

PAPER • OPEN ACCESS

Fast state tomography with optimal error bounds

To cite this article: M Gu *et al* 2020 *J. Phys. A: Math. Theor.* **53** 204001

View the [article online](#) for updates and enhancements.



IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

Fast state tomography with optimal error bounds

M Guță^{1,4} , J Kahn², R Kueng³ and J A Tropp³

¹ School of Mathematical Sciences, University of Nottingham, United Kingdom

² Institut de Mathématiques de Toulouse, Toulouse, France

³ California Institute of Technology, Pasadena, United States of America

E-mail: Madalin.guta@nottingham.ac.uk

Received 31 December 2019, revised 15 March 2020

Accepted for publication 18 March 2020

Published 28 April 2020



CrossMark

Abstract

Projected least squares is an intuitive and numerically cheap technique for quantum state tomography: compute the least-squares estimator and project it onto the space of states. The main result of this paper equips this point estimator with rigorous, non-asymptotic convergence guarantees expressed in terms of the trace distance. The estimator's *sample complexity* is comparable to the strongest convergence guarantees available in the literature and—in the case of the uniform POVM—saturates fundamental lower bounds. Numerical simulations support these competitive features.

Keywords: quantum tomography, projected least squares estimator, low rank states, concentration bounds, optimal estimation

(Some figures may appear in colour only in the online journal)

1. Introduction

Quantum state tomography is the task of reconstructing a quantum state from experimental measurement data. Among various estimation methods, maximum-likelihood (ML) [1, 2] is a universal approach which produces point estimators with good *average error* properties and is extensively used in practice [3]. However, ML is computationally expensive and the practice of using ML in combination with bootstrap to produce *error bars* is not always theoretically grounded. Indeed such methods can be used to produce asymptotic confidence regions in scenarios where the theory of local asymptotic normality applies (e.g. for fully mixed states), but are inconsistent for rank-deficient states [4]. This has spurred the development of alternatives, such as Bayesian [5–7] and region [8–10] estimators. However, these methods have other drawbacks, such as comparatively high computational cost and weak (or implicit) convergence guarantees. In parallel, the estimation of low-rank states from a small number of observables

⁴Author to whom any correspondence should be addressed.



has been studied in compressed sensing [11–14], and in a series of papers by Koltchinskii and coauthors [15–18].

In this work we present an in depth analysis of an alternative method, the *projected least squares* (PLS) estimator, and show that it improves on the status quo in several significant directions. PLS is obtained by performing a standard linear inversion [19] followed by a projection onto the space of density matrices. It is numerically ‘cheap’ in the sense that its computational cost is dominated by forming the least-squares estimator. Indeed, numerical simulations show that PLS is much faster than prominent alternatives like maximum-likelihood estimation [3] and compressed sensing [11].

PLS is also *statistically efficient*. Our main theoretical results show that PLS achieves fundamental lower bounds [20] for tomography with *separate* measurements: to reconstruct an arbitrary d dimensional state ρ of rank- r with accuracy ϵ in trace distance it suffices to measure $r^2 d \epsilon^{-2} \log d$ independent samples with a two-design measurement, or $r^2 d \epsilon^{-2}$ samples with a covariant measurement. This *sampling rate* improves upon existing results [21] and is competitive with the most powerful techniques in the literature [13, 22]. The statistical performance of the PLS estimator has also been studied in [18] in the context of a different statistical model involving noisy observations of expectations of randomly chosen observables. In this context the authors show that up to logarithmic factors, the PLS estimator achieves general lower bounds for several loss functions.

Note that *correlated* measurements over several copies of the state can achieve even lower sampling rates [20, 23, 24], but are much more challenging to implement. Here we focus on the practically more relevant case of measuring each copy of ρ separately. The complementary analysis in [25] shows the PLS has a tractable *asymptotic* behavior which may provide the route to devising simple, tight, and theoretically justified error bars.

For a direct access to results and techniques, the paper is organised according to the onion peeling principle. We first introduce concepts and results and gradually reveal the technical details, with most proofs relegated to the appendices. In section 2 we introduce the PLS estimator. In section 3 we present the main theoretical results which establish concentration bounds for the PLS for several important classes of measurements: two-designs, Pauli bases, Pauli observables and covariant measurements. Section 4 provides the algorithmic details necessary for implementing the PLS estimator for the different measurement classes, including the explicit expression of the LS estimator and the projection onto the space of states. The detailed calculations and the proofs of the main theorems can be found in the appendices. Finally, in section 5 we present simulation results showing that the statistical performance of PLS is comparable to ML and compressed sensing for a range of states and measurements.

2. The projected least squares estimator

A measurement on a d -level system is described by $d \times d$ Hermitian matrices $M_1, \dots, M_m \in \mathbb{H}_d$ that are positive semidefinite and obey $\sum_{i=1}^m M_i = \mathbb{I}$. All measurements will be assumed to be *informationally complete*, so that for every pair of distinct states $\rho \neq \sigma$, there exists $i \in [m] = \{1, \dots, m\}$ such that $\text{tr}(M_i \rho) \neq \text{tr}(M_i \sigma)$. Measuring a quantum state ρ results in one of m outcomes, indexed by $i \in [m]$. The probability of observing outcome i depends on ρ and is described by Born’s rule:

$$[p]_i = \Pr [i|\rho] = \text{tr}(M_i \rho) \quad \text{for } i \in [m]. \quad (1)$$

These probabilities can be estimated by the *empirical frequencies*: measure n identical copies of the state separately, and set

Table 1. Summary of the estimation technique.

 Projected least squares (PLS) estimator $\hat{\rho}_n$

- Estimate probabilities by frequencies (2),
 - Compute the least squares estimator (3),
 - Project it onto the set of quantum states (4).
-

$$[f_n]_i = \frac{n_i}{n} \quad \text{for } i \in [m], \quad (2)$$

where n_i is the number of times outcome i was observed. Indeed by the law of large numbers, $[f_n]_i \rightarrow [p]_i$ as $n \rightarrow \infty$. The *least-squares (LS) estimator* is the solution to the least-squares problem that results from replacing the true probabilities in Born's rule by frequencies (2):

$$\hat{L}_n = \operatorname{argmin}_{X \in \mathbb{H}_d} \sum_{i=1}^m ([f_n]_i - \operatorname{tr}(M_i X))^2. \quad (3)$$

This optimization inverts the m linear equations specified by Born's rule, and the solution has the closed form expression (7) which will be discussed in more detail in section 4. In general, \hat{L}_n can have negative eigenvalues, so it may *fail* to be a quantum state. Several ways to overcome this drawback have been proposed [17, 21, 26–29].

In this paper we consider the *projected least squares (PLS) estimator*, which is obtained by *projecting* \hat{L}_n onto the convex set of all quantum states with respect to the Frobenius distance

$$\hat{\rho}_n = \operatorname{argmin}_{\sigma \text{ is a quantum state}} \|\hat{L}_n - \sigma\|_2. \quad (4)$$

The procedure for computing PLS is summarised in table 1.

The simplicity of PLS is a key advantage for both the analysis and the actual computation. The least-squares estimator (3) admits a closed-form expression, while the projection onto the set of quantum states (11) can be computed efficiently using a simple thresholding algorithm which adjusts the eigenvalues of the LS estimator while leaving the eigenvectors unchanged [27], as will be detailed in section 4. Hence, computing PLS requires considerably less storage and arithmetic than existing techniques that are based on more complicated optimization problems.

Below we discuss the statistical performance of PLS, and show that it exhibits a rank-dependent risk reduction with respect to least squares which matches fundamental lower bounds of [20].

3. Results

3.1. Error bounds and confidence regions for $\hat{\rho}_n$

We analyze several important and practically relevant measurement schemes: *structured POVMs* (e.g. symmetrically informationally complete (SIC)–POVMs, mutually unbiased bases (MUBs) and stabilizer states), *global Pauli observables*, *Pauli basis measurements*, and the *uniform/covariant POVM*. For details on these measurement classes, and the computation of the LS and PLS estimators we refer to section 4. Here we focus on explaining the novel results and the key techniques used in establishing them. The following theorem shows that the PLS estimator $\hat{\rho}_n$ converges to the true state ρ in trace distance $\|\cdot\|_1$ and

provides rank-dependent concentration bounds for each measurement class and arbitrary sample size n .

Theorem 1. (Error bound for $\hat{\rho}_n$) Let $\rho \in \mathbb{H}_d$ be a state and fix a number of samples $n \in \mathbb{N}$. Then, for each of the aforementioned measurements, the PLS estimator (table 1) obeys

$$\Pr [\|\hat{\rho}_n - \rho\|_1 \geq \epsilon] \leq de^{-\frac{n\epsilon^2}{43g(d)r^2}} \quad \text{for } \epsilon \in [0, 1],$$

where $r = \min \{\text{rank}(\rho), \text{rank}(\hat{\rho}_n)\}$ and $g(d)$ specifies dependence on the ambient dimension:
 $g(d) = 2d$ for structured POVMs; see equation (8) for \hat{L}_n ,
 $g(d) = d^2$ for Pauli observables; see equation (9) for \hat{L}_n ,
 $g(d) \simeq d^{1.6}$ for Pauli basis measurements; see equation (10) for \hat{L}_n .

The following immediate corollary endows $\hat{\rho}_n$ with rigorous non-asymptotic error-bars in trace distance.

Corollary 1. The trace-norm ball of size $\text{rank}(\hat{\rho}_n)\sqrt{43g(d)n^{-1} \log(d/\delta)}$ around $\hat{\rho}_n$ (intersected with state space) is a δ -confidence region for the true state ρ .

We emphasize the following aspects of this result:

(a) (Almost) optimal sampling rate: theorem 1 highlights that

$$n \geq 43g(d)\frac{\text{rank}(\rho)^2}{\epsilon^2} \log\left(\frac{d}{\delta}\right) \tag{5}$$

samples suffice to ensure $\|\hat{\rho}_n - \rho\|_1 \leq \epsilon$ with probability at least $1 - \delta$. For structured POVMs and Pauli observables, this *sampling rate* is comparable to the *best* theoretical bounds for alternative tomography algorithms [11, 30]. Moreover, fundamental lower bounds in [13, 20] indicate that this scaling is optimal up to a single $\log(d)$ -factor, so it cannot be improved substantially.

(b) *Implicit exploitation of (approximate) low rank*: the number of samples required to achieve a good estimator scales quadratically in the rank, rather than the ambient dimension d . This behavior extends to the case where ρ , or $\hat{\rho}_n$, is well-approximated by a rank- r matrix; see theorem 4 in appendix E. These results are comparable with guarantees for compressed sensing methods [13] that are specifically designed to exploit low-rank. For numerical confirmation we refer to figure 2, and other simulation results in section 5.

Proof sketch for theorem 1. The least-squares estimator \hat{L}_n can be viewed as a sum of n independent random matrices. To illustrate this, consider a single measurement of the structured POVM type. Then \hat{L}_1 defined in (8) is an instance of the random matrix $X = (d + 1)|v_k\rangle\langle v_k| - \mathbb{I}$, where $k \in [m]$ occurs with probability $\frac{d}{m}\langle v_k|\rho|v_k\rangle$ (Born’s rule). This generalizes to $\hat{L}_n = \frac{1}{n}\sum_{i=1}^n X_i$, where the matrices X_i are statistically independent. Such sums of random matrices concentrate sharply around their expectation value $\mathbb{E}\hat{L}_n = \rho$, and matrix concentration inequalities [31] quantify this convergence:

$$\Pr [\|\hat{L}_n - \rho\|_\infty \geq \tau] \leq de^{-\frac{3n\tau^2}{8g(d)}} \quad \tau \in [0, 1]. \tag{6}$$

This bound induces a similar operator norm bound for the PLS estimator: the projection is a contraction in operator norm and the shift in eigenvalues is bounded by τ . The resulting bound $\|\hat{\rho}_n - \rho\|_\infty \leq 2\tau$ may be transformed into a trace-norm bound, as shown in

appendix D. This comparison only depends on the (approximate) rank of the states involved. We refer to appendix E for full details of the proof. \square

3.2. Optimal performance guarantee for the uniform POVM

The upper bound in theorem 1 involves a dimension factor d that may be extraneous. This, in turn, introduces an additional $\log(d)$ -gap between equation (5) and existing lower bounds [20]. The dimensional factors emerge because we employ matrix-valued concentration inequalities in the proof. Our second main result shows that we can remove this factor for the *uniform POVM*, which encompasses all rank-one projectors:

Theorem 2. (Convergence of $\hat{\rho}_n$ for the uniform POVM) For uniform POVM measurements, the PLS estimator obeys

$$\Pr \left[\|\hat{\rho}_n - \rho\|_1 \geq \epsilon \right] \leq e^{2.2d - \frac{\epsilon^2 n}{480r^2}} \quad \text{for } \epsilon > 0,$$

where $r = \min \{\text{rank}(\rho), \text{rank}(\hat{\rho}_n)\}$.

