

PAPER • OPEN ACCESS

Near-optimal quantum tomography: estimators and bounds

To cite this article: Richard Kueng and Christopher Ferrie 2015 *New J. Phys.* **17** 123013

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

Related content

- [Spectral thresholding quantum tomography for low rank states](#)
Cristina Butucea, Mdlin Gu and Theodore Kyraios
- [Error regions in quantum state tomography: computational complexity caused by geometry of quantum states](#)
Daniel Suess, ukasz Rudnicki, Thiago O maciel et al.
- [Practical Bayesian tomography](#)
Christopher Granade, Joshua Combes and D G Cory

Recent citations

- [Pseudo-Bayesian quantum tomography with rank-adaptation](#)
The Tien Mai and Pierre Alquier
- [Practical and Reliable Error Bars in Quantum Tomography](#)
Philippe Faist and Renato Renner



PAPER

Near-optimal quantum tomography: estimators and bounds

OPEN ACCESS

RECEIVED

25 August 2015

ACCEPTED FOR PUBLICATION

4 November 2015

PUBLISHED

10 December 2015

Content from this work
may be used under the
terms of the [Creative
Commons Attribution 3.0
licence](#).

Any further distribution of
this work must maintain
attribution to the
author(s) and the title of
the work, journal citation
and DOI.

Richard Kueng^{1,2,3,5} and Christopher Ferrie^{1,4}¹ Centre for Engineered Quantum Systems, School of Physics, The University of Sydney, Sydney, NSW, Australia² Institute for Theoretical Physics, University of Cologne, Germany³ Institute for Physics, University of Freiburg, Germany⁴ Center for Quantum Information and Control, University of New Mexico, Albuquerque, NM, 87131-0001, USA⁵ Author to whom any correspondence should be addressed.E-mail: richard@kueng.at**Keywords:** quantum tomography, quantum information, average fidelity, Bayesian tomography**Abstract**

We give bounds on the average fidelity achievable by any quantum state estimator, which is arguably the most prominently used figure of merit in quantum state tomography. Moreover, these bounds can be computed online—that is, while the experiment is running. We show numerically that these bounds are quite tight for relevant distributions of density matrices. We also show that the Bayesian mean estimator is ideal in the sense of performing close to the bound without requiring optimization. Our results hold for all finite dimensional quantum systems.

1. Introduction

Inferring a quantum mechanical description of a physical system is equivalent to assigning it a quantum state—a process referred to as *tomography*. Tomography is now a routine task for designing, testing and tuning qubits in the quest of building quantum information processing devices [1]. In determining how ‘good’ one is performing this task, a figure of merit must be reported. By far the most commonly used figure of merit for quantum states is *fidelity* [2, 3]. Nowadays, fidelity is used to compare quantum states and processes in a wide variety of tasks, from quantum chaos to quantum control to the continuous monitoring of quantum systems [4–10]. The only known case where the optimal estimator for fidelity is known is a single qubit [6]. We generalize this result to find upper and lower bounds on the average fidelity of any estimator for any dimension.

For d -dimensional state space,

$$\mathcal{S} := \left\{ \sigma \in L(\mathbb{C}^d) : \sigma \geq 0, \text{Tr}(\sigma) = 1 \right\}, \quad (1)$$

the *fidelity* between two states $\rho, \sigma \in \mathcal{S}$ is defined to be [2, 3],

$$F(\rho, \sigma) := \|\sqrt{\rho}\sqrt{\sigma}\|_1^2 = \left[\text{Tr} \sqrt{\sqrt{\rho}\sigma\sqrt{\rho}} \right]^2. \quad (2)$$

Define the *average fidelity* with respect to some measure $d\rho$ as $\mathbb{E}_\rho[F(\rho, \sigma)]$ ⁶. We want the average of this to be as large as possible. Thus, the problem can be succinctly stated as follows:

$$\begin{aligned} & \text{maximize } \mathbb{E}_\rho[F(\rho, \sigma)] \\ & \text{subject to } \text{Tr}(\sigma) = 1, \\ & \sigma \geq 0. \end{aligned} \quad (3)$$

In the context of tomography, we think of ρ as the ‘true state’ and σ as the estimated state. An *estimator* is a function from the space of data to quantum states σ : $\text{data} \mapsto \sigma(\text{data}) \in \mathcal{S}$, where data are the results of a sequence of quantum measurements. Since both the true state and data are unknown, we take the expected value with respect to the joint distribution of (ρ, data) to obtain the average fidelity:

⁶ Expectation values will always be denoted with a subscript which specifies the implicit distribution of variables being averaged over.

$$f(\sigma) = \mathbb{E}_{\rho, \text{data}}[F(\rho, \sigma(\text{data}))]. \quad (4)$$

We want this to be as large as possible. The estimator which maximizes this quantity is equivalent to the estimator maximizing the following posterior average fidelity for every data set:

$$f(\sigma | \text{data}) = \mathbb{E}_{\rho | \text{data}}[F(\rho, \sigma(\text{data}))]. \quad (5)$$

An estimator which maximizes this is called a *Bayes estimator*⁷. Bayes estimators are useful both to understand Bayesian optimality and to provide upper bounds for the worst case performance.

Now here is the subtle and important point: the measurements performed, the data themselves and the distribution from which they were generated are not important once the posterior distribution has been calculated. If we know the solution for every measure $d\rho$, then we know the solution for the posterior measure $d\rho | \text{data}$. For brevity, then, we will drop this conditional information from now on and the problem reduces again to (3).

2. Summary of results

In this work, we provide absolute benchmarks for the average fidelity performance of any tomographic estimation strategy by way of upper and lower bounds. This is important because, in the field of quantum tomography, a common theme is to compare estimators. Up to date many options are available: linear inversion [1], maximum likelihood [12], Bayesian mean [13], hedged maximum likelihood [14], and compressed sensing [15, 16]—to name a few. Often estimators are compared by simulating measurements on ensembles of states drawn according to some measure and averaging the fidelity. This can only provide conclusions about the relative performance of estimators. Thus, our bounds can be used to benchmark the fidelity performance of other candidate estimators.

We complement our theoretical findings with numerical experiments. These demonstrate the relative tightness of our bounds and, in particular, reveal that the Bayesian mean estimator is an excellent choice—owing to its near-optimal performance and ease of implementation. Importantly, both the mean of the distribution and our bounds can be computed *online*—that is, the estimator and its performance can be computed while data is being taken. In the context of Bayesian quantum information theory [13], our findings lend credence to the standard approach of using the mean of the posterior distribution as an estimator is a near-optimal one.

We note that this problem has been solved for the case of a single qubit ($d = 2$). Bagan *et al* [6] have given the optimal estimator (and measurement!) for any isotropic prior measure. Unfortunately, by making heavy use of the Bloch representation of a qubit, the methods do not generalize. Whereas, our bound holds for all distributions of states in any dimension and coincides with the results of [6] for the case of a single qubit.

2.1. Ensembles of pure states

We first present the analytically soluble case of measures supported only on pure states. Such a case is common in theoretical studies which average the performance of their protocols over the popular choice of the unique Haar invariant measure on pure states. The solution is organized into the following theorem:

Theorem 1. *Choose an arbitrary dimension d and assume that the integration measure $d\rho$ is supported only on pure states. Then, the state which solves the optimization problem (3) is the eigenvector of $\mathbb{E}_{\rho}[\rho]$ with maximal eigenvalue. It achieves a maximal fidelity of $\|\mathbb{E}_{\rho}[\rho]\|_{\infty}$.*

The proof is a simple exercise in linear programming. When ρ is a pure state, the fidelity simplifies to $F(\rho, \sigma) = \text{Tr}(\rho\sigma)$. Linearity allows us to bring the expectation inside the trace so that the problem becomes

$$\begin{aligned} & \text{maximize } \text{Tr}(\mathbb{E}_{\rho}[\rho]\sigma) \\ & \text{subject to } \text{Tr}(\sigma) = 1, \\ & \sigma \geq 0. \end{aligned} \quad (6)$$

The solution can be found in many textbooks covering linear programming—e.g. [17]. This solution also coincides with the one noted for a distribution supported on two states in [18].

2.2. General measures on mixed states

For measures with support on mixed states, the situation is markedly different. Our main technical contribution are new upper bounds for this case. We obtain them by replacing the fidelity function—which is notoriously

⁷ The terminology and objective functions used here can be seen as standard generalizations of those familiar in decision theory. See, e.g., [11].

difficult to grasp—in the main optimization problem (3) by quantities that are easier to handle in full generality. One rather straightforward approach to do so is to relate the fidelity function $f(\rho, \sigma)$ between arbitrary states $\rho, \sigma \in \mathcal{S}$ to corresponding Schatten- p -norm distances

$$\|\rho - \sigma\|_p = \left(\text{Tr}(|\rho - \sigma|^p) \right)^{1/p},$$

with $1 \leq p \leq \infty$ and $|X| = \sqrt{X^*X}$ for any $X \in L(\mathbb{C}^d)$. This can be done by employing the well-known and often used Fuchs–van de Graaf inequalities [19]

$$1 - \sqrt{F(\rho, \sigma)} \leq \frac{1}{2} \|\rho - \sigma\|_1 \leq \sqrt{1 - F(\rho, \sigma)} \quad \forall \rho, \sigma \in \mathcal{S}.$$

