be controlled by varying N_0, the number of measurements per iteration. We find that the asymptotic scaling of infidelity is independent of N_0 [10], [29] for both pure and mixed qubit states. The initial advantage offered by smaller values of N_0 is due to the adaptive scheme. As the number of iterations increase, all curves in the inset of Fig. 11 converge.

Moreover, we have also provided a quantum circuit that works hybridly with our particle filter. Although it estimates ρ to a fair extent, it would be best if it is taken as only a proof of concept. The circuit in Fig. 6 uses the simplest techniques to prepare ρ and rotate measurement operators. More efficient circuits can be used that reduce the decoherence of the prepared state and in turn help us improve our estimates of the true state [30], [31]. Moreover, this technique applies a new measurement in each iteration. Although this requirement is fundamental for initial iterations, more attention can be afforded to its actual efficacy later on. If we can know that after certain n iterations, QST estimates $\hat{\rho}_n$ with measurement configuration k and after $n+1$ it estimates $\hat{\rho}_{n+1}$ with measurement configuration $k+1$ where the trace distance $\delta\left(\hat{\rho}_n, \hat{\rho}_{n+1}\right) \approx 0$, the computational overhead in calculating and changing to configuration $k+1$ can be avoided without significant difference to the scaling of our estimate.

5 Conclusion

We proposed an adaptive Bayesian QST technique for multiple qubits that changes the measurement basis in each iteration such that it seamlessly uses the prior as an effective first step. We have reported the numerical and experimental infidelities in our estimates and their uncertainties for arbitrary two-dimensional states. Furthermore, we have provided a comparison of infidelities with popular Bayesian particle filter methods used for QST and demonstrated our advantage in estimation over them. One prospective work can be to use the empirically derived prior to produce novel adaptive quantum tomography methods which take advantage of the maximum possible L1 error of the first approximated state. Maximum L1 error of the first estimate $\hat{\rho}$ is the maximum absolute distance between the corresponding elements of the Bloch vectors of ρ and $\hat{\rho}$. An advantage of such a method is that it allows us to evaluate lower bounds of infidelity with respect to N_0 and total number of iterations.

References

1. Raynal, P., Lü, X., Englert, B.G.: Mutually unbiased bases in six dimensions: The four most distant bases. Phys. Rev. A **83**(6) (2011) 062303
2. Yan, F., Yang, M., Cao, Z.L.: Optimal reconstruction of the states in qutrit systems. Phys. Rev. A **82**(4) (2010) 044102
3. Adamson, R., Steinberg, A.M.: Improving quantum state estimation with mutually unbiased bases. Phys. Rev. Lett. **105**(3) (2010) 030406
4. Huszár, F., Houlsby, N.M.: Adaptive Bayesian quantum tomography. Phys. Rev. A **85**(5) (2012) 052120
5. Pereira, L., Zambrano, L., Cortés-Vega, J., Niklitschek, S., Delgado, A.: Adaptive quantum tomography in high dimensions. Phys. Rev. A **98**(1) (2018) 012339
6. Mahler, D.H., Rozema, L.A., Darabi, A., Ferrie, C., Blume-Kohout, R., Steinberg, A.: Adaptive quantum state tomography improves accuracy quadratically. Phys. Rev. Lett. **111**(18) (2013) 183601
7. Hou, Z., Zhu, H., Xiang, G.Y., Li, C.F., Guo, G.C.: Achieving quantum precision limit in adaptive qubit state tomography. npj Quantum Inf. **2**(1) (2016) 1–5
8. Blume-Kohout, Robin: Optimal, reliable estimation of quantum states. New J. Phys **12**(4) (2010) 043034
9. Lambert, B.: A Student's Guide to Bayesian Statistics. SAGE, London (2018)
10. Ferrie, C.: High posterior density ellipsoids of quantum states. New J. Phys **16**(2) (2014) 023006
11. Ferrie, C.: Self-guided quantum tomography. Phys. Rev. Lett. **113**(19) (2014) 190404
12. Granade, C., Ferrie, C., Flammia, S.T.: Practical adaptive quantum tomography. New J. Phys **19**(11) (2017) 113017
13. Abraham, H., et al.: Qiskit: An Open-source Framework for Quantum Computing (2019)
14. Munro, W., James, D., White, A., Kwiat, P.: Measurement of qubits. Phys. Rev. A **64**(030302) (2001)
15. Doucet, A., Godsill, S., Andrieu, C.: On sequential Monte Carlo sampling methods for Bayesian filtering. Volume 10. (2000)
16. Liu, J., West, M.: Combined parameter and state estimation in simulation-based filtering. Springer (2001)
17. Granadeand, C., Combes, J., Cory, D.: Practical Bayesian tomography. New J. Phys **18**(3) (2016) 033024
18. Gonçalves, D.S., Gomes-Ruggiero, M.A., Lavor, C., Jiménez, O.F., Ribeiro, P.S.: Local solutions of maximum likelihood estimation in quantum state tomography. Quantum Info. Comput. **12**(9-10) (2012) 775–790
19. Bickel, P., Li, B., Bengtsson, T., et al.: Sharp failure rates for the bootstrap particle filter in high dimensions. (2008)
20. Granade, C., Ferrie, C., Hincks, I., Casagrande, S., Alexander, T., Gross, J., Kononenko, M., Sanders, Y.: Qinfer: Statistical inference software for quantum applications. Quantum **1** (2017) 5
21. Johansson, J.R., Nation, P.D., Nori, F.: QuTiP: An open-source Python framework for the dynamics of open quantum systems. Comput. Phys. Commun **183**(8) (2012) 1760–1772
22. Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information. 10th anniversary edn. Cambridge University Press, New York (2010)
23. Granade, C.E., Ferrie, C., Wiebe, N., Cory, D.G.: Robust online Hamiltonian learning. New J. Phys **14**(10) (2012) 103013
24. IBM Quantum. https: //quantum-computing.ibm.com/, 2021.

25. Struchalin, G.I., Pogorelov, I.A., Straupe, S.S., Kravtsov, K.S., Radchenko, I.V., Kulik, S.P.: Experimental adaptive quantum tomography of two-qubit states. Phys. Rev. A **93**(1) (2016) 012103
26. Granade, C., Ferrie, C., Flammia, S.: Practical adaptive quantum tomography: Supplementary material (Apr 2016)
27. D.Gill, R., Massar, S.: State estimation for large ensembles. Phys. Rev. A **61** (2000) 042312
28. Zhu, H.: Quantum State Estimation and Symmetric Informationally Complete POMs. PhD Thesis, National University of Singapore (2012)
29. Rambach, M., Qaryan, M., Kewming, M., Ferrie, C., White, A.G., Romero, J.: Robust and Efficient High-Dimensional Quantum State Tomography. Phys. Rev. Lett. **126**(10) (2021) 100402
30. Chapman, R.J., Ferrie, C., Peruzzo, A.: Experimental demonstration of self-guided quantum tomography. Phys. Rev. Lett. **117**(4) (2016) 040402
31. Zulehner, A., Paler, A., Wille, R.: An efficient methodology for mapping quantum circuits to the IBM QX architectures. IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst **38**(7) (2018) 1226–1236