

Efficient Quantum State Tracking in Noisy Environments

Markus Rambach,^{1,*} Akram Youssry,^{2,3} Marco Tomamichel,⁴ and Jacqueline Romero¹

¹*Australian Research Council Centre of Excellence for Engineered
Quantum Systems & School of Mathematics and Physics,
University of Queensland, QLD 4072, Australia.*

²*University of Technology Sydney, Centre for Quantum Software and Information, Ultimo NSW 2007, Australia.*

³*Quantum Photonics Laboratory and Centre for Quantum Computation and Communication Technology,
RMIT University, Melbourne, VIC 3000, Australia*

⁴*Department of Electrical and Computer Engineering & Centre for Quantum Technologies,
National University of Singapore, Singapore 119077, Singapore.*

Quantum state tomography, which aims to find the best description of a quantum state—the density matrix, is an essential building block in quantum computation and communication. Standard techniques for state tomography are incapable of tracking changing states and often perform poorly in the presence of environmental noise. Although there are different approaches to solve these problems theoretically, experimental demonstrations have so far been sparse. Our approach, matrix-exponentiated gradient tomography, is an online tomography method that allows for state tracking, updates the estimated density matrix dynamically from the very first measurements, is computationally efficient, and converges to a good estimate quickly even with noisy data. The algorithm is controlled via a single parameter, its learning rate, which determines the performance and can be tailored in simulations to the individual experiment. We present an experimental implementation of matrix-exponentiated gradient tomography on a qutrit system encoded in the transverse spatial mode of photons. We investigate the performance of our method on stationary and evolving states, as well as significant environmental noise, and find fidelities of around 95% in all cases.

INTRODUCTION

Characterising quantum systems becomes increasingly important as quantum technologies begin to scale up. Experiments often require verification of the prepared quantum state and detection of errors as the state evolves, e.g., through deliberate evolution or environmental perturbations. Thus, quantum state tomography—finding the density matrix that best describes a quantum system—is a task central to quantum information processing. However, quantum state tomography (QST) is notoriously resource-intensive [1]. Reconstructing the d -by- d density matrix of a d -dimensional quantum system up to a fixed precision requires $O(d^2)$ parameters that can only be obtained after measuring at least $O(d^3)$ copies of the quantum state [2]. Most common algorithms for QST (e.g. maximum likelihood estimation [3] and least squares regression [4]) require a tomographically complete set of measurements and are done in post-processing, thus precluding real-time control of experiments based on the QST results.

Current experimental realisations of quantum devices often suffer from perturbations that change over time which might happen faster than a complete set of measurements can be performed. It might also be of interest to the experimentalist to observe the evolution of states in an experiment. Both these scenarios require online QST—learning algorithms that continuously update the estimate of the state description [5], and thus enable real-time control and diagnosis of errors. Recent proposals for this include Bayesian approaches [6, 7], adaptive measurements [8–11], and machine learning techniques [11–15]. All these techniques are also closely related to continuous learning [16, 17]—weak measurements over time to characterise a systems evolution, motivated by feedback control to correct errors—and Hamiltonian identification/learning [18, 19]—algorithms to determine Hamiltonian parameters governing the dynamics or unknown structures in the system, motivated by distinguishing and quantification of errors.

Moreover, arbitrary noise from the environment or from imperfect measurements can easily surpass the signal strength in experiments. This degrades our ability to accurately estimate the quantum state,

* m.rambach@uq.edu.au

hence the robustness of techniques is crucial. Usual efforts to circumvent this issue experimentally try to improve the signal-to-noise ratio (SNR), which is not always feasible or indeed possible. Some theoretical work exist, e.g., Ref. [20, 21] where the noise is formalised into the measurement, or Ref. [22] where an algorithm robust to noisy data is designed. However, none of these consider dynamical states and we are not aware of any experimental demonstrations.

Here, we focus on quantum state tomography using the matrix-exponentiated gradient method (MEG) [13, 14]. The reduced computational complexity per iteration of MEG, $O(d^3)$ in contrast to maximum likelihood and least squares regression which are both $O(d^4)$, makes it a good candidate for online quantum state tomography. The MEG algorithm is also efficient—it does not require a full set of measurements for each state estimate—and robust—simulations show convergence even with very noisy measurements. Additionally, the MEG update rule ensures that the estimated state is physical (i.e., a positive-semi definite density matrix), which is generally not guaranteed in most tomography methods unless the estimate is projected back into the physical space, causing a bias towards low-rank states in the estimator [23].