This inequality together with the hierarchy of Schatten- p -norms assures

$$F(\rho, \sigma) \leq 1 - \frac{1}{4} \|\rho - \sigma\|_1^2 \leq 1 - \frac{1}{4} \|\rho - \sigma\|_2^2, \tag{7}$$

for any two quantum states $\rho, \sigma \in \mathcal{S}$. Replacing the objective function in the central optimization problem (3) by such an upper bound results in a different optimization which admits a general analytic solution. Clearly, such a relaxed optimum bounds the original figure of merit from above and allows us to establish our second main result.

Theorem 2. *For any finite dimension d and any distribution $d\rho$, the maximal average fidelity achieved by any estimator $\sigma \in \mathcal{S}$ obeys*

$$\max_{\sigma \in \mathcal{S}} \mathbb{E}_\rho[F(\rho, \sigma)] \leq 1 - \frac{1}{4} \text{Tr} \left(\mathbb{E}_\rho[\rho^2] - \mathbb{E}_\rho[\rho]^2 \right). \tag{8}$$

Note that the expression on the right-hand side of (8) can be interpreted as a non-commutative generalization of the variance of a probability distribution. Having already outlined the main ideas necessary to establish such a result, we refer to section 4.2 for a complete proof.

Another way of establishing upper bounds on the average fidelity involves the concept of *super-fidelity*, which provides the following upper bound on the fidelity [20]:

$$F(\rho, \sigma) \leq \text{Tr}(\rho\sigma) + \sqrt{1 - \text{Tr}(\rho^2)} \sqrt{1 - \text{Tr}(\sigma^2)}. \tag{9}$$

Although more involved, we shall see that such an approach yields strictly better bounds than the ones presented in Theorem 2. For brevity, we define $\hat{\rho} := \mathbb{E}_\rho[\rho]$ and $p_\rho := \mathbb{E}_\rho[\sqrt{1 - \text{Tr}(\rho^2)}]$, such that inequality (9) assures

$$\max_{\sigma \in \mathcal{S}} \mathbb{E}_\rho[F(\rho, \sigma)] \leq \max_{\sigma \in \mathcal{S}} \left(\text{Tr}(\hat{\rho}\sigma) + p_\rho \sqrt{1 - \text{Tr}(\sigma^2)} \right), \tag{10}$$

for any distribution $d\rho$. Although more tractable than the original problem, the optimization on the right hand side still requires solving a non-commutative maximization over all quantum states $\sigma \in \mathcal{S}$. However, applying a corollary of the famous Birkhoff–von Neumann theorem—see e.g. [21, theorem 8.7.6]—allows for restricting this optimization to density operators σ that commute with the distribution’s mean $\hat{\rho}$ —see lemma 1 below. If $\hat{\tau}_1, \dots, \hat{\tau}_d$ denote the eigenvalues of $\hat{\rho}$ such a restriction assures that solving the right-hand side of (10) is equivalent to

$$\begin{aligned} & \text{maximize} && \sum_{i=1}^d \hat{\tau}_i s_i + p_\rho \sqrt{1 - \sum_{i=1}^d s_i^2}, \\ & \text{subject to} && \sum_{i=1}^d s_i = 1, \\ & && s_i \geq 0, \quad 1 \leq i \leq d, \end{aligned} \tag{11}$$

which is a commutative convex optimization problem. We refer to lemma 1 below for a detailed proof of this assertion. Note that, if the measure $d\rho$ is supported exclusively on pure states, p_ρ vanishes and (11) reduces to theorem 1 which is tight.

In order to obtain analytical bounds for mixed states, we further relax (11) by replacing the non-negativity constraints ($s_i \geq 0$) by the weaker demand that the optimization vector $(s_1, \dots, s_d)^T \in \mathbb{R}^d$ is contained in the Euclidean unit ball—i.e. $\sum_{i=1}^d s_i^2 \leq 1$. As we shall show in section 5, such a simplification is the tightest possible ellipsoidal relaxation of (11) and allows us to apply the method of Lagrangian multipliers in a straightforward fashion. Doing so results in the main theoretical statement of this paper.