This result exactly reproduces the best existing performance guarantees for tomography from independent measurements [22]. The bound follows from standard techniques from high-dimensional probability theory.

Proof sketch of theorem 2. Similarly to the proof of theorem 1, we first establish a (different) concentration bound for the LS estimator \hat{L}_n with respect to the operator norm, and then use the same arguments as in theorem 1 to convert this into a concentration bound for the PLS with respect to the trace norm. The operator norm has a variational formulation:

$$\|\hat{L}_n - \rho\|_\infty = \max_{z \in \mathbb{S}^d} |\langle z | \hat{L}_n - \rho | z \rangle|.$$

The optimization over the unit sphere may be replaced by a maximization over a finite point set, called a *covering net*, whose cardinality scales exponentially in d . For any $z \in \mathbb{S}^d$, $\langle z | \hat{L}_n - \rho | z \rangle$ is a sum of n i.i.d. random variables that exhibit subexponential tail decay. (Measuring the uniform POVM allows us to draw this conclusion.) Standard concentration inequalities yield a tail bound that decays exponentially in the number n of samples. Applying a union bound over all points z_i in the net then ensures $\Pr \left[\|\hat{L}_n - \rho\|_\infty \geq \tau \right] \leq 2e^{c_1 d - c_2 n \tau^2}$. Subsequently, closeness in operator norm for \hat{L}_n may be converted into closeness in trace-norm for $\hat{\rho}_n$ at the cost of an additional (effective) rank factor, cf appendix D. The details of the concentration bound for \hat{L}_n can be found in appendix F. \square

4. Algorithmic considerations

4.1. Explicit solutions for the least squares estimator (3)

Tomographically complete measurements can be viewed as injective linear maps $\mathcal{M} : \mathbb{H}_d \rightarrow \mathbb{R}^m$ with components $[\mathcal{M}(X)]_i = \text{tr}(M_i X)$ specified by Born's rule (1). It is well known that the least-squares problem (3) admits the closed-form solution:

$$\hat{L}_n = (\mathcal{M}^\dagger \mathcal{M})^{-1} (\mathcal{M}^\dagger (f_n)). \quad (7)$$

We evaluate this formula for different measurements and content ourselves with sketching key steps and results. The details can be found in appendix B.

4.1.1. Structured POVMs and the uniform POVM. Also known as two-designs, these systems include highly structured, rank-one POVMs $\{\frac{d}{m}|v_i\rangle\langle v_i|\}_{i=1}^m$, such as symmetric informationally complete POVMs [32], maximal sets of mutually unbiased bases [33], the set of all stabilizer states [34, 35], as well as the uniform POVM. By definition, for $X \in \mathbb{H}_d$, all of the above systems obey

$$\mathcal{M}^\dagger \mathcal{M}(X) = \frac{d^2}{m} \sum_{i=1}^m \langle v_i | X | v_i \rangle |v_i\rangle\langle v_i| = \frac{md}{d+1} (X + \text{tr}(X)\mathbb{I}).$$

These equations can readily be inverted, and equation (7) simplifies to

$$\hat{L}_n = (d+1) \sum_{i=1}^m [f_n]_i |v_i\rangle\langle v_i| - \mathbb{I}. \quad (8)$$

4.1.2. Pauli observables. Fix $d = 2^k$ (k qubits), and let $W_1, \dots, W_{d^2} \in \mathbb{H}_d$ be the set of Pauli observables, comprising all possible k -fold tensor products of the elementary 2×2 Pauli matrices. We can approximate the expectation value $\text{tr}(W_i \rho)$ of each Pauli observable by the empirical mean $\hat{\mu}_i = [f_{n/d^2}^+]_i - [f_{n/d^2}^-]_i$ of the two-outcome POVM $P_i^\pm = \frac{1}{2}(\mathbb{I} \pm W_i)$. Pauli matrices form a unitary operator basis, and the evaluation of equation (7) is simple:

$$\hat{L}_n = \frac{1}{d} \sum_{i=1}^{d^2} \hat{\mu}_i W_i = \frac{1}{d} \sum_{i=1}^{d^2} \left([f_{n/d^2}^+]_i - [f_{n/d^2}^-]_i \right) W_i. \quad (9)$$

4.1.3. Pauli basis measurements. Rather than approximating (global) expectation values, it is possible to perform different combinations of local Pauli measurements. For $d = 2^k$, there are 3^k potential combinations in total. Each of the settings $\vec{s} \in \{x, y, z\}^k$ corresponds to a basis measurement $|b_{\vec{s}}^{(\vec{s})}\rangle\langle b_{\vec{s}}^{(\vec{s})}|$, where $\vec{s} \in \{\pm 1\}^k$ labels the 2^k potential outcomes. The union \mathcal{M} of all 3^k bases obeys $(\mathcal{M}^\dagger \mathcal{M})(X) = 3^k \mathcal{D}_{1/3}^{\otimes k}(X)$, where $\mathcal{D}_{1/3}(X) = \frac{1}{3}\rho + (1 - \frac{1}{3})\frac{\text{tr}(X)}{2}\mathbb{I}$ denotes a single-qubit depolarizing channel. Evaluating equation (7) yields

$$\hat{L}_n = \frac{1}{3^k} \sum_{\vec{s}, \vec{\sigma}} [f_{n/3^k}]_{\vec{\sigma}}^{(\vec{s})} \left(\mathcal{D}_{1/3}^{\otimes k} \right)^{-1} \left(|b_{\vec{\sigma}}^{(\vec{s})}\rangle\langle b_{\vec{\sigma}}^{(\vec{s})}| \right). \quad (10)$$

Finally, we point out that all of these explicit solutions are guaranteed to have unit trace: $\text{tr}(\hat{L}_n) = 1$. They are not positive semidefinite in general.

4.2. Explicit solutions for the projection step (4)

The PLS estimator is defined to be the state closest in Frobenius norm to the least-squares estimator \hat{L}_n . The search (4) admits a simple, *analytic* solution [27]. Define the all-ones vector $\vec{1} \in \mathbb{R}^d$ and the thresholding function $[\cdot]^+$ with components $[\vec{y}]_i^+ = \max\{[\vec{y}]_i, 0\}$. Let $\hat{L}_n = U \text{diag}(\vec{\lambda}) U^\dagger$ be an eigenvalue decomposition. Then

$$\hat{\rho}_n = U \text{diag} \left(\left[\vec{\lambda} - x_0 \vec{1} \right]^+ \right) U^\dagger, \quad (11)$$

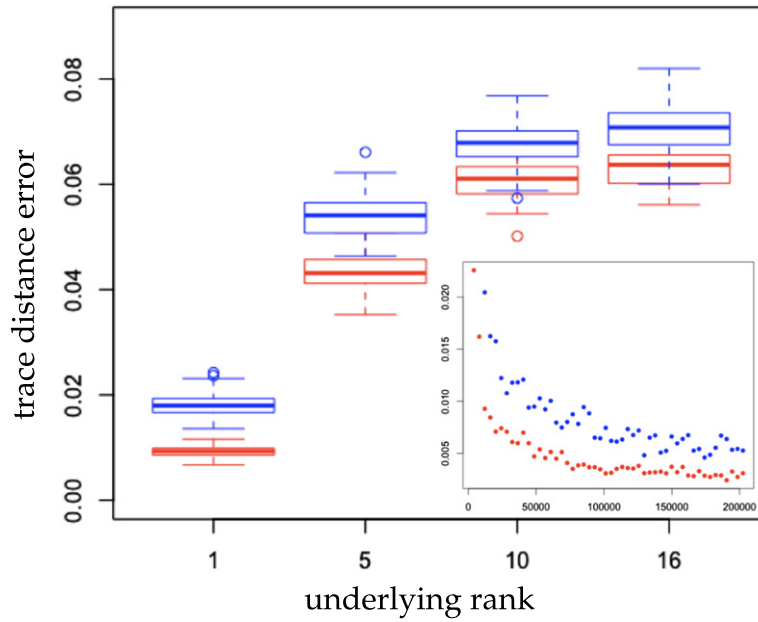


Figure 1. Comparison between PLS (blue) and maximum likelihood (red) for 4-qubit Pauli basis measurements: boxplots of trace distance error for ML vs PLS for 100 datasets generated with random states of rank 1, 5, 10 and 16, and 200 repetitions per setting. Inset: trace distance error as a function of sample size for a pure target state.

where $x_0 \in \mathbb{R}$ is chosen so that $\text{tr}(\hat{\rho}_n) = 1$. The fact that \hat{L}_n itself has unit trace ensures that this solution to equation (4) is unique. The number x_0 may be determined by applying a root-finding algorithm to the non-increasing function $f(x) = 2 + \text{tr}(\hat{L}_n) - dx + \sum_{i=1}^d |\lambda_i - x|$.

4.3. Runtime analysis

The two steps discussed here are inherently *scalable*: just count frequencies to determine the LS estimators (8)–(10) at a total cost of (at most) $\min\{m, n\}$ matrix additions. The subsequent projection onto state-space is a particular type of soft-thresholding. The associated computational cost is dominated by the eigenvalue decomposition and has runtime (at most) $\mathcal{O}(d^3)$.

In summary, forming \hat{L}_n is the dominant cost of a naïve implementation. However, the high degree of structure may allow us to employ techniques from randomized linear algebra [36] to further reduce the cost.

5. Numerical experiments

We numerically compare the performance of PLS to *maximum likelihood* (ML) and *compressed sensing* (CS), respectively. Additional numerical studies for MUBs can be found at the end of this section.

5.1. PLS versus ML

Figure 1 compares ML and PLS for Pauli basis measurements in dimension $d = 2^4$. The trace-norm error incurred by PLS is within a factor of two of ML for low-rank states. A more




Figure 2. Comparison between PLS (blue) and Compressed Sensing (red): 5-qubit Pauli observables and a (random) pure target state. CS was implemented with $m = 256$ randomly selected Pauli observables.

exhaustive simulation study comparing PLS with other ‘truncated’ estimators for a range of states can be found in [25]. This confirms the statistical efficiency of PLS for the family of low rank states which are the focus of this paper.

Computationally, the PLS is significantly faster than ML (implemented using Hradil’s algorithm [1, 2]). While an exhaustive comparative analysis of the computational complexity is beyond the scope of the paper, we can get a rough idea by noting that both PLS and ML involve computing a sum of matrices over all outcomes (the LS estimator in the case of PLS and the ‘ R ’ operator in the case of ML). For both estimators, this constitutes the most time-consuming subroutine. However, this is done only once in the case of PLS (followed by the projection which is relatively fast), while the ML needs to iterate the step a large number of times to achieve convergence to the maximum. Another advantage of PLS is that the LS part could in principle be computed online while the data is gathered, thus reducing the extra time to that of computing the final projection step.

5.2. PLS versus CS

CS is a natural benchmark for low-rank tomography. The papers [11, 12] apply to Pauli observables, and they show that a random choice of $m \geq Crd \log^6(d)$ Pauli observables is sufficient to reconstruct *any* rank- r state. The actual reconstruction is performed by solving a convex optimization problem, e.g. the least-squares fit over the set of quantum states [37, 38]. Numerical studies from [13] suggest that $m = 256$ is appropriate for $d = 2^5$ and $r = 1$. Figure 2 shows that PLS consistently outperforms the CS estimator in this regime. Importantly, PLS was also much faster to evaluate than both, ML and CS.




Figure 3. PLS convergence with MUB measurements: $(d + 1)$ different MUB basis measurements are used to estimate a pure target state. Error decay in the number of samples per basis s seems to be almost independent of the ambient problem dimension d (different colors). *Inset:* ordinary plot of the same data.