In this work, we experimentally demonstrate real-time, online quantum state tomography based on MEG using a high-dimensional quantum system encoded in the transverse spatial mode of single photons. We achieve fidelities of up to 95% even in the presence of significant noise due to statistical fluctuations (e.g. very low count rates) and environmental effects (e.g. an ambient light source). To our knowledge, this is the first experimental $O(d^3)$ online tomography implementation to date. A similar $O(d^3)$ online algorithm has been proposed recently [15], however, an experiment is still lacking.

RESULTS

Overview

Matrix-exponentiated gradient tomography (MEG) [13, 14] is an online algorithm that can estimate and track quantum states, adapted from machine learning techniques. The idea is to construct an iterative procedure, in which an estimate $\hat{\rho}_t$ of the (mixed or pure) quantum state at the iteration t is updated to a more accurate estimate $\hat{\rho}_{t+1}$, given a new measurement performed on the unknown quantum state. The matrix-exponentiated gradient update rule [14] for the estimate $\hat{\rho}_{t+1}$ is given by

$$\hat{\rho}_{t+1} = \frac{\exp(\log(\hat{\rho}_t) - \eta_t \nabla L_t)}{\text{tr} \exp(\log(\hat{\rho}_t) - \eta_t \nabla L_t)}, \quad (1)$$

where η_t is the learning rate, and L_t is the loss function,

$$L_t = (\text{tr}(\hat{\rho}_t X_t) - y_t)^2, \quad (2)$$

and its gradient is given by

$$\nabla L_t = 2(\text{tr}(\hat{\rho}_t X_t) - y_t) X_t. \quad (3)$$

The pair (X_t, y_t) denotes the measurement record where X_t is the measurement operator and y_t is the experimentally obtained average value of measuring the state with operator X_t . The measurements have to be informationally-complete as a requirement for performing any quantum tomography procedure. Therefore, at each iteration, a measurement operator X_t is chosen at random from a complete set of bases, and together with the measurement result y_t , the estimate of the state $\hat{\rho}_t$ is updated to $\hat{\rho}_{t+1}$. The learning rate η_t determines the weight given to the new obtained information at each iteration. For tomography on stationary states, η_t can be chosen to decrease continuously with the number of iterations while for online state tracking it needs to be constant or adaptive, depending on former measurement outcomes.

Compared to the originally proposed MEG in Ref. [14], we made a few changes and extensions to reflect the conditions in our experiment (more details in the METHODS section). First, we extended MEG to work with general higher-dimensional systems $d > 2$, beyond qubit systems and local Pauli measurements, by measuring mutually unbiased bases (MUBs) [24]. However, not all MUBs are known

for arbitrary dimensions so we introduced a second scheme which uses generalised Pauli operator measurements [25]. These operators form an orthonormal informationally-complete basis set and can be used to expand MEG to arbitrary high dimensions. Second, we modified the update rule in order to make better use of the experimental data. MEG originally drew on the probability of one measured state per iteration to update its estimate. Experimentally, we measure photon counts and need the counts of all d states of a given basis to calculate probabilities. This means that in each iteration we can calculate d probabilities which are now all employed to update the estimate. Third, we also adapted MEG to account for the fact that the prepared states in our experiment are very close to pure states—true for systems that have very high control over their preparation. In each iteration t , we find the eigenvalues and eigenvectors of the estimated density matrix $\hat{\rho}_t$ and pick the eigenvector with the largest eigenvalue as our estimated state $|\phi_t\rangle$ (density matrix $\hat{\rho}_\phi$) for benchmarking. We then calculate the infidelity $1 - f(\hat{\rho}_\phi, \hat{\Omega}_\psi) = 1 - \left(\text{tr} \sqrt{\sqrt{\hat{\rho}_\phi} \hat{\Omega}_\psi \sqrt{\hat{\rho}_\phi}} \right)^2$, where $f(\hat{\rho}_\phi, \hat{\Omega}_\psi)$ is the fidelity and $\hat{\Omega}_\psi$ is the density matrix of the theoretical prepared state $|\psi_t\rangle$. Our results show convergence to the unknown state within experimental limitations, most prominently mode-dependent losses, in all investigated situations and even for excessive experimental noise.