Theorem 3. For any finite dimension d and any distribution $d\rho$ over states, the fidelity achieved by any estimator $\sigma \in \mathcal{S}$ is bounded from above by

$$\mathbb{E}_\rho[F(\rho, \sigma)] \leq \frac{1}{d} \left(1 + \sqrt{d-1} \sqrt{d \left(\mathbb{E}_\rho \left[\sqrt{1 - \text{Tr}(\rho^2)} \right]^2 + \text{Tr}(\mathbb{E}_\rho[\rho]^2) \right) - 1} \right). \quad (12)$$

The matrix achieving this optimum corresponds to

$$\sigma^\sharp = \frac{1}{d} \mathbb{1} + \sqrt{\frac{d-1}{d(p_\rho^2 + \text{Tr}(\hat{\rho}^2)) - 1}} \left(\hat{\rho} - \frac{1}{d} \mathbb{1} \right), \quad (13)$$

where $\mathbb{1} \in L(\mathbb{C}^d)$ denotes the identity matrix.

Again, we content ourselves here with outlining the proof architecture necessary to establish such a result and refer to section 4 for a detailed analysis.

Note that since we relaxed the maximization constraints, σ^\sharp in general fails to be positive-semidefinite and is thus not a valid density operator, though we do not use it as such. In particular, the bound is not tight when $d\rho$ is supported only on pure states—as might be evident from the possibility of non-positive states arising from the $(\hat{\rho} - \frac{1}{d}\mathbb{1})$ term in (13). On the other hand, the distribution is known and thus in the case of a distribution supported only on pure states, one should consult the exact solution in theorem 1.

Conversely, if σ^\sharp happens to be a state, it also solves the optimization (11) and the analytical bound (12) exactly reproduces an *a priori* tighter one. In all of our numerical experiments, some of which are presented below, this was indeed the case.

It is also worthwhile to point out that super-fidelity—the bound in (9)—and the actual fidelity coincide for one qubit, i.e. for $d = 2$ [20]. Also replacing positive semidefiniteness by bounded purity yields the same feasible set for that particular case. Consequently the bound (12) reproduces one of the main results in [6]:

Corollary 1. In the single-qubit case (i.e. $d = 2$) the bound (12) exactly reproduces the maximum average fidelity in [6, equation (2.9)] and σ^\sharp is the optimal estimator.

Finally, we want to emphasize that establishing bounds on the average fidelity by using the super-fidelity instead of the Fuchs–van de Graaf inequalities leads to strictly better results:

Corollary 2. Let the dimension d and the distribution $d\rho$ over states be arbitrary. Then, the bound presented in section 2 (Fuchs van–de Graaf inequality) is either trivial—i.e. equal to one—or it strictly majorizes the one presented theorem 3 (super-fidelity).

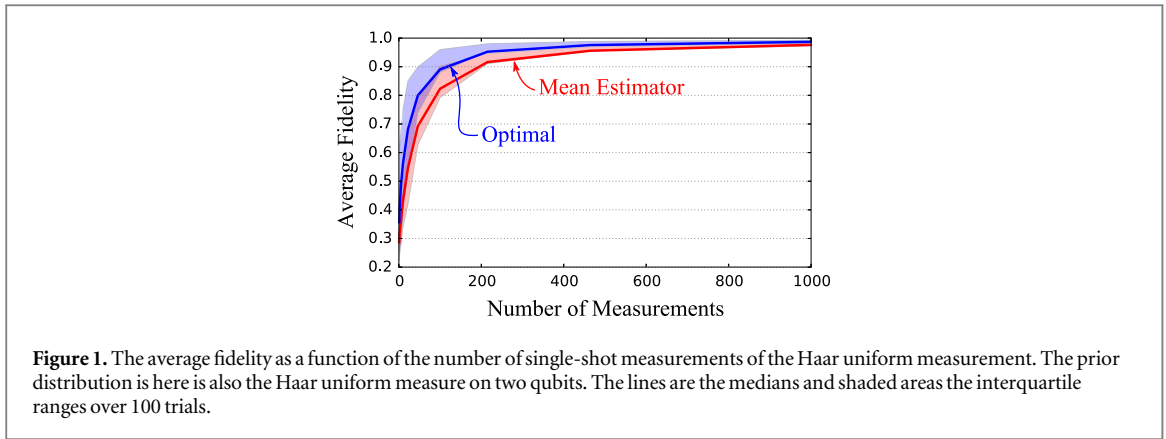
3. Numerical experiments

Note that fidelity achieved by *any* estimator is a lower bound on the one achieved by the optimal estimator. A particularly convenient and generally well motivated [18] estimator is the mean of the distribution $\hat{\rho} = \mathbb{E}_\rho[\rho]$. Our findings underline that for distributions of states relevant to tomography, the mean is very near-optimal. In the context of tomography the mean is furthermore arguably the most convenient estimator, since every other quantity of interest requires its calculation anyway.