5.3. PLS for mutually unbiased bases

Maximal sets of mutually unbiased bases (MUBs) form a structured POVM (two-design) that lends itself to numerical investigation. Efficient algebraic constructions of MUBs exist for dimensions $d = p^k$ where p is a prime number and k is an arbitrary integer [39–41]. To further underline the implicit advantage of low-rank we fix a prime dimension d and choose a pure state uniformly from the Haar measure on the complex unit sphere in d dimensions. We compute the outcome probabilities for each of the $d + 1$ different MUB measurements. We then sample outcomes from each distribution a total of $s = \frac{n}{d+1}$ times and compute the estimator $\hat{\rho}_n$ associated with the total frequency statistics. Figure 3 shows the relation between reconstruction error (in trace distance) and the number of samples per basis on a log–log -scale for a range of different prime dimensions between $d = 100$ and $d = 200$. The significant overlap between the graphs for different dimensions suggests that the rate of convergence scales as d/n in terms of total sample size $n = s(d + 1)$ and dimension d ; if true, this would mean that the additional $\log(d)$ -factor in the norm-one upper bound resulting from theorem 1 may not be necessary, similarly to the upper bound in theorem 2 for uniform measurements. Whether this is the case remains an open question.

6. Conclusion and outlook

Linear inversion is one of the oldest and simplest approaches to solve the practically important task of quantum state tomography. In this work, we focused on related method called *projected least squares* (PLS) that projects the least-squares estimator onto the set of all quantum states. Not only is this estimator numerically cheap, but it comes with strong, non-asymptotic convergence guarantees. These results are derived using concentration inequalities for sums of random matrices, and they exploit the randomness inherent in quantum experiments.

We showed that PLS is competitive, both in theory and in practice. For a variety of measurements, the results match the *best* existing theoretical results for the sampling rate of other tomography methods. In particular, for the uniform POVM, an order of $\frac{r^2 d}{\epsilon^2}$ samples suffice to reconstruct any rank- r state up to accuracy ϵ in trace distance. This result also saturates existing lower bounds [20] on the minimal sampling rate required for *any* tomographic procedure with independent measurements. Numerical studies underline these competitive features.

6.1. Outlook

Corollary 1 is not (yet) optimal. *Bootstrapping* could be used to obtain tighter confidence regions, and the low computational cost of PLS may speed up this process considerably. It also seems fruitful to combine the ideas presented here with recent insights from [42]. The proof of theorem 1 indicates that PLS is stable with respect to time-dependent state generation (*drift*). Moreover, the general principle of PLS can be extended to the related problem of process tomography. We intend to address these points in future work.

Acknowledgments

The authors thank Philippe Faist, Matthias Kleinmann, Anirudh Acharya and Theodore Kypraios for fruitful discussions. Martin Kliesch provided very helpful comments regarding an earlier version of the draft. RK and JT are supported by ONR Award No. N00014-17-12146. RK also acknowledges funding provided by the Institute of Quantum Information and Matter, an NSF Physics Frontiers Center (NSF Grant PHY-1733907).

Appendix A. Organisation of appendices

At the heart of this work is *projected least squares* (PLS)—a simple point estimator for quantum state tomography from informationally complete measurements $\{M_1, \dots, M_m\} \subset \mathbb{H}_d$. PLS is a three-step procedure:

- (a) Estimate outcome probabilities by frequencies.
- (b) Construct the least squares (linear inversion) estimator:

$$\hat{L}_n = \operatorname{argmin}_{X \in \mathbb{H}_d} \sum_{i=1}^m (f_i - \operatorname{tr}(M_i X))^2. \quad (\text{A.1})$$

- (c) Project onto the set of all quantum states:

$$\hat{\rho}_n = \operatorname{argmin}_{\sigma \text{ is a quantum state}} \|\hat{L}_n - \sigma\|_2. \quad (\text{A.2})$$

We analyze the performance of PLS for a variety of concrete measurement scenarios: *structured POVMs*, *Pauli observables*, *Pauli basis measurements* and the *uniform POVM*. For each of them, $\hat{\rho}_n$ may be equipped with rigorous non-asymptotic confidence regions in trace distance. In this appendix, we complement the rather succinct presentation in the main text with additional explanations, motivations and more detailed arguments.

Outline: in appendix B we provide explicit least squares solutions (A.1) for the different measurements. We also review essential features and properties of the individual scenarios to provide context.

Appendix C contains the main conceptual insight of this work: least squares estimators may be interpreted as sums of independent random matrices—the randomness is due to the fundamental laws of quantum mechanics (Born’s rule). This allows us to apply strong matrix-valued concentration inequalities to show that, with high probability, \hat{L}_n is close to the true target state in operator norm.

Appendix D is devoted to showing that closeness of \hat{L}_n in operator norm implies closeness of $\hat{\rho}_n$ in trace norm.

We combine these two insights in appendix E to arrive at the main result of this work: convergence guarantees for the PLS estimator in trace norm. The result derived there is a strict generalization of the main result quoted in the main text. It extends to the notion of *effective rank* which may be beneficial in concrete applications. We illustrate this potential benefit with a caricature of a faulty state preparation apparatus.

In appendix F yet stronger convergence guarantees for the uniform POVM are derived. The proof technique is completely different and we believe that it may be of independent interest to the community.

Appendix B. Closed-form expressions for least squares estimators

As outlined in the main text, any POVM measurement can be viewed as a linear map $\mathcal{M} : \mathbb{H}_d \rightarrow \mathbb{R}^m$, defined component-wise as $[\mathcal{M}(X)]_i = \text{tr}(M_i X)$ for $i \in [m]$. We will focus our attention on injective, or tomographically complete, measurements. In this case, the least squares estimator (A.1) admits a unique solution:

$$\hat{L}_n = (\mathcal{M}^\dagger \mathcal{M})^{-1} \mathcal{M}^\dagger (f_n),$$

where $f_n \in \mathbb{R}^m$ subsumes the individual frequency estimates. In this section, we evaluate this formula explicitly for different types of prominent measurements.

B.1. The uniform POVM and two-designs

The *uniform/covariant* POVM in d -dimensions corresponds to the union of all (properly re-normalized) rank-one projectors: $\{d|v\rangle\langle v|dv\}_{v \in \mathbb{S}^d}$. Here, dv denotes the unique, unitarily invariant, measure on the complex unit sphere induced by the Haar measure (over the unitary group $U(d)$). Its high degree of symmetry allows for analyzing this POVM by means of powerful tools from representation theory. This is widely known, see e.g. [43, 44], but we include a short presentation here to be self-contained. Define the *frame operator of order* k : $F_{(k)} = \int_{\mathbb{S}^d} (|v\rangle\langle v|)^{\otimes k} dv \in \mathbb{H}_d^{\otimes k}$. Unitary invariance of dv implies that this frame operator commutes with every k -fold tensor product of a unitary matrix $U \in U(d)$:

$$U^{\otimes k} F_{(k)} = \int_{\mathbb{S}^d} (U|v\rangle\langle v|)^{\otimes k} dv = \int_{\mathbb{S}^d} (|\tilde{v}\rangle\langle \tilde{v}|U)^{\otimes k} d\tilde{v} = F_{(k)} U^{\otimes k}.$$

Here, we have used a change of variables ($\tilde{v} = Uv$) together with the fact that dv is unitarily invariant ($d\tilde{v} = dv$). Schur’s lemma—one of the most fundamental tools in representation theory—states that any matrix that commutes with every element of a given group representation must be proportional to a sum of the projectors onto the associated irreducible representations (irreps). For the task at hand, the representation of interest is the diagonal representation

of the unitary group: $U \mapsto U^{\otimes k}$ for all $U \in U(d)$. This representation affords, in general, many irreps that may be characterized using Schur–Weyl duality. The *symmetric subspace* $\text{Sym}^{(k)} \subset (\mathbb{C}^d)^{\otimes k}$ is one of them and corresponds to the subspace of all vectors that are invariant under permuting tensor factors. Crucially, $F_{(k)}$ is an average over rank-one projectors onto vectors $|v\rangle^{\otimes k} \in \text{Sym}^{(k)}$ and, therefore, its range must be contained entirely within $\text{Sym}^{(k)}$. Combining this with the assertion of Schur’s lemma then yields

$$F_{(k)} = \int_{\mathbb{S}^d} (|v\rangle\langle v|)^{\otimes k} dv = \binom{d+k-1}{k}^{-1} P_{\text{Sym}^{(k)}} \quad k \in \mathbb{N}, \tag{B.1}$$

The pre-factor $\binom{d+k-1}{k}^{-1} = \dim(\text{Sym}^{(k)})^{-1}$ follows from the fact that $F_{(k)}$ has unit trace.

This closed-form expression is very useful. In particular, it implies that the uniform POVM $\{d|v\rangle\langle v|\}_{v \in \mathbb{S}^d}$ is almost an isometry. Fix $X \in \mathbb{H}_d$ and compute

$$\begin{aligned} (d+1) \int_{\mathbb{S}^d} d\langle v|X|v\rangle|v\rangle\langle v|dv &= (d+1)d\text{tr}_1 \left(\int_{\mathbb{S}^d} (|v\rangle\langle v|)^{\otimes 2} dv (X \otimes \mathbb{I}) \right) \\ &= 2\text{tr}_1 (P_{\text{Sym}^{(2)}}(X \otimes \mathbb{I})), \end{aligned} \tag{B.2}$$

where $\text{tr}_1(A \otimes B) = \text{tr}(A)B$ denotes the partial trace over the first tensor factor. The projector onto the totally symmetric subspace of two parties has an explicit representation: $P_{\text{Sym}^{(2)}} = \frac{1}{2}(\mathbb{I} + \mathbb{F})$, where \mathbb{F} denotes the *flip operator*, i.e. $\mathbb{F}|x\rangle \otimes |y\rangle = |y\rangle \otimes |x\rangle$ for all $|x\rangle, |y\rangle \in \mathbb{C}^d$ and extend it linearly to the entire tensor product. Inserting this explicit characterization into equation (B.2) yields

$$(d+1) \int_{\mathbb{S}^d} d\langle v|X|v\rangle|v\rangle\langle v|dv = \text{tr}_1((\mathbb{I} + \mathbb{F})(X \otimes \mathbb{I})) = X + \text{tr}(X)\mathbb{I}. \tag{B.3}$$

We emphasize that the full symmetry of the uniform POVM is not required to derive this formula: equation (B.1) for $k = 2$ is sufficient. This motivates the following definition:

Definition 1 (Two-design). A (finite) set of m rank-one projectors $\{|v_i\rangle\langle v_i|\}_{i=1}^m$ is called a (*complex-projective*) *two-design* if

$$\frac{1}{m} \sum_{i=1}^m (|v_i\rangle\langle v_i|)^{\otimes 2} = \binom{d+1}{2}^{-1} P_{\text{Sym}^{(2)}}.$$

Each two-design is proportional to a POVM $\mathcal{M} = \{\frac{d}{m}|v_i\rangle\langle v_i|\}_{i=1}^m$. Clearly, each element is positive semidefinite and $\frac{d}{m} \sum_{i=1}^m |v_i\rangle\langle v_i| = \mathbb{I}$ follows from taking the partial trace of the defining property. Viewed as maps $\mathcal{M} : \mathbb{H}_d \rightarrow \mathbb{R}^m$, these POVMs obey

$$\mathcal{M}^\dagger \mathcal{M}(X) = \frac{d^2}{m^2} \sum_{i=1}^m \langle v_i|X|v_i\rangle|v_i\rangle\langle v_i| = \frac{d(X + \text{tr}(X)\mathbb{I})}{(d+1)m} \quad \text{for } X \in \mathbb{H}_d.$$

This can be readily inverted

$$(\mathcal{M}^\dagger \mathcal{M})^{-1}(X) = \frac{m}{d} ((d+1)X - \text{tr}(X)\mathbb{I}). \tag{B.4}$$

and the least-squares estimator becomes

$$\begin{aligned}\hat{L}_n &= (\mathcal{M}^\dagger \mathcal{M})^{-1} (\mathcal{M}^\dagger(f)) = (\mathcal{M}^\dagger \mathcal{M})^{-1} \left(\frac{d}{m} \sum_{i=1}^m f_i |v_i\rangle \langle v_i| \right) \\ &= \sum_{i=1}^m f_i ((d+1)|v_i\rangle \langle v_i| - \text{tr}(|v_i\rangle \langle v_i|) \mathbb{I}) = (d+1) \sum_{i=1}^m f_i |v_i\rangle \langle v_i| - \mathbb{I} \quad (\text{B.5})\end{aligned}$$

for any frequency vector $f_n \in \mathbb{R}^m$. Mathematically, this is a consequence of the fact that two-design POVMs ‘almost’ form a tight frame on \mathbb{H}_d . The close connection to well-behaved, tomographically complete, rank-one POVMs has spurred considerable interest in the identification of two-designs. Over the past decades, the following concrete examples have been identified:

- (a) *Equiangular lines (SIC POVMs)*: a family of m unit vectors $|v_1\rangle, \dots, |v_m\rangle \in \mathbb{S}^d$ is equiangular, if $|\langle v_i, v_j \rangle|^2$ is constant for all $i \neq j$. The minimal cardinality of such a set is $m = d^2$ in which case the angle must be fixed: $|\langle v_i, v_j \rangle|^2 = \frac{1}{d+1}$. Such maximal sets of equiangular lines are known to form two-designs [32] and have been termed *symmetric, informationally complete (SIC) POVMs*. This nomenclature underlines the importance of equation (B.3) for the original quantum motivation of the study of equiangular lines. While several explicit constructions of SIC POVMs exist, the general question of their existence remains an intriguing open problem.
- (b) *Mutually unbiased bases (MUBs)*: two orthonormal bases $\{|b_i\rangle\}_{i=1}^d$ and $\{|c_i\rangle\}_{i=1}^d$ of \mathbb{C}^d are *mutually unbiased* if $|\langle b_i, c_j \rangle|^2 = \frac{1}{d}$ for all $1 \leq i, j \leq d$. The study of such mutually unbiased bases (MUBs) has a rich history in quantum mechanics that dates back to Schwinger [45]. It is known that at most $(d+1)$ pairwise mutually unbiased bases can exist in dimension d and explicit algebraic constructions are known for prime power dimensions ($d = p^k$). Klappenecker and Roettler [33] showed that maximal sets of MUBs are guaranteed to form two-designs.
- (c) *Stabilizer states (STABS)*: the stabilizer formalism is one of the cornerstones of quantum computation, fault tolerance and error correction, see e.g. [46]. Let \mathcal{P}_k be the Pauli group on k qubits ($d = 2^k$), i.e. the group generated by k -fold tensor products of the elementary Pauli matrices, \mathbb{I} and $i\mathbb{I}$. It is then possible to find maximal abelian subgroups $\mathcal{S} \subset \mathcal{P}_k$ of size $d = 2^k$. Since all matrices $W \in \mathcal{S}$ commute, they can be simultaneously diagonalized and determine a single unit vector which is the joint eigenvector with eigenvalue $+1$ of all the matrices in \mathcal{S} (provided that $-\mathbb{I} \notin \mathcal{S}$). Such vectors are called *stabilizer states* (STAB) and the group $\mathcal{S} \subset \mathcal{P}_k$ is its associated *stabilizer group*. A total of $m = 2^k \prod_{i=0}^{k-1} (d^i + 1) = 2^{\frac{1}{2}k^2 + o(k)}$ different stabilizer states can be generated this way. The union of all of them is actually known to form a three-design [47–49] and, therefore, also a two-design. The latter is also a consequence of earlier results [34, 35]

B.2. Pauli observables

For $d = 2^k$, the Pauli matrices $W_1, \dots, W_{d^2} \in \mathbb{H}_d$ arise from all possible k -fold tensor products of elementary Pauli matrices $\{\mathbb{I}, \sigma_x, \sigma_y, \sigma_z\} \subset \mathbb{H}_2$. They are well-known to form a unitary operator basis:

$$X = \frac{1}{d} \sum_{i=1}^{d^2} \text{tr}(W_i X) W_i, \quad (\text{B.6})$$

for all $X \in \mathbb{H}_d$. While they do constitute observables, Pauli matrices by themselves are not POVMs. However, every observable W_i may be associated with a two-outcome POVM $\mathcal{M}_i = \{P_i^\pm\} = \{\frac{1}{2}(\mathbb{I} \pm W_i)\}$. The union $\bigcup_{i=1}^{d^2} \mathcal{M}_i$ of all these two-outcome POVMs constitutes a linear map $\mathcal{M} : \mathbb{H}_d \rightarrow \mathbb{R}^{2^m}$ that obeys

$$\mathcal{M}^\dagger \mathcal{M}(X) = \sum_{i=1}^{d^2} \sum_{o=\pm} \text{tr}(P_i^o X) P_i^o = \sum_{i=1}^{d^2} \frac{1}{2} (\text{tr}(X)\mathbb{I} + \text{tr}(W_i X)W_i) = \frac{1}{2} (d^2 \text{tr}(X)\mathbb{I} + X),$$

where the equality follows from equation (B.6). Once more, this expression can be readily inverted:

$$(\mathcal{M}^\dagger \mathcal{M})^{-1}(X) = \frac{2}{d} X - \frac{2 \text{tr}(X)}{d^2 + 1} \mathbb{I}.$$

Before we continue, we note that one Pauli matrix is equal to the identity, say $W_1 = \mathbb{I}$, and the associated POVM is trivial. Hence, we suppose that n copies of ρ are distributed equally among all $d^2 - 1$ non-trivial two-Outcome POVMs \mathcal{M}_i . We denote the resulting frequencies by $[f]_i^\pm$ and suppress the dependence on the number of samples. Then, the explicit solution to the least squares problem becomes

$$\begin{aligned} \hat{L}_n &= (\mathcal{M}^\dagger \mathcal{M})^{-1} (\mathcal{M}^\dagger(f_n)) = (\mathcal{M}^\dagger \mathcal{M})^{-1}(\mathbb{I}) + \sum_{i=2}^{d^2} \sum_{o=\pm} [f]_i^o (\mathcal{M}^\dagger \mathcal{M})^{-1}(P_i^o) \\ &= \frac{2}{d(d^2 + 1)} \mathbb{I} + \sum_{i=2}^{d^2} [f]_i^+ \left(\frac{1}{d} (\mathbb{I} + W_i) - \frac{d}{d^2 + 1} \mathbb{I} \right) + \sum_{i=2}^{d^2} [f]_i^- \left(\frac{1}{d} (\mathbb{I} - W_i) - \frac{d}{d^2 + 1} \mathbb{I} \right) \\ &= \frac{1}{d} \sum_{i=2}^{d^2} ([f]_i^+ - [f]_i^-) W_i + \frac{2 + \sum_{i=2}^{d^2} ([f]_i^+ + [f]_i^-)}{d(d^2 + 1)} \mathbb{I}. \end{aligned}$$

We can simplify this expression further by noticing that each two-outcome POVM is dichotomic: either $+$ or $-$ is observed for every run. This implies $[f]_i^+ + [f]_i^- = 1$ and, by extension, $\sum_{i=2}^{d^2} ([f]_i^+ + [f]_i^-) = d^2 - 1$. Hence,

$$\hat{L}_n = \frac{1}{d} \sum_{i=2}^{d^2} ([f]_i^+ - [f]_i^-) W_i + \frac{1}{d} \mathbb{I} = \frac{1}{d} \sum_{i=1}^{d^2} ([f]_i^+ - [f]_i^-) W_i, \quad (\text{B.7})$$

because $[f]_1^+ = 1$. This is the formula from the main text and has a compelling interpretation: the difference $\hat{\mu}_i = [f]_i^+ - [f]_i^-$ is an empirical estimate for the expectation value $\mu_i = \text{tr}(W_i X)$ of the i th Pauli observable. Finally, note that this estimator is again unbiased with respect to random fluctuations in the sample statistics:

$$\mathbb{E}[\hat{L}_n] = \frac{1}{d} \sum_{i=1}^d \text{tr}(W_i \rho) W_i = \rho. \quad (\text{B.8})$$

B.3. Pauli basis measurements

Before considering the general case, we find it instructive to consider the single qubit case in more detail. For now, fix $d = 2$ and note that there are three non-trivial Pauli matrices σ_s

with $s \in \{x, y, z\}$. We may associate each σ_s with a two-outcome POVM that is also a basis measurement: $\frac{1}{2}(\mathbb{I} \pm \sigma_s) = |b_{\pm}^{(s)}\rangle\langle b_{\pm}^{(s)}|$. For $s \neq s'$

$$\left| \langle b_{\pm}^{(s)}, b_{o'}^{(\pm)} \rangle \right|^2 = \frac{1}{4} \text{tr}((\mathbb{I} \pm \sigma_s)(\mathbb{I} \pm \sigma_{s'})) = \frac{1}{2},$$

because $\sigma_s, \sigma_{s'}$ and $\sigma_s \sigma_{s'}$ have vanishing trace. This implies that the six vectors $|b_o^{(s)}\rangle\langle b_o^{(s)}|$ with $o \in \{\pm 1\}$ form a maximal set of $3 = (d + 1)$ mutually unbiased bases. Such vector sets form spherical two-designs and equation (B.3) ensures for any $X \in \mathbb{H}_d$

$$\sum_{s,o} \langle b_o^{(s)}|X|b_o^{(s)}\rangle |b_o^{(s)}\rangle\langle b_o^{(s)}| = X + \text{tr}(X)\mathbb{I} = 3 \left(\frac{1}{3}X + \frac{2}{3} \frac{\text{tr}(X)}{2} \mathbb{I} \right) = 3\mathcal{D}_{1/3}(X). \tag{B.9}$$

Here, $\mathcal{D}_{1/3}(X) = \frac{1}{3}X + (1 - \frac{1}{3})\frac{\text{tr}(X)}{2}\mathbb{I}$ denotes a single-qubit depolarizing channel with loss parameter $p = \frac{1}{3}$.

This behavior extends to multi-qubit systems, i.e. $d = 2^k$. Suppose that we perform k local (single-qubit) Pauli measurements on a k -qubit state $\rho \in \mathbb{H}_d$. Then, there are a total of k potential combinations that we label by a string $\vec{s} = (s_1, \dots, s_k) \in \{x, y, z\}^k$. Each of them corresponds to a POVM $\mathcal{M}^{(\vec{s})}$ with $2^k = d$ outcomes that we label by $\vec{o} = (o_1, \dots, o_k) \in \{\pm 1\}^k$. The POVM element associated with index \vec{s} and outcome \vec{o} has an appealing tensor-product structure: $|b_{\vec{o}}^{(\vec{s})}\rangle\langle b_{\vec{o}}^{(\vec{s})}| = \bigotimes_{i=1}^k |b_{o_i}^{(s_i)}\rangle\langle b_{o_i}^{(s_i)}|$. Let $\mathcal{M} = \bigcup_{\vec{s}} \mathcal{M}^{(\vec{s})} : \mathbb{H}_d \rightarrow \mathbb{R}^{3^k} \times \mathbb{R}^{2^k}$ denote the union of all such basis measurements. Then, the following formula is true for tensor product matrices $X = \bigotimes_{i=1}^k X_i$ and $X_i \in \mathbb{H}_2$:

$$\begin{aligned} \mathcal{M}^\dagger \mathcal{M}(X) &= \sum_{\vec{s}, \vec{o}} \langle b_{\vec{o}}^{(\vec{s})}|X|b_{\vec{o}}^{(\vec{s})}\rangle |b_{\vec{o}}^{(\vec{s})}\rangle\langle b_{\vec{o}}^{(\vec{s})}| = \bigotimes_{i=1}^k \left(\sum_{s_i, o_i} \langle b_{o_i}^{(s_i)}|X_i|b_{o_i}^{(s_i)}\rangle |b_{o_i}^{(s_i)}\rangle\langle b_{o_i}^{(s_i)}| \right) \\ &= 3^k \bigotimes_{i=1}^k \mathcal{D}(X_i) = 3^k \mathcal{D}_{1/3}^{\otimes k}(X), \end{aligned} \tag{B.10}$$

where we have used equation (B.9). Linear extension ensures that this formula remains valid for arbitrary matrices $X \in \mathbb{H}_d$. Since the single qubit depolarizing channel is invertible, the same is true for its k -fold tensor product and we conclude

$$(\mathcal{M}^\dagger \mathcal{M})^{-1}(X) = \frac{1}{3^k} \left(\mathcal{D}_{1/3}^{\otimes k} \right)^{-1}(X).$$

Inserting this explicit expression into the closed-form expression for the least squares estimator yields

$$\hat{L}_n = (\mathcal{M}^\dagger \mathcal{M})^{-1} \left(\mathcal{M}^\dagger(\vec{f}) \right) = \frac{1}{3^k} \sum_{\vec{s}, \vec{o}} [f]_{\vec{o}}^{(\vec{s})} \left(\mathcal{D}_{1/3}^{\otimes k} \right)^{-1} \left(|b_{\vec{o}}^{(\vec{s})}\rangle\langle b_{\vec{o}}^{(\vec{s})}| \right),$$

as advertised in the main text. Here, $[f]_{\vec{o}}^{(\vec{s})}$ is assumed to be a frequency approximation to $p_{\vec{o}}^{(\vec{s})} = \langle b_{\vec{o}}^{(\vec{s})}|\rho|b_{\vec{o}}^{(\vec{s})}\rangle$.