State Estimation and Online Tracking

We investigated the qutrit ($d = 3$) state estimation and online tracking performance of MEG using mutually unbiased bases (MUBs) and generalised Pauli operator measurements, which makes the algorithm applicable to higher dimensions. We also investigated MEG performance in the presence of small and large statistical and environmental noise. In all cases we continuously evolve the prepared state in time following

$$|\psi_t\rangle = \exp(-i\sigma\omega t) |\psi_0\rangle, \quad (4)$$

with σ a Hermitian matrix and $\omega = \frac{1.3}{t_{\text{tot}}}$ the rate of change dependent on the total amount of iterations t_{tot} . The rate of change was chosen to allow the state to evolve to a minimum fidelity compared to the initial state and then back to its initial state. We studied three possible cases of σ : $0_{3,3}$ which is the 3-dimensional zero matrix, σ_z which is the generalised 3-dimensional *Pauli Z* matrix given by

$$\sigma_z = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -2 \end{pmatrix}, \quad (5)$$

and σ_r which is a general random Hermitian matrix. The zero matrix results in the identity evolution leaving the initial state stationary, while the other two matrices simulate a rotation in the Hilbert space of a qutrit. To ensure a fair comparison, the learning rate—a crucial parameter for performance and convergence of the algorithm—is kept constant at $\eta = 5$, which shows good performance in all cases albeit not optimal for $\sigma = 0_{3,3}$. The learning rate can be optimised individually, dependent on investigated scenario and physical system.

An initial baseline measurement using MUBs and high signal rates ($N \approx 10^6$ photons per measurement) is shown in Fig. 1(a). The high count rate ensures that statistical Poissonian distributed counting noise ($\Delta N = \sqrt{N} \approx 10^3$), background ($N_{\text{back}} \approx 50$) and dark counts ($N_{\text{dark}} \approx 100$), are negligible and do not deteriorate the performance. For this case, the algorithm finds density matrices with a median purity above 99.6% in all investigated scenarios, as to be expected for our system where we have a high degree of control over the state preparation. Since the fidelities in the experiments are very high, we plot the infidelity, $1 - f(\hat{\rho}_\phi, \hat{\Omega}_\psi)$, for clarity. The solid lines in Fig. 1(a) are the median infidelities of 50 randomly chosen states of qutrits (according to Haar measure), with the shaded regions bounded by the upper and lower quartile (50 ± 25)%.

In theory, MEG can estimate any quantum state with arbitrary precision and accuracy given an infinite amount of iterations. However, experimental imperfections such as misalignment and mode-dependent loss significantly limit the minimum infidelities achievable by QST algorithms (not just




Figure 1. MEG on qutrits for (a) high signal rates ($N \approx 10^6$ photons per measurement) using MUBs, (b) low signal rates ($N \approx 10^2$ photons per measurement) using MUBs, and (c) low signal rates using generalised Pauli operators. In each case, three scenarios according to different time evolutions of the states (Eq. 4) are investigated: stationary (blue), σ_z (red), and a random unitary σ_r (black). Solid lines are median performances of the algorithm out of 50 randomly chosen states (according to Haar measure). The shaded regions are the interquartile ranges.

MEG). The infidelities for MEG listed in Tab. 1 are similar to those obtained using other quantum state tomography algorithms applied to our system, e.g. root-approach and maximum-likelihood estimators [26] where we set the rank to 1.

We next show the robustness of MEG to noise using MUBs, reducing the count rate per iteration to $N \approx 10^2$. This means that the values for Poissonian, dark, and background noise become comparable and play a significant role in state estimation and tracking. Robustness to noise in the limit of a small number of repeated preparations of the unknown state—here the number of photons—is especially interesting for systems like ions and superconducting qubits, as the preparation is considerably more time consuming compared to our system. Again, we benchmark the performance of MEG by calculating the infidelity in each iteration, see Fig. 1(b). State estimation and online tracking are in excellent agreement with an overall mean infidelity $\sim 5.5\%$, slightly higher than in the high count rate case, but well within the uncertainty bars. This already shows the robustness of the modified algorithm to noise. Infidelities below 10% are achieved within around 15 iterations, four times more than in the baseline measurement with high count rates (see Tab. 1 for details).

We furthermore studied the performance of MEG under high noise, but using generalised Pauli operator measurements rather than MUBs, shown in Fig. 1(c). As generalised Pauli operators can be mathematically defined using them instead of MUBs is a powerful tool to describe states far beyond the qubit. We observe that state estimation and online tracking are again in excellent agreement and the overall mean infidelity $\sim 5.1\%$ is slightly better than in the MUB case, but overlapping within uncertainty bars. This demonstrates that generalised Pauli measurements is a viable path to extend MEG towards arbitrary high dimensions in realistic experimental settings. The initial convergence towards the unknown state is slower with around 25 iterations necessary to reach an infidelity of 10%. This is to be expected as MUBs span the qutrit state space more efficiently compared to the Pauli operator measurements. Tab. 1 shows details of the performance.