Finding an analytical expression for the posterior distribution is a very challenging problem, let alone performing the multidimensional integrals required for the calculation of the expectations above. Thus, we turn to numerics. In particular, we use the sequential Monte Carlo (SMC) algorithm, which has been successfully applied to quantum statistical problems in the context of dynamical parameter estimation [22–24] and quantum state estimation [25–27]. Also, this algorithm is available as an open-source implementation in python [28].

Employing SMC allows us to perform the Bayesian updating and averaging. A complete and detailed discussion of the algorithm appears in [23] and thus we will not repeat the details here, but we will sketch the idea. The algorithm starts with a set of quantum states $\{\rho_j\}_{j=0}^n$, the elements of which are called *particles*. Here, $n = |\{\rho_j\}|$ is the number of particles and controls the accuracy of the approximation. By approximating the prior distribution by a weighted sum of Dirac delta-functions,

$$\text{Pr}(\rho) \approx \sum_{j=1}^n w_j \delta(\rho - \rho_j), \quad (14)$$



Bayes' rule then becomes

$$w_j \mapsto \Pr(\text{data} | \rho_j) w_j, \quad (15)$$

followed by a normalization step. The SMC algorithm is designed to approximate expectation values, such that

$$\mathbb{E}_\rho[f(\rho)] \approx \sum_{j=1}^n w_j f(\rho_j), \quad (16)$$

for any function f . In other words, the SMC algorithm allows us to efficiently compute the multidimensional integrals with respect to the measure defined by the posterior probability distribution. We use this algorithm, as implemented by [28], to numerically compute averages arising in simulated tomography experiments. By doing so, we explore the efficacy of our claims for a variety of distributions relevant to practice and found natural in experimentation.

Recall the sharp distinction between measures supported on pure states and those with full support. We use the fact that theorem 1 provides us with the optimal estimator in the former case to lend support to the claim that the mean estimator is a good candidate for a computationally simple, yet still near-optimal, alternative to solving the optimization problem in general. In figure 1, we present the results of numerical simulations on two qubits. Plotted is the average fidelity achieved by the optimal estimator (see theorem 1) and the mean estimator $\mathbb{E}_\rho[\rho]$. The average is taken with respect to a distribution that begins as the Haar invariant measure on pure states and is updated through simulated measurement data, where the measurement is the 'uniform POVM' consisting of all pure states, distributed uniformly according to the Haar measure. For independent measurements—i.e. local, non-adaptive ones—this measurement is optimal [29, theorem 3.1]. We see that the mean estimator's fidelity tracks the optimal fidelity quite well.

In figure 2, we plot the average fidelity of the mean estimator against our bound (12) for measures supported also on mixed quantum states. Again, we simulate measurement data to get an accurate sense of how well the average fidelity of the mean estimator performs with respect to our bound for distributions relevant to tomography. In this case, the *prior* distribution is either the Hilbert–Schmidt measure (left column), or the *arcsine* and *Bures* distributions [30] for two qubits (right column). In each case, many other natural distributions appear as we update our prior through Bayes' rule. We see again that the mean estimator is a 'good' estimator in that it comes close to the bound on the optimal fidelity and is the easiest non-trivial average quantity to evaluate.

4. Proofs

In this section we provide detailed derivations and proofs of the statements presented in section 2.