We conclude this section with a single-qubit observations that allows for characterizing this expression in a more explicit fashion. Note that one may rewrite $\mathcal{D}_{1/3}(X)$ as

$\frac{\text{tr}(X)}{2}\mathbb{I} + \frac{1}{6}\sum_s \text{tr}(\sigma_s X)\sigma_s$. This facilitates the computation of the single-qubit inverse:

$$\mathcal{D}_{1/3}^{-1}(X) = \frac{\text{tr}(X)}{2}\mathbb{I} + \frac{3}{2}\sum_s \text{tr}(\sigma_s X)\sigma_s$$

and, in particular

$$\mathcal{D}_{1/3}^{-1}\left(|b_{\pm}^{(s)}\rangle\langle b_{\pm}^{(s)}|\right) = \frac{1}{2}\left(\mathcal{D}^{-1}(\mathbb{I}) \pm \mathcal{D}^{-1}(\sigma_s)\right) = \frac{1}{2}(\mathbb{I} \pm 3\sigma_s) = 3|b_{\pm}^{(s)}\rangle\langle b_{\pm}^{(s)}| - \mathbb{I}.$$

This ensures

$$\hat{L}_n = \frac{1}{3^k} \sum_{\vec{s}, \vec{o}} [f]_{\vec{o}}^{(\vec{s})} \bigotimes_{i=1}^k (3|b_{o_i}^{(s_i)}\rangle\langle b_{o_i}^{(s_i)}| - \mathbb{I}). \tag{B.11}$$

which, again, is an unbiased estimator with respect to random fluctuations in the sample statistics.

Finally, we also point out another consequence that will be important later on:

$$\left(\mathcal{D}_{1/3}^{-1}\left(|b_o^{(s)}\rangle\langle b_o^{(s)}|\right)\right)^2 = 5\mathcal{D}_{3/5}\left(|b_o^{(s)}\rangle\langle b_o^{(s)}|\right), \tag{B.12}$$

where $\mathcal{D}_{3/5}$ is another single-qubit depolarizing channel.

Appendix C. The matrix Bernstein inequality and concentration in operator norm

Scalar concentration inequalities provide sharp bounds on the probability of a sum of independent random variables deviating from their mean. Classical examples include Hoeffding’s, Chernoff’s and Bernstein’s inequality—all of which have found widespread use in a variety of scientific disciplines. The main results of this work are based on a matrix generalizations of these classical statements—in particular the *matrix Bernstein inequality* developed by one of the authors, see [31, theorem 1.4].

Theorem 3 (Matrix Bernstein inequality). *Consider a sequence of n independent, Hermitian random matrices $A_1, \dots, A_n \in \mathbb{H}_d$. Assume that each A_i satisfies*

$$\mathbb{E}[A_i] = 0 \quad \text{and} \quad \|A_i\|_{\infty} \leq R \text{ almost surely.}$$

Then, for any $t > 0$

$$\Pr \left[\left\| \sum_{i=1}^n (A_i - \mathbb{E}[A_i]) \right\|_{\infty} \geq t \right] \leq \begin{cases} d \exp\left(-\frac{3t^2}{8\sigma^2}\right) & t \leq \frac{\sigma^2}{R}, \\ d \exp\left(-\frac{3t}{8R}\right) & t \geq \frac{\sigma^2}{R}, \end{cases}$$

where $\sigma^2 = \left\| \sum_{i=1}^n \mathbb{E}[A_i^2] \right\|_{\infty}$.

First results of this kind originate in Banach space theory [50–53] and were later independently developed in quantum information theory [54, 55]. Further advances by Oliveira [56] and one of the authors [31] led to the result that we employ here. We refer to the monograph [31] for a detailed exposition of related work and history.

Similar to the scalar Bernstein inequality, the tail behavior in theorem 3 consists of two regimes. Small deviations are suppressed in a subgaussian fashion, while larger deviations follow a subexponential decay. The ratio $\frac{\sigma^2}{R}$ marks the transition from one regime into the other. We also note in passing that this result recovers the scalar Bernstein inequality for $d = 1$ ($\mathbb{H}_1 \simeq \mathbb{R}$).

C.1. Concentration for structured POVM measurements

For structured measurements (two-designs) we may rewrite the (plain) least squares estimator (B.5) as

$$\hat{L}_n = (d + 1) \sum_{i=1}^m [f_n]_i |v_i\rangle\langle v_i| - \mathbb{I} = \frac{1}{n} \sum_{i=1}^n X_i,$$

where each X_i is an i.i.d. copy of the random matrix $X \in \mathbb{H}_d$ that assumes $(d + 1)|v_k\rangle\langle v_k| - \mathbb{I}$ with probability $\frac{d}{m} \langle v_k | \rho | v_k \rangle$ for all $k \in [m]$. Unbiasedness with respect to the sample statistics ensures $\mathbb{E} [\hat{L}_n] = \mathbb{E} [X] = \rho$. Hence, $\hat{L}_n - \rho$ is a sum of i.i.d., centered random matrices $\frac{1}{n} (X_i - \mathbb{E} [X_i])$. These obey

$$\frac{1}{n} \|X_i - \mathbb{E} [X_i]\|_\infty = \frac{1}{n} \|(d + 1)|v_k\rangle\langle v_k| - \mathbb{I} - \rho\|_\infty \leq \frac{d}{n} =: R,$$

where $k \in [m]$ is arbitrary. Next, note that the random matrix X obeys

$$\mathbb{E} [(X - \mathbb{E} [X])^2] = \mathbb{E} [X^2] - \mathbb{E} [X]^2 = \mathbb{E} [X^2] - \rho^2$$

and also

$$\begin{aligned} \mathbb{E} [X^2] &= \sum_{k=1}^m \frac{d}{m} \langle v_k | \rho | v_k \rangle ((d + 1)|v_k\rangle\langle v_k| - \mathbb{I})^2 = \frac{d(d^2 - 1)}{m} \sum_{k=1}^m \langle v_k | \rho | v_k \rangle |v_k\rangle\langle v_k| + \mathbb{I} \\ &= (d - 1)(\rho + \mathbb{I}) + \mathbb{I}, \end{aligned}$$

according to equation (B.3). This allows us to bound the variance parameter:

$$\left\| \sum_{i=1}^n \mathbb{E} \left[\frac{1}{n} (X - \mathbb{E} [X]) \right]^2 \right\|_\infty = \frac{1}{n} \|(d - 1)\rho + d\mathbb{I} - \rho^2\|_\infty \leq \frac{2d}{n} =: \sigma^2.$$

The ratio $\frac{\sigma^2}{R} = 2$ indicates that any choice of $\tau \in [0, 2]$ will fall into the subgaussian regime of the matrix Bernstein inequality and theorem 3 yields

$$\Pr \left[\|\hat{L}_n - \rho\|_\infty \geq \tau \right] = \Pr \left[\left\| \frac{1}{n} \sum_{i=1}^n (X_i - \mathbb{E} [X_i]) \right\|_\infty \geq \tau \right] \leq d e^{-\frac{3\tau^2 n}{16d}}. \quad (\text{C.1})$$

C.2. Concentration for (global) Pauli observables

We assume that the total number of samples n is distributed equally among the d^2 different Pauli measurements. Similar to before, unbiasedness (B.8) and the explicit characterization of the LI estimator (B.7) allow us to write

$$\hat{L}_n - \rho = \sum_{k=1}^{d^2} \frac{1}{n} \sum_{i=1}^{n/d^2} \left(X_i^{(k)} - \mathbb{E} \left[X_i^{(k)} \right] \right).$$

Here, each $X_i^{(k)}$ is an independent instance of the random matrix $X^{(k)} = \pm dW_k$ with probability $\frac{1}{2} (1 \pm \text{tr}(W_k \rho))$ each. This is a sum of centered random matrices that are independent, but in general not identically distributed. However, independence alone suffices for applying theorem 3. We note in passing that this would not be the case for earlier (weaker) versions of the matrix Bernstein inequality. Bound

$$\frac{1}{n} \left\| X_i^{(k)} - \mathbb{E} \left[X_i^{(k)} \right] \right\|_{\infty} = \frac{d}{n} \left\| (1 \pm \text{tr}(W_i \rho)) \mathbb{I} \right\|_{\infty} \leq \frac{2d}{n} =: R$$

and use $\mathbb{E} \left[(X - \mathbb{E}[X])^2 \right] \leq \mathbb{E} \left[X^2 \right]$ (in the positive semidefinite order) to considerably simplify the variance computation:

$$\frac{d^2}{n} \left\| \sum_{k=1}^{d^2} \mathbb{E} \left[(X^{(k)})^2 \right] \right\|_{\infty} = \frac{1}{n} \left\| \sum_{k=1}^{d^2} \mathbb{I} \right\|_{\infty} = \frac{d^2}{n} =: \sigma^2,$$

because $(X^{(k)})^2 = \frac{1}{d^2} \mathbb{I}$. Applying theorem 3 yields

$$\Pr \left[\left\| \hat{L}_n - \rho \right\|_{\infty} \geq \tau \right] \leq d e^{-\frac{3\tau^2 \tilde{n}}{8}} \quad \tau \in [0, d/2].$$

C.3. Concentration for Pauli-basis measurements

Once more, we assume that the total budget of samples n is distributed equally among all 3^k Pauli basis choices. Unbiasedness of the LI estimator together with the explicit description (B.11) allows us to once more interpret $\hat{L}_n - \rho$ as a sum of independent, centered random matrices:

$$\hat{L}_n - \rho = \sum_{\vec{s}} \frac{1}{n} \sum_{i=1}^{n/3^k} \left(X_i^{(\vec{s})} - \mathbb{E} \left[X_i^{(\vec{s})} \right] \right).$$

For each $\vec{s} \in \{x, y, z\}^k$, $X_i^{(\vec{s})}$ is an independent copy of the random matrix

$$X^{(\vec{s})} = \bigotimes_{i=1}^k \left(3 |b_{\vec{o}_i}^{(s_i)}\rangle \langle b_{\vec{o}_i}^{(s_i)}| - \mathbb{I} \right)$$

with probability $\langle b_{\vec{o}}^{(\vec{s})} | \rho | b_{\vec{o}}^{(\vec{s})} \rangle$ for each $\vec{o} \in \{\pm 1\}^k$. Jensen's inequality implies

$$\frac{1}{n} \left\| X^{(\vec{s})} - \mathbb{E} \left[X^{(\vec{s})} \right] \right\|_{\infty} \leq \frac{2}{n} \left\| X^{(\vec{s})} \right\|_{\infty} = \frac{2}{n} \prod_{i=1}^k \left\| 3 |b_{\vec{o}_i}^{(s_i)}\rangle \langle b_{\vec{o}_i}^{(s_i)}| - \mathbb{I} \right\|_{\infty} = \frac{2^{k+1}}{n} =: R.$$