Since a changing state leads to changing probabilities of finding a system in a particular state, we also show that we are able to follow these probabilities. Fig. 2 illustrates one example of these probabilities $p_i = |\langle i | \psi_t \rangle|^2$ for the logical basis states $|i\rangle, i = 0, 1, 2$, out of the 50 randomly chosen qutrits analysed in Fig. 1. The theoretically expected probabilities in the experiment (dashed lines) are followed nicely by online state tracking of the deterministic evolution in the experiment (solid lines). This confirms that the low infidelities directly correspond to the capability to predict measurement outcomes on the states well. We also highlight that MEG quickly converges to a useful estimate (after 3-4 iterations), even though the initial guess is a completely mixed qutrit ($p_i = 1/3 \forall i$) to avoid bias in the results. The small offset between theory and experiment could be further decreased by using




Figure 2. Tracking of a quantum state under randomly chosen Hermitian evolution σ_r . MEG can track the theoretical (dashed) deterministic evolution of a state experimentally (solid). The three probability components p_i for the logical basis states $i = 0, 1, 2$ (blue, red, black) follow the theory nicely. Starting at an initial guess of a completely mixed state, useful predictions for p_i are obtained very effectively after 3-4 iterations.

an adaptive learning rate.

Noisy MEG

We tested the efficacy of MEG under significant environmental noise experimentally—by adding a light bulb close to the detector while keeping the signal rate at $N \approx 10^2$ photons per measurement—and in simulations—adding Poissonian distributed noise to the raw data. The strength of the noise in the experiment is controlled via the distance and angle of the light bulb to the single photon detector. As a proof of principle demonstration, we only looked at stationary states using MUBs in our experiment, yet the simulations show similar behaviour for all estimation and online tracking scenarios investigated in the previous section. The median infidelities of 50 randomly chosen states for different methods (solid—experiment, dashed—simulation) and noise strengths (colours) are shown in Fig. 3. Shaded uncertainty ranges similar to Fig. 1 are omitted for clarity. As expected, the achievable infidelity and number of iterations to get there both become higher as the noise increases, however, MEG is showing

Table 1. MEG performance indicators. Top: iterations needed to reach 10% infidelity. Bottom: mean infidelity. All values achieved with a constant learning rate $\eta = 5$. Uncertainties are the boundaries of the interquartile ranges.

Signal rate (Hz)	MUB			Pauli operators		
	$\sigma = 0_{3,3}$	$\sigma = \sigma_z$	$\sigma = \sigma_r$	$\sigma = 0_{3,3}$	$\sigma = \sigma_z$	$\sigma = \sigma_r$
Iterations for Infidelity < 10%						
10^2	14 $^{+17}_{-10}$	13 $^{+26}_{-6}$	16 $^{+26}_{-12}$	30 $^{+32}_{-21}$	23 $^{+68}_{-14}$	24 $^{+38}_{-17}$
10^6	4 $^{+2}_{-1}$	4 $^{+3}_{-1}$	4 $^{+4}_{-1}$			
Mean Infidelity						
10^2	5.3 $^{+3.2}_{-2.2}$	5.4 $^{+3.4}_{-2.4}$	5.6 $^{+3.4}_{-2.5}$	4.9 $^{+2.8}_{-2.0}$	5.4 $^{+3.5}_{-2.4}$	5.0 $^{+3.3}_{-2.1}$
10^6	3.4 $^{+1.8}_{-1.2}$	5.1 $^{+2.2}_{-1.9}$	4.7 $^{+2.3}_{-1.9}$			




Figure 3. MEG performance under excessive noise in experiment (exp, solid) and simulation (sim, dashed) with low signal rates ($\sim 10^2$ photons per measurement), using MUBs. Median performance out of 50 randomly chosen states $|\psi_0\rangle$ (experiment) and 20 runs each (total of $50 \times 20 = 1000$ traces) of added Poissonian distributed noise (simulation). Uncertainty ranges omitted for clarity.

exceptional robustness.