4.1. A detailed proof of theorem 2

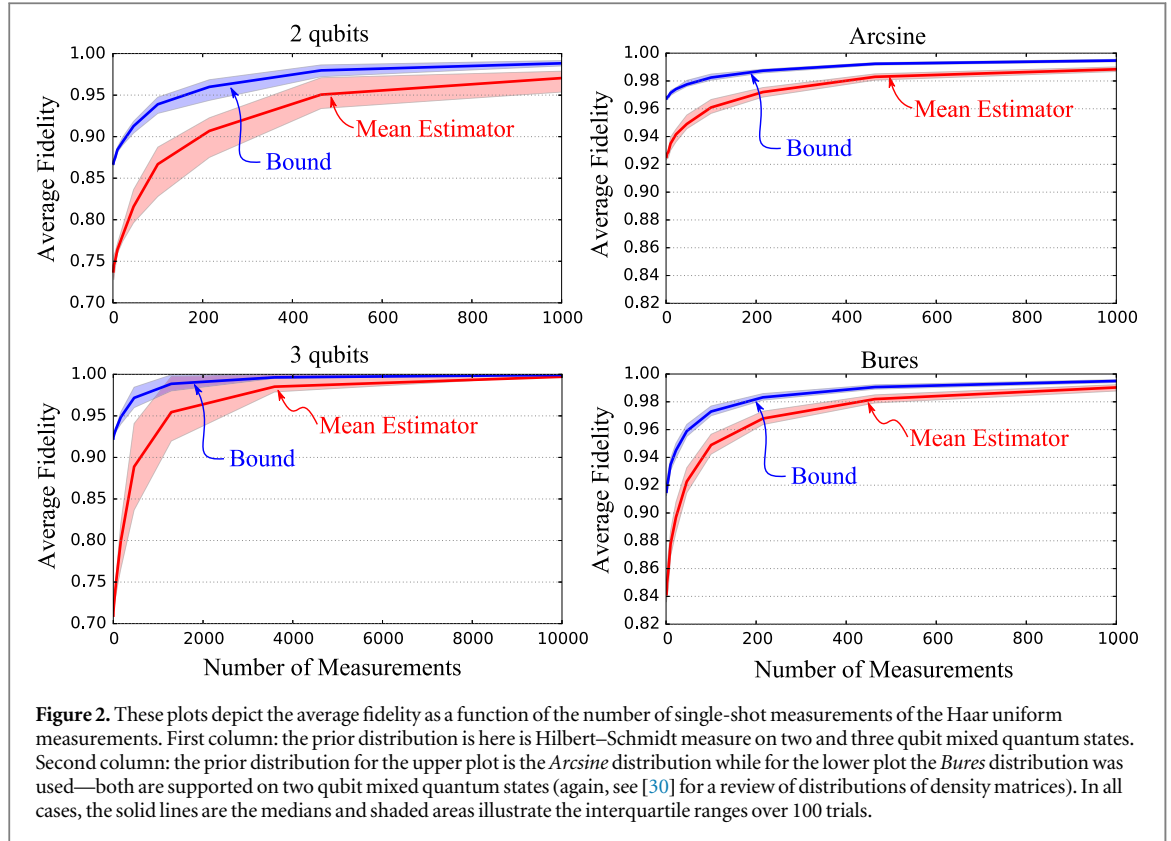
Recall that in theorem 2 we have claimed that the bound

$$\max_{\sigma \in \mathcal{S}} \mathbb{E}_\rho[F(\rho, \sigma)] \leq 1 - \frac{1}{4} \text{Tr} \left(\mathbb{E}_\rho[\rho^2] - \mathbb{E}[\rho]^2 \right), \quad (17)$$

is valid for any prior distribution $d\rho$. In order to derive such a statement, we start with inequality (7)

$$F(\rho, \sigma) \leq 1 - \frac{1}{4} \|\rho - \sigma\|_2^2,$$

which is a direct combination of the Fuchs–van de Graaf inequalities and the norm inequality $\|\cdot\|_2 \leq \|\cdot\|_1$. As such it is valid for any two states $\rho, \sigma \in \mathcal{S}$ which in turn assures that it remains valid upon taking expectations



over $d\rho$ on both sides:

$$\mathbb{E}_\rho[F(\rho, \sigma)] \leq 1 - \frac{1}{4} \mathbb{E}_\rho[\|\rho - \sigma\|_2^2]. \quad (18)$$

Moreover, we can optimize over σ on both sides to obtain

$$\max_{\sigma \in \mathcal{S}} \mathbb{E}_\rho[F(\rho, \sigma)] \leq 1 - \frac{1}{4} \min_{\sigma \in \mathcal{S}} \mathbb{E}_\rho[\|\rho - \sigma\|_2^2]. \quad (19)$$

The minimum on the right-hand side can in fact be calculated analytically. To this end, we define the function

$$f(\sigma) := \mathbb{E}_\rho[\|\rho - \sigma\|_2^2] = \text{Tr}(\mathbb{E}_\rho[\rho^2]) - 2 \text{Tr}(\mathbb{E}_\rho[\rho]\sigma) + \text{Tr}(\sigma^2).$$

Note that $f(\sigma)$ is convex, because it corresponds to a weighted average of convex norm-functions $\|\sigma - \rho\|_2^2$ and its matrix-valued derivative corresponds to

$$f'(\sigma) = -2\mathbb{E}_\rho[\rho] + 2\sigma. \quad (20)$$

This derivative vanishes if and only if $\sigma^\# = \mathbb{E}_\rho[\rho]$ holds and convexity of $f(\sigma)$ implies that this critical state corresponds to the unique minimum. The corresponding function value amounts to

$$f(\sigma^\#) = \text{Tr}(\mathbb{E}_\rho[\rho^2]) - \text{Tr}(\mathbb{E}_\rho[\rho]^2) \quad (21)$$

and reinserting this global minimum into (19) yields the desired bound (17).