For the variance, we once more use $\mathbb{E} \left[(X - \mathbb{E}[X])^2 \right] \leq \mathbb{E} \left[X^2 \right]$ and compute

$$\begin{aligned}
 \frac{1}{n} \sum_{\vec{s}} \mathbb{E} (X^{(\vec{s})})^2 &= \sum_{\vec{s}, \vec{o}} \langle b_{\vec{o}}^{(\vec{s})} | \rho | b_{\vec{o}}^{(\vec{s})} \rangle \bigotimes_{i=1}^k \left(\mathcal{D}^{-1} \left(|b_{\vec{o}}^{(\vec{s})}\rangle \langle b_{\vec{o}}^{(\vec{s})}| \right) \right)^2 \\
 &= \sum_{\vec{s}} \sum_{\vec{o}} \langle b_{\vec{o}}^{(\vec{s})} | \rho | b_{\vec{o}}^{(\vec{s})} \rangle \bigotimes_{i=1}^k \frac{5}{3} \mathcal{D}_{3/5} \left(|b_{\vec{o}}^{(\vec{s})}\rangle \langle b_{\vec{o}}^{(\vec{s})}| \right) \\
 &= 5^k \mathcal{D}_{3/5}^{\otimes k} \left(\frac{1}{3^k} \sum_{\vec{s}, \vec{o}} \langle b_{\vec{o}}^{(\vec{s})} | \rho | b_{\vec{o}}^{(\vec{s})} \rangle |b_{\vec{o}}^{(\vec{s})}\rangle \langle b_{\vec{o}}^{(\vec{s})}| \right) \\
 &= 5^k \mathcal{D}_{3/5}^{\otimes k} \left(\mathcal{D}_{1/3}^{\otimes k} (\rho) \right) = \frac{5^k}{n} \mathcal{D}_{1/5}^{\otimes k} (\rho),
 \end{aligned}$$

where we have used equation (B.12) and the fact that the combination of two depolarizing channels is again a depolarizing channel. This expression can be evaluated explicitly. For $\alpha \subset [k]$, let $\text{tr}_\alpha(\rho)$ denote the partial trace over all indices contained in α . Then, for $X = \bigotimes_{j=1}^k X_j \in \mathbb{H}_d$

$$5^k \mathcal{D}_{1/5}^{\otimes k} \left(\bigotimes_{j=1}^k X_j \right) = \bigotimes_{j=1}^k (\text{tr}(X_j) \mathbb{I} + X_j) = \sum_{\alpha \subset [k]} 2^{|\alpha|} \text{tr}_\alpha(X) \otimes \mathbb{I}^{\otimes \alpha}$$

and this extends linearly to all of $\mathbb{H}_d \simeq \mathbb{H}_2^{\otimes k}$. Consequently,

$$\begin{aligned}
 \left\| \frac{1}{n} \sum_{\vec{s}} \mathbb{E} \left[(X^{(\vec{s})})^2 \right] \right\|_\infty &= \frac{1}{n} \left\| \sum_{\alpha \subset [k]} 2^{|\alpha|} \text{tr}_\alpha(\rho) \otimes \mathbb{I}^{\otimes \alpha} \right\|_\infty \leq \frac{1}{n} \sum_{\alpha \subset [k]} 2^{|\alpha|} \|\text{tr}_\alpha(\rho)\|_\infty \|\mathbb{I}^{\otimes \alpha}\|_\infty \\
 &\leq \frac{1}{n} \sum_{\alpha \subset [k]} 2^{|\alpha|} = \frac{1}{n} \sum_{j=0}^k \binom{k}{j} 2^j = \frac{(2+1)^k}{n} = \frac{3^k}{n} =: \sigma^2.
 \end{aligned}$$

This estimate is actually tight for pure product states of the form $\rho = (|\psi\rangle\langle\psi|)^{\otimes k}$. We may now apply theorem 3 to conclude

$$\Pr \left[\|\hat{L}_n - \rho\|_\infty \geq \tau \right] \leq d \exp \left(-\frac{3n\tau^2}{8 \times 3^k} \right) \quad \tau \in [0, 1].$$

Appendix D. Conversion of confidence regions from operator norm to trace norm

The final ingredient for the framework presented in this manuscript is a reliable way to transform operator-norm closeness of the (plain) least squares estimator \hat{L}_n into a statement about closeness of the PLS estimator $\hat{\rho}_n$ in trace distance. Recall that the optimization problem (A.2) admits an analytic solution [27]. Let $U \text{diag}(\vec{\lambda}) U^\dagger$ be an eigenvalue decomposition of \hat{L}_n . Then,

$$\hat{\rho}_n = U \text{diag} \left(\left[\vec{\lambda} - x_0 \vec{1} \right]^+ \right) U^\dagger, \tag{D.1}$$

where x_0 is chosen such that $\text{tr}(\hat{\rho}_n) = 1$ and $[\vec{y}]_i^+ = \max \{ [\vec{y}]_i, 0 \}$ denotes thresholding on non-negative components. This solution is *unique*, provided that $\text{tr}(\hat{L}_n) = 1$, which is the case for all the least squares estimators we consider.

The conversion from closeness in operator norm to closeness in trace norm will introduce a factor that is proportional to the *effective rank* of the density matrix ρ , rather than a full dimensional factor. For $r \in \mathbb{N}$, we define the best rank- r approximation ρ_r of a quantum state $\rho \in \mathbb{H}_d$ as the optimal feasible point of

$$\sigma_r(\rho) = \underset{\text{rank}(Z) \leq r}{\text{minimize}} \quad \|\rho - Z\|_1. \tag{D.2}$$

This problem can be solved analytically. Let $\rho = \sum_{i=1}^d \lambda_i |x_i\rangle\langle x_i|$ be an eigenvalue decomposition with eigenvalues arranged in non-increasing order. Then,

$$Z^\sharp = \sum_{i=1}^r \lambda_i |x_i\rangle\langle x_i|, \quad \text{and} \quad \sigma_r(\rho) = \sum_{i=r+1}^d \lambda_i = 1 - \text{tr}(\rho_r),$$

highlighting that the best rank- r approximation is simply a truncation onto the r largest contributions in the eigenvalue decomposition. This truncated description is accurate if the residual error $\sigma_r(\rho)$ is small. If this is the case it is reasonable to say that ρ is well approximated by a rank- r matrix Z^\sharp and has *effective rank* r .

Proposition 1. *Suppose that $\hat{L}_n \in \mathbb{H}_d$ obeys $\text{tr}(\hat{L}_n) = 1$ and $\|\hat{L}_n - \rho\|_\infty \leq \tau$ for some quantum state $\rho \in \mathbb{H}_d$ and $\tau \geq 0$. Then, for any $r \in \mathbb{N}$ the PLS estimator $\hat{\rho}_n$ obeys*

$$\|\hat{\rho}_n - \rho\|_1 \leq 4r\tau + 2 \min \{ \sigma_r(\rho), \sigma_r(\hat{\rho}_n) \},$$

where $\sigma_r(\rho)$ is defined in equation (D.2).

The statement readily follows from combining two auxiliary results. The first one states that the threshold value x_0 in the analytic solution of $\hat{\rho}_n$ must be small if \hat{L}_n is operator-norm close to a quantum state.

Lemma 1. *Instantiate the assumptions from proposition 1. Then, the threshold value in equation (D.1) obeys $x_0 \in [0, \tau]$.*

Proof. By assumption \hat{L}_n has unit trace. If it is in addition psd, $\hat{\rho}_n = \hat{L}_n$, because \hat{L}_n is already a quantum state and the projection is trivial ($x_0 = 0$). Otherwise, \hat{L}_n is indefinite and unit trace ensures that the positive part dominates. Hence, x_0 must be strictly positive to enforce $\text{tr}(\hat{\rho}_n) = 1$.

For the upper bound, let $P \in \mathbb{H}_d$ denote the orthogonal projection onto the range of $\hat{\rho}_n$. Then, $\hat{\rho}_n = P(\hat{L}_n - x_0\mathbb{I})P = P\hat{L}_nP - x_0P$, according to equation (D.1). In semidefinite order, this implies

$$\hat{\rho}_n = P(\hat{L}_n - \rho)P - x_0P + P\rho P \leq (\|\hat{L}_n - \rho\|_\infty - x_0)P + P\rho P \leq (\tau - x_0)P + P\rho P,$$

where the last line follows from the assumption $\|\hat{L}_n - \rho\|_\infty \leq \tau$. The trace preserves semidefinite order and we conclude

$$0 \leq \text{tr}(\hat{\rho}_n) - \text{tr}(P\rho P) \leq (\tau - x_0) \text{tr}(P)$$

which implies an upper bound of τ , because $\text{tr}(P) > 0$. □

The second technical lemma generalizes a result that is somewhat folklore in quantum information theory: the ‘effective rank’ of a difference of two quantum states is proportional to the minimal rank of the two density operators involved.

Lemma 2. Fix $r \in \mathbb{N}$ and let $\rho, \sigma \in \mathbb{H}_d$ be quantum states. Then,

$$\|\rho - \sigma\|_1 \leq 2r\|\rho - \sigma\|_\infty + 2 \min \{\sigma_r(\rho), \sigma_r(\sigma)\},$$

where the residual error $\sigma_r(\cdot)$ was defined in equation (D.2).

Proof. We can without loss of generality assume $\sigma_r(\rho) \leq \sigma_r(\sigma)$. Decompose ρ into $\rho_r + \rho_c$, where ρ_r is the best rank- r approximation (D.2) and $\rho_c = \rho - \rho_r$ denotes the ‘tail’. By construction, both ρ_r and ρ_c are positive semidefinite matrices that obey $\sigma_r(\rho) = \text{tr}(\rho_c) = 1 - \text{tr}(\rho_r)$. The triangle inequality then implies

$$\|\rho - \sigma\|_1 \leq \|\rho_r - \sigma\|_1 + \sigma_r(\rho),$$

because $\|\rho_c\|_1 = \sigma_r(\rho)$. Next, let $P_+, P_- \in \mathbb{H}_d$ be the projections onto the positive and non-positive ranges of $\rho_r - \sigma$. By construction, P_+ has rank at most $\text{rank}(\rho_r) = r$ and the trace norm equals

$$\|\rho_r - \sigma\|_1 = \text{tr}(P_+(\rho_r - \sigma)) - \text{tr}(P_-(\rho_r - \sigma)).$$

On the other hand,

$$\sigma_r(\rho) = \text{tr}(\sigma - \rho_r) = -\text{tr}(P_+(\rho_r - \sigma)) - \text{tr}(P_-(\rho_r - \sigma)),$$

because $P_+ + P_- = \mathbb{I}$. Combining both relations yields

$$\begin{aligned} \|\rho_r - \sigma\|_1 &= 2\text{tr}(P_+(\rho_r - \sigma)) + \sigma_r(\rho) \leq 2\text{tr}(P_+(\rho - \sigma)) + \sigma_r(\rho) \\ &\leq 2\|P_+\|_\infty\|\rho - \sigma\|_1 + \sigma_r(\rho), \end{aligned}$$

where we have used $\text{tr}(P_+\rho_c) \geq 0$ and Hoelder’s inequality. Finally, note that $\|P_+\|_1 = \text{rank}(P_+) = r$ by construction and the claim follows. \square

The main result of this section is a rather straightforward combination of these two technical statements.

Proof of proposition 1. Fix $r \in \mathbb{N}$ and use, lemma 2 to conclude

$$\|\hat{\rho}_n - \rho\|_1 \leq 2r\|\hat{\rho}_n - \rho\|_\infty + 2 \min \{\sigma_r(\rho), \sigma_r(\hat{\rho}_n)\}.$$

Next, note that according to (D.1), $\hat{\rho}_n$ may be viewed as the positive definite part of the matrix $\hat{L}_n - x_0\mathbb{I}$. Such a restriction to the positive part can never increase the operator norm distance to another positive semidefinite matrix. Hence,

$$\|\hat{\rho}_n - \rho\|_\infty \leq \|\hat{L}_n - \rho\|_\infty + |x_0|\|\mathbb{I}\|_\infty \leq 2\tau,$$

where the last inequality follows from lemma 1. \square

Appendix E. Proof of the main result

By now we have everything in place to provide a complete proof of the main result of this work.

Theorem 4. Let $\rho \in \mathbb{H}_d$ be a state. Suppose that we either perform n structured POVM measurements (set $g(d) = 2d$), n Pauli observable measurements (set $g(d) = d^2$), or n Pauli

basis measurements (set $g(d) = d^{1.6}$). Then, for any $r \in \mathbb{N}$ and $\epsilon \in [0, 1]$, the PLS estimator $\hat{\rho}_n$ (A.2) obeys

$$\Pr \left[\|\hat{\rho}_n - \rho\|_1 \geq \epsilon + 2 \min \{ \sigma_r(\rho), \sigma_r(\hat{\rho}_n) \} \right] \leq de^{-\frac{n\epsilon^2}{43g(d)r^2}},$$

where $\sigma_r(\rho), \sigma_r(\hat{\rho}_n)$ denote the residual error of approximating ρ and $\hat{\rho}_n$ by a rank- r matrix (D.2).