We observe that up to $N_{back} \approx 2$ kHz of experimental and simulated noise (not shown explicitly in the figure), the infidelities are comparable to the *no noise* case within uncertainty bars. This corresponds to a signal-to-noise ratio $\text{SNR} \approx 0.05$, so only one in 20 photons carry the actual information on the unknown quantum state prepared in the experiment. It is noteworthy that even at 0 kHz of artificially introduced noise, the naturally occurring counting noise together with background and dark count rates are almost double the signal count rate, i.e., ~ 170 compared to ~ 100 photons per measurement. Experimental and simulated data for $N_{back} = [1, 2.5]$ kHz show good overlap within uncertainty bars, indicating the validity of the noise model in simulation. Therefore we expect to achieve infidelities of around 10% up to $\text{SNR} \approx 0.01$ (orange dashed line). The seemingly smoother performance of the simulations stems from 20 runs of randomly added shot-to-shot noise for each of the 50 investigated states. In other words, the results from simulations reflect the median infidelity of $50 \times 20 = 1000$ traces, compared to 50 in the experiment.

DISCUSSION

We have presented an efficient tomography algorithm to track the continuous evolution of quantum states despite significant environmental noise. Our approach, matrix-exponentiated gradient tomography (MEG), uses machine learning techniques to iteratively update an estimate of a quantum state in real time. We have extended the original online tomography proposal [14] to arbitrary high-dimensional systems (qudits) and experimentally demonstrated the algorithm with a qutrit encoded in the shape of photons. We have shown that MEG is highly resilient to noise: we are able to estimate the quantum state with high fidelities, and accurately predict probabilities of measurement outcomes for different types of evolution and signal strengths. Our technique will enable applications that benefit from efficient online tomography, e.g., quantum process tomography [27, 28] and quantum error mitigation [29]/correction [30].

Our results can be extended to process tomography through the Choi-Jamiołkowski isomorphism [31, 32], which allows us to represent a channel using a quantum state that lives in a d^2 -

dimensional Hilbert space. In this case, the process estimation becomes state estimation, and MEG could be utilised for the procedure. This provides a more efficient solution compared to standard process tomography which requires computational complexity $O(d^8)$ by using a tomographically complete set of states in preparation and measurement, even $O(d^{10})$ for an over-complete set [33]. The bottleneck of the standard procedure comes from the matrix-vector multiplications required for the estimation procedure. On the other hand, for the MEG update rule, the bottleneck operation is the matrix exponential which is computationally more efficient (see [14] and references within).

METHODS

Experimental Setup and Measurement

The quantum states are encoded in the transverse spatial mode—the shape—of single photons. We describe these modes in the Laguerre-Gaussian basis $\{|l_i, p_i\rangle\}$ [34], where each randomly chosen state is given by the superposition $|\psi\rangle = \sum_i c_i |l_i, p_i\rangle$ (with $\sum_i |c_i|^2 = 1$). For the investigated qutrits in the experiment we use the three logical basis states $\{|0\rangle \equiv |-1, 0\rangle, |1\rangle \equiv |0, 2\rangle, |2\rangle \equiv |1, 0\rangle\}$, keeping the Gouy phase of the same order and therefore preventing the shape from rotating as the photons propagate [35]. The experiment is conducted with highly attenuated CW laser light (Thorlabs, mean photon number $|\alpha|^2 = 0.01$) together with two spatial light modulators (SLM, Meadowlark Optics) and a single photon detector (Perkin-Elmer, ~ 100 Hz dark count rate). The first SLM is used to prepare the *unknown* quantum state $|\psi_t\rangle$ via a hologram that changes phase and amplitude of the initial state $|l = 0, p = 0\rangle$ from the laser source. A second hologram displayed on the other SLM is then measuring the state $|\psi_m\rangle$ from one randomly chosen basis set (see *MEG Estimation Extensions* section for details). Finally, a single-mode fibre and the detector are acting as a local filter with the number of detected photons proportional to the overlap $|\langle\psi_m|\psi_t\rangle|^2$ between the prepared and measured mode. The schematic, details, and characterisation of the simple but effective prepare-and-measure experiment are described extensively in Ref. [36] and its supplemental material.