4.2. A detailed derivation of theorem 3

Our main theoretical statement—theorem 3—follows from a three step procedure which was already briefly outlined in section 2.

The first step invokes the concept of super-fidelity [20] which assures

$$\max_{\sigma \in \mathcal{S}} \mathbb{E}_\rho[F(\rho, \sigma)] \leq \max_{\sigma \in \mathcal{S}} \left(\text{Tr}(\hat{\rho}\sigma) + p_\rho \sqrt{1 - \text{Tr}(\sigma^2)} \right),$$

with $\hat{\rho} = \mathbb{E}_\rho[\rho]$ and $p_\rho = \mathbb{E}_\rho[\sqrt{1 - \text{tr}(\rho^2)}]$ for any distribution $d\rho$. As it turns out, the optimization on the right-hand side of this equation is much more tractable than the original problem on the left-hand side. This is manifested by the following technical statement which is a direct consequence of the celebrated Birkhoff–von Neumann theorem.

Lemma 1. Fix any $p_\rho \geq 0$ and suppose that $\hat{\rho} \in \mathcal{S}$ is an arbitrary density operator with eigenvalue decomposition $\hat{\rho} = \sum_{i=1}^d \hat{r}_i |b_i\rangle\langle b_i|$. Then the optimization

$$\begin{aligned} & \underset{\sigma \in L(\mathbb{C}^d)}{\text{maximize}} && \text{Tr}(\hat{\rho}\sigma) + p_\rho \sqrt{1 - \text{Tr}(\sigma^2)}, \\ & \text{subject to} && \sigma \geq 0, \text{Tr}(\sigma) = 1 \end{aligned} \tag{22}$$

is equivalent to solving

$$\begin{aligned} & \underset{s_1, \dots, s_d \in \mathbb{R}}{\text{maximize}} && \sum_{i=1}^d \hat{r}_i s_i + p_\rho \sqrt{1 - \sum_{i=1}^d s_i^2}, \\ & \text{subject to} && \sum_{i=1}^d s_i = 1, \\ & && s_i \geq 0 \quad 1 \leq i \leq d. \end{aligned} \tag{23}$$

Moreover, there is a one-to-one correspondence between any feasible array (s_1, \dots, s_d) of this problem and the density operator $\tilde{\sigma} = \sum_{i=1}^d s_i |b_i\rangle\langle b_i|$.

Proof. At the heart of this statement is an immediate corollary of the Birkhoff–von Neumann theorem—see e.g. [21, theorem 8.7.6]. For $d \times d$ Hermitian matrices ρ, σ this corollary assures

$$\text{Tr}(\rho\sigma) \leq \sum_{i=1}^d r_i s_i, \tag{24}$$

where r_i and s_i denote the eigenvalues of ρ and σ , respectively, arranged in non-increasing order. If $\hat{\rho}$ has eigenvalue decomposition $\hat{\rho} = \sum_{i=1}^d \hat{r}_i |b_i\rangle\langle b_i|$, the right-hand side of (24) corresponds to $\text{Tr}(\hat{\rho}\tilde{\sigma})$ where $\tilde{\sigma} = \sum_{i=1}^d s_i |b_i\rangle\langle b_i|$. Clearly, if $\sigma \in \mathcal{S}$ was a quantum state to begin with, so is $\tilde{\sigma}$, because the spectra of σ and $\tilde{\sigma}$ coincide. Moreover, such a definition assures that both states have equal purity, i.e. $\text{Tr}(\sigma^2) = \text{Tr}(\tilde{\sigma}^2)$. Consequently, for any feasible point σ of the optimization (22), there is a $\tilde{\sigma}$ of the above form which admits a larger value in the optimization. Inserting the particular form of $\tilde{\sigma}$ into this program results in (23). \square