Note that theorem 1 in the main text is an immediate consequence of this more general result: simply set $r = \min \{ \text{rank}(\rho), \text{rank}(\hat{\rho}_n) \}$ which in turn ensures $\min \{ \sigma_r(\rho), \sigma_r(\hat{\rho}_n) \} = 0$.

However, unlike this specification, theorem 4 does feature an additional degree of freedom. The parameter $r \in \mathbb{N}$ allows for interpolating between small values (small sampling rate, but a potentially large reconstruction error) and large values (high sampling rate, but low reconstruction error). This tradeoff is particularly benign for quantum states that are approximately low-rank. Due to experimental imperfections, such states arise naturally in many experiments that aim at generating a pure quantum state. We illustrate this by means of the following caricature of a faulty state preparation protocol. Suppose that an apparatus either produces a target state $|\psi\rangle\langle\psi|$ perfectly, or fails completely, in the sense that it outputs a maximally mixed state. Then, the resulting state is

$$\rho = (1 - p)|\psi\rangle\langle\psi| + \frac{p}{d}\mathbb{I},$$

where $p \in [0, 1]$ denotes the probability of failure. This state has clearly full rank. Theorem 1 in the main text requires at least $n \geq 43 \frac{g(d)d^2}{\epsilon^2} \log(d/\delta)$ samples to estimate it up to trace-norm accuracy ϵ with high probability. In contrast, theorem 4 ensures that already $n \geq 43 \frac{g(d)}{\epsilon^2} \log(d/\delta)$ samples suffice to ensure that, with high probability, the PLS estimator obeys $\|\hat{\rho}_n - \rho\|_1 \leq \epsilon + 2p$. For sufficiently high success probabilities/low accuracy ($\epsilon \geq 2p/d$) this outperforms the original statement.

Proof of theorem 4. We illustrate the proof for structured POVMs—the other settings are completely analogous. Fix $\epsilon \in [0, 1]$, $r \in \mathbb{N}$ and set $\tau = \frac{\epsilon}{4r}$. Then, the main result of appendix C.1—equation (C.1)—ensures that the least squares estimator \hat{L}_n obeys

$$\|\hat{L}_n - \rho\|_\infty \leq \tau = \frac{\epsilon}{4r} \tag{E.1}$$

with probability of failure bounded by $de^{-\frac{\epsilon^2 n}{86dr^2}}$. Assuming that this condition is true, proposition 1 readily yields $\|\hat{\rho}_n - \rho\|_1 \leq \epsilon + 2 \min \{ \sigma_r(\rho), \sigma_r(\hat{\rho}_n) \}$. \square

Appendix F. Improved convergence guarantees for the uniform POVM

All the convergence results derived so far feature an additional $\log(d)$ -factor. This is a consequence of the matrix Bernstein inequality that proved instrumental in deriving these results. One can show that such an additional factor necessarily features in *all* matrix concentration inequalities that are based on exclusively first and second moments of the random matrices in question [57].

However, the following question remains: is this $\log(d)$ -factor an artifact of the proof technique, or is it an intrinsic feature?

In this section we rule out the second possibility: a different proof technique allows for avoiding this $\log(d)$ -factor, provided that the POVM is sufficiently symmetric and well-behaved. More precisely, we re-visit the uniform POVM $\{d|v\rangle\langle v|dv\}_{v\in\mathbb{S}^d}$ and exploit the fact that equation (B.1) completely characterizes *all* moments of the resulting outcome distribution. This opens the door for applying very strong proof techniques from large dimensional probability theory that found wide-spread applications in a variety of subjects, including compressed sensing [58] and, more recently, quantum information theory [59]. We believe that this technique may be of independent interest and find it therefore worthwhile to present it in a self-contained fashion. Roughly speaking, it is based on the following steps:

- (a) *Reformulation*: the operator norm of a random Hermitian matrix $A \in \mathbb{H}_d$ admits a variational definition:

$$\|A\|_\infty = \max_{y \in \mathbb{S}^d} |\langle y|A|y \rangle|. \tag{F.1}$$

- (b) *Discretization*: replace the maximization over the entire complex unit sphere \mathbb{S}^d by a maximization over a finite point set \mathcal{N} that covers \mathbb{S}^d to sufficiently high accuracy (*covering net*).
- (c) *Concentration*: fix $y \in \mathcal{N}$ and show that the scalar random variable $s_y = \langle y|A|y \rangle$ concentrates sharply around its expectation value.
- (d) *Union bound*: apply a union bound over all $|\mathcal{N}|$ random variables s_y to obtain an upper bound on the operator norm $\|A\|_\infty$.

Ideally the tail bound from (d) is sharp enough to ‘counter-balance’ the $|\mathcal{N}|$ -pre-factor that results from the union bound in step (iv). Should this be not the case, more sophisticated methods, like generic chaining [60], may still allow for drawing non-trivial conclusions. Fortunately, for the task at hand, this turns out to not be necessary and the rather naive strategy sketched above suffices to achieve a result that is (provably) optimal up to a constant factors:

Theorem 5. *Suppose that we perform n independent uniform POVM measurements on a quantum state $\rho \in \mathbb{H}_d$. Then, the associated least squares estimator \hat{L}_n obeys*

$$\Pr \left[\|\hat{L}_n - \rho\|_\infty \geq \tau \right] \leq 2 \exp \left(c_1 d - c_2 n \tau^2 \right),$$

In particular, $n \geq C \frac{d}{\tau^2} \log(1/\delta)$ suffices to ensure $\|\hat{L}_n - \rho\|_\infty \leq \tau$ with probability at least $1 - \delta$. Here $c_1, c_2, C > 0$ denote constants of sufficient size.

No effort has been made to optimize the constants. The proof presented here yields $c_1 = 2 \log(3)$ and $c_2 = \frac{1}{480}$ which could be further improved by a more careful analysis. Importantly, the second part of this statement can be combined with proposition 1 to readily deduce the last technical result of the main text:

Corollary 2. *(Re-statement of theorem 5). For any rank- r state ρ , a number of $n \geq C \frac{r^2 d}{\epsilon^2} \log(1/\delta)$ uniform POVM measurements suffice to ensure $\|\hat{\rho}_{(n)}^\sharp - \rho\|_1 \leq \epsilon$ with probability at least $1 - \delta$.*

Not only does this statement reproduce the best known sampling rates for tomography with independent measurements [22], it also exactly matches lower bounds on the minimal sample complexity associated with *any* tomographic procedure that may apply in this setting [20, table 1].

The remainder of this section is dedicated to proving theorem 5. For the sake of accessibility, we will divide this proof into three subsections that contain the steps summarized above.

F.1. Step I: reformulation and discretization

Suppose that we perform n uniform POVM measurements on a fixed quantum state $\rho \in \mathbb{H}_d$. Then, the least squares estimator is equivalent to a sum of i.i.d. random matrices:

$$\hat{L}_n = \frac{1}{n} \sum_{i=1}^n X_i.$$

Each X_i is an independent copy of the random matrix X that assumes the value $(d + 1)|v\rangle\langle v| - \mathbb{I}$ with probability $d\langle v|\rho|v\rangle dv$ and v may range over the entire complex unit sphere. Unbiasedness of this estimator in turn implies

$$\|\hat{L}_n - \rho\|_\infty = \left\| \frac{1}{n} \sum_{i=1}^n (X_i - \mathbb{E}[X_i]) \right\|_\infty = \max_{y \in \mathbb{S}^d} \left| \langle y | \frac{1}{n} \sum_{i=1}^n (X_i - \mathbb{E}[X_i]) | y \rangle \right|.$$

Next, we employ a result that is somewhat folklore in random matrix theory, see e.g. [61, lemma 5.3]. It states that the maximum over the entire unit sphere may be replaced by a maximum over certain finite point sets, called *covering nets*: a covering-net of \mathbb{S}^d with fineness $\theta > 0$ is a finite set of unit vectors $\{z_j\}_{j=1}^N \subseteq \mathbb{S}^d$ that covers the entire (complex) unit sphere in the sense that every $y \in \mathbb{S}^d$ is at least θ -close to a point in the net.

Lemma 3. *Let $\mathcal{N}_\theta = \{z_j\}_{j=1}^N$ be a covering net of \mathbb{S}^d with fineness θ . Then, for any matrix $A \in \mathbb{H}_d$:*

$$\max_{j \in [N]} |\langle z_j | A | z_j \rangle| \leq \|A\|_\infty \leq \frac{1}{1 - 2\theta} \max_{j \in [N]} |\langle z_j | A | z_j \rangle|.$$

This result highlights that already a rather coarse net suffices to get reasonable approximations to the operator norm. Here, we choose $\theta = \frac{1}{4}$ which, while certainly not optimal, simplifies exposition. In particular,

$$\|\hat{L}_n - \rho\|_\infty \leq 2 \max_{j \in [N]} \left| \frac{1}{n} \sum_{i=1}^n \langle z_j | X_i - \mathbb{E}[X_i] | z_j \rangle \right|, \tag{F.2}$$

where the maximization is over a covering net of fineness $\theta = \frac{1}{4}$.

F.2. Step II: concentration

Note that the right hand side of equation (F.2) corresponds to a maximum over N different random variables—each of them labeled by a unit vector z_j in the net. Let $z \in \mathbb{S}^d$ be such a vector. Then, the associated random variable itself corresponds to an empirical average of n i.i.d. variables:

$$s_z = \langle z | X - \mathbb{E}[X] | z \rangle.$$

Clearly, s_z obeys $\mathbb{E}[s_z] = 0$ and, more importantly, has sub-exponential moment growth. While this follows directly from the fact that s_z is bounded, the following result highlights that this tail-behavior is actually independent of the ambient dimension.

Lemma 4. *Fix $z \in \mathbb{S}^d$. Then for any integer $p \geq 2$, the random variable s_z obeys*

$$\mathbb{E}[|s_z|^p] \leq 27 \times 6^{p-2} p!$$

We divert the proof of this statement to the end of this section and content ourselves with emphasizing that the closed form expression of the frame operator (B.1) is essential for bounding all moments simultaneously. More relevant to the task at hand is that such a moment behavior ensures that the tails of the distribution of s_z follow an exponential decay: $\Pr [|s_z| \geq t] \leq e^{-ct}$, where c is a constant independent of the dimension d . Strong classical concentration inequalities apply for sums of i.i.d. random variables that exhibit such sub-exponential behavior. We choose to apply a rather general version of the classical Bernstein inequality, see e.g. [58, theorem 7.30].

Theorem 6. *Let $s_1, \dots, s_n \in \mathbb{R}$ i.i.d. copies of a mean-zero random variable s that obeys $\mathbb{E} [|s|^p] \leq p!R^{p-2}\sigma^2/2$ for all integers $p \geq 2$, where $R, \sigma^2 > 0$ are constants. Then, for all $t > 0$,*

$$\Pr \left[\left| \sum_{i=1}^n s_i \right| \geq t \right] \leq 2 \exp \left(-\frac{t^2/2}{n\sigma^2 + Rt} \right).$$

Lemma 4 ensures that the random variable s_z meets this requirement with $\sigma^2 = 54$ and $R = 6$. Hence, the following corollary is an immediate consequence of theorem 6.