MEG Estimation Extensions

Measurement Schemes

We use two measurement schemes to perform our state estimation and online tracking. The first scheme is utilising Mutually Unbiased Bases (MUBs) [24], measurements that are convenient to perform on the photonic system in our experiments. For a d -dimensional quantum system, we have $d + 1$ basis sets, each set consists of d states. However, complete sets of MUBs are not known for every dimension and so, we also introduced a second scheme using generalised Pauli operators [25]. These operators are given by the set $\{u_{jk}, v_{jk}, w_l\}$

$$u_{jk} = |j\rangle\langle k| + |k\rangle\langle j|, \quad 0 \leq j \leq k \leq d \quad (6a)$$

$$v_{jk} = -i|j\rangle\langle k| + i|k\rangle\langle j|, \quad 0 \leq j \leq k \leq d \quad (6b)$$

$$w_l = \sqrt{\frac{2}{l+1}} \left(-l|l\rangle\langle l| + \sum_{j=0}^l |j\rangle\langle j| \right). \quad 0 \leq l \leq d \quad (6c)$$

The operators reduce to the Pauli and Gell-Mann matrices for $d = 2$ and $d = 3$, respectively. The total number of operators is $d^2 - 1$, and for each operator, there are d eigenstates that could be measured. The generalised Pauli operators form an orthonormal basis set, and thus are informationally-complete.

Modified update rule

The learning rate is usually chosen to be adaptive to guarantee convergence in the case of noisy measurements as shown in [14]. However, upon exploring the effect of choosing different learning

rates, we found that a constant learning rate is sufficient to obtain an estimate with the maximum possible accuracy obtainable given our experimental capabilities. Thus, we fix the learning rate to be $\eta_t = \eta = 5$.

The measurement schemes we use provide us with more information at each iteration. In the case of MUB measurements, we randomly select a basis at each iteration. In order to calculate the probabilities, we need to projectively measure all states of that basis for normalisation of the measurement counts. The same holds when we use generalised Pauli measurements, we still need to measure all the eigenstates of a given Pauli operator. Therefore, we modify the loss function to include multiple measurement results obtained at a given iteration as follows.

$$L_t = \sum_{i=1}^d \left(\text{tr}(\hat{\rho}_t X_t^{(i)}) - y_t^{(i)} \right)^2, \quad (7)$$

with gradient

$$\nabla L_t = 2 \sum_{i=1}^d \left(\text{tr}(\hat{\rho}_t X_t^{(i)}) - y_t^{(i)} \right) X_t^{(i)}, \quad (8)$$

where $X_t^{(i)} = |\psi_i\rangle\langle\psi_i|$ is the projector of the i^{th} state of the MUB (or Pauli operator) basis that was randomly chosen at iteration t out of $d + 1$ (or $d^2 - 1$) possible bases, and $y_t^{(i)}$ is the measured probability.

Since we are restricting the experiments to pure states in our investigations, we can increase the estimation accuracy using this prior information. The method we use is to project the state into the subspace of its largest eigenvalue, or equivalently, finding the pure state closest to our estimate in fidelity. We illustrate this method with an example: we can write the estimated density matrix $\hat{\rho}_t$ of a pure quantum state $|\phi_t\rangle$ in the form

$$\hat{\rho}_t = \lambda |\phi_t\rangle\langle\phi_t| + (1 - \lambda)\hat{\rho}_{\text{noise}}, \quad (9)$$

where the parameter $\lambda \in [0, 1]$ indicates the noise strength, and $\hat{\rho}_{\text{noise}} = I_d/d$ is the completely mixed state representing the noise in the estimate, with I_d the d dimensional identity matrix. The special cases $\lambda = 0$ and $\lambda = 1$ correspond to the noiseless and completely noisy states respectively. Now, let $|\gamma_t\rangle$ be an eigenstate of $\hat{\rho}_t$ with eigenvalue γ_t , i.e., $\hat{\rho}_t |\gamma_t\rangle = \gamma_t |\gamma_t\rangle$. Then, we can write

$$\gamma_t = \langle\gamma_t|\hat{\rho}_t|\gamma_t\rangle \quad (10a)$$

$$= \lambda \langle\gamma_t|\phi_t\rangle \langle\phi_t|\gamma_t\rangle + (1 - \lambda) \frac{\langle\gamma_t|I_d|\gamma_t\rangle}{d} \quad (10b)$$

$$= \lambda |\langle\gamma_t|\phi_t\rangle|^2 + \frac{1 - \lambda}{d}. \quad (10c)$$

From there, we can see that the maximum eigenvalue is

$$\gamma_{t,\text{max}} = \lambda + \frac{1 - \lambda}{d}, \quad (11)$$

which is obtainable when there is maximum overlap between γ_t and ϕ_t , in the best case $|\langle\gamma_t|\phi_t\rangle|^2 = 1$, i.e., $|\gamma_t\rangle = |\phi_t\rangle$. This shows that the true pure state $|\phi_t\rangle$ is an eigenvector of the density matrix $\hat{\rho}_t$ corresponding to the maximum eigenvalue, in other words, closest to our estimator in fidelity. So, after each iteration we calculate the eigenvalues of the state estimate and select the corresponding eigenvector to be the estimate of the pure state $|\phi_t\rangle$.

DATA AVAILABILITY

All the generated datasets used in this study are publicly available on Figshare at https://figshare.com/projects/Efficient_Tracking_of_Noisy_Quantum_States/130433.

CODE AVAILABILITY

The source code for the proposed theoretical methods is publicly available on github at https://github.com/akramyousry/MEG_state_tracking.

ACKNOWLEDGEMENTS

This research was supported by the Australian Research Council Centre of Excellence for Engineered Quantum Systems (EQUS, CE170100009) and Discovery Project (DP200102273). AY is supported by an Australian Government Research Training Program Scholarship, and the Australian Research Council under the Centre of Excellence Scheme No. CE170100012. MT is supported by the National Research Foundation, Prime Minister's Office, Singapore and the Ministry of Education, Singapore under the Research Centres of Excellence programme, and by NUS startup grants (R-263-000-E32-133 and R-263-000-E32-731). JR is supported by a Westpac Bicentennial Foundation Research Fellowship.

AUTHOR CONTRIBUTIONS

MR performed the optical experiments and acquired and evaluated the data with inputs from JR. MR and AY performed the numerical experiments. AY and MT designed and AY implemented the algorithm based on feedback from MR and JR. MT and JR supervised the project. MR led the writing of the article with contributions from all authors.

COMPETING INTERESTS

The authors declare no competing financial or non-financial interests.

REFERENCES

- [1] M. Paris and J. Řeháček, eds., *Quantum State Estimation*, Lecture Notes in Physics, Vol. 649 (Springer Berlin, Heidelberg, 2004).
- [2] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, Sample-optimal tomography of quantum states, *IEEE Transactions on Information Theory* **63**, 5628 (2017).
- [3] G. M. D'Ariano, M. G. Paris, and M. F. Sacchi, Quantum Tomography, in *Advances in Imaging and Electron Physics*, Vol. 128, edited by P. W. Hawkes (Elsevier, 2003) pp. 205–308.
- [4] T. Opatrný, D.-G. Welsch, and W. Vogel, Least-squares inversion for density-matrix reconstruction, *Phys. Rev. A* **56**, 1788 (1997).
- [5] S. Aaronson, X. Chen, E. Hazan, S. Kale, and A. Nayak, Online learning of quantum states, *J. Stat. Mech. Theory Exp.* **2019**, 124019 (2019).
- [6] C. Granade, J. Combes, and D. G. Cory, Practical Bayesian tomography, *New J. Phys.* **18**, 033024 (2016).
- [7] C.-M. Lin, Y.-M. Hsu, and Y.-H. Li, *An Online Algorithm for Maximum-Likelihood Quantum State Tomography* (2020), [arXiv:2012.15498](https://arxiv.org/abs/2012.15498).
- [8] B. Qi, Z. Hou, Y. Wang, D. Dong, H.-S. Zhong, L. Li, G.-Y. Xiang, H. M. Wiseman, C.-F. Li, and G.-C. Guo, Adaptive quantum state tomography via linear regression estimation: Theory and two-qubit experiment, *npj Quantum Inf.* **3**, 19 (2017).
- [9] Y. Chen and X. Wang, *More Practical and Adaptive Algorithms for Online Quantum State Learning* (2020), [arXiv:2006.01013](https://arxiv.org/abs/2006.01013).
- [10] S. Nohara, R. Okamoto, A. Fujiwara, and S. Takeuchi, Adaptive quantum state estimation for dynamic quantum states, *Phys. Rev. A* **102**, 030401 (2020).
- [11] Y. Quek, S. Fort, and H. K. Ng, Adaptive quantum state tomography with neural networks, *npj Quantum Inf.* **7**, 105 (2021).
- [12] C. Ferrie, Self-Guided Quantum Tomography, *Phys. Rev. Lett.* **113**, 190404 (2014).
- [13] Y.-H. Li and V. Cevher, Convergence of the Exponentiated Gradient Method with Armijo Line Search, *J. Optim. Theory Appl.* **181**, 588 (2019).