Corollary 3. *Fix $z \in \mathbb{S}^d$. Then, for any $t \in [0, 1]$*

$$\Pr \left[\left| \frac{1}{n} \sum_{i=1}^n \langle z | X_i - \mathbb{E} [X_i] | z \rangle \right| \geq t \right] \leq 2e^{-\frac{nt^2}{120}}.$$

F.3. Step III: union bound

Recall that equation (F.2) upper-bounds $\|\hat{L}_n - \rho\|_\infty$ by a maximum over finitely many random variables, each of which is controlled by the strong exponential tail inequality from corollary 3. To exploit this, we fix $\tau \in [0, 1]$ and apply a union bound (also known as Boole’s inequality) over all these different random variables to obtain

$$\begin{aligned} \Pr [\|\hat{L}_n - \rho\|_\infty \geq \tau] &\leq \Pr \left[\max_{j \in [M]} \left| \frac{1}{n} \sum_{i=1}^n \langle z_j | X_i - \mathbb{E} [X_i] | z_j \rangle \right| \geq \frac{\tau}{2} \right] \\ &\leq N \max_{j \in [M]} \Pr \left[\left| \frac{1}{n} \sum_{i=1}^n \langle z_j | X_i - \mathbb{E} [X_i] | z_j \rangle \right| \geq \frac{\tau}{2} \right] \leq 2Ne^{-\frac{n\tau^2}{480}}, \end{aligned}$$

where the last line is due to corollary 3 Here, $N = |\mathcal{N}_{\frac{1}{4}}|$ denotes the cardinality of a covering net for the complex unit sphere \mathbb{S}^d with fineness $\theta = \frac{1}{4}$. The complex unit sphere admits an isometric embedding into the real-valued unit sphere in $2d$ -dimensions: map real- and imaginary parts of each complex vector component onto two distinct real parameters. This map preserves Euclidean lengths and, by extension, also the geometry of the unit sphere. Volumetric upper bounds on the cardinality of covering nets for the $2d$ -dimensional real-valued unit sphere are widely known, see e.g. [58, proposition C.3] and [61, lemma 5.2]: $|\mathcal{N}_\theta| \leq \left(1 + \frac{2}{\theta}\right)^{2d}$. Since a fineness of $\theta = \frac{1}{4}$ suffices for our purpose, we can conclude $N \leq 3^{2d}$ and consequently,

$$\Pr [\|\hat{L}_n - \rho\|_\infty \geq \tau] \leq 2 \times 3^{2d} e^{-\frac{n\tau^2}{480}} = 2e^{2 \log(3)d - \frac{n\tau^2}{480}}$$

This concludes the proof of theorem 5.

F.4. Proof of lemma 4

Recall that, by assumption, the random matrix X assumes the value $X = (d + 1)|v\rangle\langle v| - \mathbb{I}$ with probability $\langle v|\rho|v\rangle dv$, where v may range over the entire complex unit sphere \mathbb{S}^d . Moreover, $\mathbb{E}[X] = \rho$. For fixed $z \in \mathbb{S}^d$, we may therefore write

$$s_z = \langle z|X - \mathbb{E}[X]|z\rangle = (d + 1)\langle v|B|v\rangle,$$

where $B = |z\rangle\langle z| - \frac{1+\langle z|\rho|z\rangle}{d+1}\mathbb{I} \in \mathbb{H}_d$ has bounded trace norm

$$\|B\|_1 \leq 1 + (1 + \langle z|\rho|z\rangle) \leq 3. \tag{F.3}$$

Next, recall a basic identity from matrix analysis that states

$$|\langle v|B|v\rangle| = |\text{tr}(|v\rangle\langle v|B)| \leq \text{tr}(|v\rangle\langle v||B|),$$

where $|B| = \sqrt{B^2}$ denotes the absolute value of the matrix B . Also, the Schatten- p norms of matrices and their absolute values coincides, in particular $\|B\|_1 = \text{tr}(|B|) = \||B|\|_1$. We can use this trick to absorb the absolute value in the moment computation. More precisely, fix an integer $p \geq 2$ and note that $\mathbb{E}[|s_z|^p]$ obeys

$$\mathbb{E}[|(d + 1)\langle v|B|v\rangle|^p] \leq (d + 1)^p \mathbb{E}[\text{tr}(|v\rangle\langle v||B|)^p].$$

We can now include the distribution of the random matrices X , and (by extension) $|v\rangle\langle v|$, to compute

$$\begin{aligned} \mathbb{E}[|s_z|^p] &\leq (d + 1)^p \mathbb{E}[\text{tr}(|v\rangle\langle v||B|)^p] = d(d + 1)^p \int_{\mathbb{S}^d} \langle v|\rho|v\rangle \text{tr}(|v\rangle\langle v||B|)^p dv \\ &= d(d + 1)^p \text{tr}\left(\int_{\mathbb{S}^d} (|v\rangle\langle v|)^{\otimes(p+1)} \rho \otimes |B|^{\otimes p}\right) \\ &= d(d + 1)^p \binom{d + p}{p + 1}^{-1} \text{tr}\left(P_{\text{Sym}^{(p+1)}} \rho \otimes |B|^{\otimes p}\right), \end{aligned}$$

where the last equation is due to equation (B.1). Next, we note that Hoelder’s inequality implies

$$\text{tr}\left(P_{\text{Sym}^{(p+1)}} \rho \otimes |B|^{\otimes p}\right) \leq \|P_{\text{Sym}^{(p+1)}}\|_\infty \|\rho\|_1 \| |B|^{\otimes p} \|_1 \leq 3^p,$$

because $P_{\text{Sym}^{(p+1)}}$ is an orthogonal projector, ρ is a quantum state and B is bounded in trace norm (F.3). For the remaining pre-factor we use the crude bound

$$d(d + 1)^p \binom{d + p}{p + 1}^{-1} \leq (p + 1)! \leq 3 \times 2^{p-2} p!$$

to establish the statement.

ORCID iDs

M Guță  <https://orcid.org/0000-0002-9017-2222>

References

- [1] Hradil Z 1997 *Phys. Rev. A* **55** R1561–4
- [2] Řeháček J, Hradil Z and Ježek M 2001 *Phys. Rev. A* **63** 040303
- [3] Häffner H, Hänsel W, Roos C F, Benhelm J, Chek-al kar D, Chwalla M, Körber T, Rapol U D, Riebe M, Schmidt P O, Becher C, Gühne O, Dür W and Blatt R 2005 *Nature* **438** 643
- [4] Andrews D W K 2000 *Econometrica* **68** 399–405
- [5] Blume-Kohout R 2010 *New J. Phys.* **12** 043034
- [6] Granade C, Combes J and Corrie D G 2016 *New J. Phys.* **18** 033024
- [7] Granade C, Ferrie C and Flammia S T 2017 *New J. Phys.* **19** 113017
- [8] Christandl M and Renner R 2012 *Phys. Rev. Lett.* **109** 120403
- [9] Blume-Kohout R 2012 arXiv:1202.5270
- [10] Faist P and Renner R 2016 *Phys. Rev. Lett.* **117** 010404
- [11] Gross D, Liu Y K, Flammia S T, Becker S and Eisert J 2010 *Phys. Rev. Lett.* **105** 150401
- [12] Liu Y K 2011 *Universal Low-Rank Matrix Recovery from Pauli Measurements (Adv. Neural Inf. Process. Syst. vol 24)* ed J Shawe-Taylor et al (Red Hook, NJ: Curran Associates, Inc) pp 1638–46
- [13] Flammia S T, Gross D, Liu Y K and Eisert J 2012 *New J. Phys.* **14** 095022
- [14] Riofrío C, Gross D, Flammia S, Monz T, Nigg D, Blatt R and Eisert J 2017 *Nat. Commun.* **8** 15305
- [15] Koltchinskii V 2011 *Ann. Stat.* **39** 2936–73
- [16] Koltchinskii V 2013 *IMS Collections* **9** 213–26
- [17] Koltchinskii V and Xia D 2015 *J. Mach. Learn. Res.* **16** 1757–92 (<http://www.jmlr.org/papers/volume16/koltchinskii15a/koltchinskii15a.pdf>)
- [18] Xia D and Koltchinskii V 2016 *Electron. J. Stat.* **10** 1935–7524
- [19] D’Ariano G M and Perinotti P 2007 *Phys. Rev. Lett.* **98** 020403
- [20] Haah J, Harrow A W, Ji Z, Wu X and Yu N 2017 *IEEE Trans. Inf. Theory* **63** 5628–41
- [21] Sugiyama T, Turner P S and Murao M 2013 *Phys. Rev. Lett.* **111** 160406
- [22] Kueng R, Rauhut H and Terstiege U 2017 *Appl. Comput. Harmon. Anal.* **42** 88–116
- [23] Kahn J and Guță M 2009 *Commun. Math. Phys.* **289** 597–652
- [24] O’Donnell R and Wright J 2016 Efficient quantum tomography *Proc. 48th Annual ACM Symp. on Theory of Computing STOC ’16* (New York: ACM) pp 899–912
- [25] Acharya A, Kypraios T and Guță M 2019 arXiv:1901.07991
- [26] Kaznady M S and James D F V 2009 *Phys. Rev. A* **79** 022109
- [27] Smolin J A, Gambetta J M and Smith G 2012 *Phys. Rev. Lett.* **108** 070502
- [28] Alquier P, Butucea C, Hebiri M, Meziani K and Morimae T 2013 *Phys. Rev. A* **88** 032113
- [29] Butucea C, Guță M and Kypraios T 2015 *New J. Phys.* **17** 113050
- [30] Kueng R, Zhu H and Gross D 2016 arXiv:1610.08070
- [31] Tropp J A 2012 *Found. Comput. Math.* **12** 389–434
- [32] Renes J M, Blume-Kohout R, Scott A J and Caves C M 2004 *J. Math. Phys.* **45** 2171–80
- [33] Klappenecker A and Rötteler M 2005 Mutually unbiased bases are complex projective 2-designs *Int. Symp. on Information Theory, 2005 ISIT 2005 Proc.* pp 1740–4
- [34] Dankert C, Cleve R, Emerson J and Livine E 2009 *Phys. Rev. A* **80** 012304
- [35] Gross D, Audenaert K and Eisert J 2007 *J. Math. Phys.* **48** 052104
- [36] Halko N, Martinsson P and Tropp J 2011 *SIAM Rev.* **53** 217–88
- [37] Kalev A, Kosut R L and Deutsch I H 2015 *NPJ Quantum Inf.* **1** 15018
- [38] Kabanava M, Kueng R, Rauhut H and Terstiege U 2016 *Inf. Inference* **5** 405–41
- [39] Wootters W K and Fields B D 1989 *Ann. Phys.* **191** 363–81
- [40] Bandyopadhyay S, Boykin P O, Roychowdhury V and Vatan F 2002 *Algorithmica* **34** 512–28
- [41] Klappenecker A and Rötteler M 2004 *Constructions of Mutually Unbiased Bases Finite Fields and Applications* ed G L Mullen, A Poli and H Stichtenoth (Berlin: Springer) pp 137–44
- [42] Scholten T L and Blume-Kohout R 2018 *New J. Phys.* **20** 023050
- [43] Scott A J 2006 *J. Phys. A* **39** 13507
- [44] Gross D, Kraemer F and Kueng R 2015 *J. Fourier Anal. Appl.* **21** 229–66
- [45] Schwinger J 1960 *Proc. Natl. Acad. Sci. USA* **46** 570–9
- [46] Nielsen M A and Chuang I L 2011 *Quantum Computation and Quantum Information: 10th Anniversary Edition* 10th edn (New York: Cambridge University Press)
- [47] Kueng R and Gross D 2015 arXiv:1510.02767
- [48] Zhu H 2017 *Phys. Rev. A* **96** 062336

- [49] Webb Z 2015 arXiv:[1510.02769](https://arxiv.org/abs/1510.02769)
- [50] Tomczak-Jaegermann N 1974 *Stud. Math.* **50** 163–82
- [51] Lust-Piquard F 1986 *CR Acad. Sci. Paris* **303** 289–92
- [52] Pisier G and Xu Q 1997 *Commun. Math. Phys.* **189** 667–98
- [53] Rudelson M 1999 *J. Funct. Anal.* **164** 60–72
- [54] Ahlswede R and Winter A 2002 *IEEE Trans. Inf. Theor.* **48** 569–79
- [55] Gross D 2011 *IEEE Trans. Inf. Theor.* **57** 1548–66
- [56] Oliveira R I et al 2010 *Electron. Commun. Probab.* **15** 26
- [57] Tropp J A 2018 *Appl. Comput. Harmon. Anal.* **44** 700–36
- [58] Foucart S and Rauhut H 2013 *A Mathematical Introduction To Compressive Sensing* (Basel: Birkhäuser)
- [59] Lancien C and Winter A 2017 arXiv:[1711.00697](https://arxiv.org/abs/1711.00697)
- [60] Talagrand M 2006 *The Generic Chaining: Upper and Lower Bounds of Stochastic Processes* (Berlin: Springer)
- [61] Vershynin R 2012 Introduction to the non-asymptotic analysis of random matrices *Compressed Sensing: Theory and Practice* ed Y C Eldar and G Kutyniok (Cambridge: Cambridge University Press) pp 210–68