- [14] A. Youssry, C. Ferrie, and M. Tomamichel, Efficient online quantum state estimation using a matrix-exponentiated gradient method, *New J. Phys.* **21**, 033006 (2019).
- [15] K. Zhang, S. Cong, K. Li, and T. Wang, An online optimization algorithm for the real-time quantum state tomography, *Quantum Inf. Process.* **19**, 361 (2020).
- [16] A. Silberfarb, P. S. Jessen, and I. H. Deutsch, Quantum State Reconstruction via Continuous Measurement, *Phys. Rev. Lett.* **95**, 030402 (2005).
- [17] A. Shabani, J. Roden, and K. B. Whaley, Continuous Measurement of a Non-Markovian Open Quantum System, *Phys. Rev. Lett.* **112**, 113601 (2014).
- [18] J. H. Cole, S. G. Schirmer, A. D. Greentree, C. J. Wellard, D. K. L. Oi, and L. C. L. Hollenberg, Identifying an experimental two-state Hamiltonian to arbitrary accuracy, *Phys. Rev. A* **71**, 062312 (2005).
- [19] A. Anshu, S. Arunachalam, T. Kuwahara, and M. Soleimanifar, Sample-efficient learning of interacting quantum systems, *Nat. Phys.* **17**, 931 (2021).
- [20] Y. I. Bogdanov, B. I. Bantysh, N. A. Bogdanova, A. B. Kvasnyy, and V. F. Lukichev, Quantum states tomography with noisy measurement channels, in *Int. Conf. Micro- Nano-Electronics 2016*, Vol. 10224, edited by V. F. Lukichev and K. V. Rudenko (2016) p. 102242O.
- [21] V. N. Ivanova-Rohling, N. Rohling, and G. Burkard, *Optimal Quantum State Tomography with Noisy Gates* (2022), [arXiv:2203.05677](https://arxiv.org/abs/2203.05677).
- [22] A. Farooq, U. Khalid, J. ur Rehman, and H. Shin, Robust Quantum State Tomography Method for Quantum Sensing, *Sensors* **22**, 2669 (2022).
- [23] R. Blume-Kohout, Optimal, reliable estimation of quantum states, *New J. Phys.* **12**, 043034 (2010).
- [24] T. Durt, B.-G. Englert, I. Bengtsson, and K. Życzkowski, On Mutually Unbiased Bases, *Int. J. Quantum Inf.* **08**, 535 (2010).
- [25] G. Kimura, The Bloch vector for N-level systems, *Phys. Lett. A* **314**, 339 (2003).
- [26] Y. I. Bogdanov, Unified statistical method for reconstructing quantum states by purification, *J. Exp. Theor. Phys.* **108**, 928 (2009).
- [27] I. L. Chuang and M. A. Nielsen, Prescription for experimental determination of the dynamics of a quantum black box, *J. Mod. Opt.* **44**, 2455 (1997).
- [28] J. L. O'Brien, G. J. Pryde, A. Gilchrist, D. F. V. James, N. K. Langford, T. C. Ralph, and A. G. White, Quantum Process Tomography of a Controlled-NOT Gate, *Phys. Rev. Lett.* **93**, 080502 (2004).
- [29] S. Endo, S. C. Benjamin, and Y. Li, Practical Quantum Error Mitigation for Near-Future Applications, *Phys. Rev. X* **8**, 031027 (2018).
- [30] D. A. Lidar and T. A. Brun, eds., *Quantum Error Correction* (Cambridge University Press, Cambridge, 2012).
- [31] A. Jamiołkowski, Linear transformations which preserve trace and positive semidefiniteness of operators, *Reports Math. Phys.* **3**, 275 (1972).
- [32] M.-D. Choi, Completely positive linear maps on complex matrices, *Linear Algebra Appl.* **10**, 285 (1975).
- [33] G. C. Knee, E. Bolduc, J. Leach, and E. M. Gauger, Quantum process tomography via completely positive and trace-preserving projection, *Phys. Rev. A* **98**, 062336 (2018).
- [34] L. Allen, S. M. Barnett, and M. J. Padgett, *Opt. Angular Momentum* (CRC Press, Boca Raton, 2016).
- [35] N. K. Langford, R. B. Dalton, M. D. Harvey, J. L. O'Brien, G. J. Pryde, A. Gilchrist, S. D. Bartlett, and A. G. White, Measuring Entangled Qutrits and Their Use for Quantum Bit Commitment, *Phys. Rev. Lett.* **93**, 053601 (2004).
- [36] M. Rambach, M. Qaryan, M. Kewming, C. Ferrie, A. G. White, and J. Romero, Robust and Efficient High-Dimensional Quantum State Tomography, *Phys. Rev. Lett.* **126**, 100402 (2021).