In order to arrive at the bound presented in theorem 3, we employ one more relaxation which is going to allow us to solve the resulting problem analytically in full generality. To be concrete, we replace the non-negativity constraints ($s_i \geq 0$) in (23) by the weaker demand that the optimization vector $(s_1, \dots, s_d)^T \in \mathbb{R}^d$ is contained in the Euclidean unit ball—i.e. $\sum_{i=1}^d s_i^2 \leq 1$. Note that we explore the geometric properties of such a relaxation in section 5. In a nutshell it corresponds to the tightest possible elliptical relaxation of the feasible set in (22). By doing so, we arrive at the problem

$$\begin{aligned} & \underset{s_1, \dots, s_d \in \mathbb{R}}{\text{maximize}} && \sum_{i=1}^d \hat{r}_i s_i + p_\rho \sqrt{1 - \sum_{i=1}^d s_i^2}, \\ & \text{subject to} && \sum_{i=1}^d s_i = 1, \sum_{i=1}^d s_i^2 \leq 1, \end{aligned} \tag{25}$$

which can be solved analytically via the method of Lagrangian multipliers:

Lemma 2. Let $\hat{r}_1, \dots, \hat{r}_d$ denote the eigenvalues of any density operator and fix $p_\rho > 0$. Then the problem (25) has a unique solution. The optimal value corresponds to

$$\frac{1}{d} \left(1 + \sqrt{d-1} \sqrt{d(p_\rho^2 + \text{Tr}(\hat{\rho}^2)) - 1} \right)$$

and the array $(s_1^\sharp, \dots, s_d^\sharp)$ achieving this optimum corresponds to the particular matrix

$$\sigma^\sharp = \frac{1}{d} \mathbf{1} + \sqrt{\frac{d-1}{d(p_\rho^2 + \text{Tr}(\hat{\rho}^2)) - 1}} \left(\hat{\rho} - \frac{1}{d} \mathbf{1} \right). \tag{26}$$

Note that this result together with the relaxations outlined in this section immediately implies theorem 3 upon inserting the definitions of p_ρ and $\hat{\rho}$. The assumption $p_\rho > 0$ is furthermore non-critical, because, by definition, $p_\rho = 0$ if and only if ρ is supported exclusively on pure states. This particular case, however, is already fully covered by theorem 1.

Proof of lemma 2. Throughout this proof we shall represent the eigenvalues of the density operator $\hat{\rho}$ as a vector $\mathbf{r} = (\hat{r}_1, \dots, \hat{r}_d)^T \in \mathbb{R}^d$. Likewise we shall encompass the scalar optimization variables s_i in the vector $\mathbf{s} \in \mathbb{R}^d$. Furthermore, let $\mathbf{0} = (0, \dots, 0)^T$ and $\mathbf{1} = (1, \dots, 1)^T$ denote the ‘all-zeros’ and ‘all-ones’ vectors on \mathbb{R}^d , respectively. For $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, we will also make use of the standard inner product $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^d x_i y_i$ and the vectorial inequality $\mathbf{x} \geq \mathbf{y}$ shall indicate component-wise inequality, i.e. $x_i \geq y_i$ for all $1 \leq i \leq d$.

In such a vectorial form, the optimization problem (23) corresponds to

$$\begin{aligned} & \text{maximize } f(\mathbf{s}) = \langle \mathbf{r}, \mathbf{s} \rangle + p_\rho \sqrt{1 - \langle \mathbf{s}, \mathbf{s} \rangle}, \\ & \text{subject to } g(\mathbf{s}) = \langle \mathbf{1}, \mathbf{s} \rangle = 0. \\ & \langle \mathbf{s}, \mathbf{s} \rangle \leq 1. \end{aligned} \tag{27}$$

Note that (27) is a convex optimization problem, as it requires maximizing a concave function over a convex set. As such, it has a unique maximum. One way of finding this maximum is to apply standard techniques such as the Karush–Kuhn–Tucker multiplier method [17] which are designed to take into account the inequality constraint (28).

However, here we opt for a less direct but considerably more convenient and less cumbersome approach: we ignore the inequality constraint in (27) for now and employ the standard technique of Lagrangian multipliers (for equality constraints) in order to find the unique critical point $\mathbf{s}^\#$ of the optimization. In a second step, we are going to verify that this vector strictly obeys the additional inequality constraint, we have ignored so far, i.e. $\langle \mathbf{s}^\#, \mathbf{s}^\# \rangle < 1$. This in turn implies that said inequality constraint is not active at the critical point which in retrospect confirms that we were in fact right to ignore it in the first place. Finally, the fact that we face a convex optimization problem assures that this unique critical point indeed yields the sought for global maximum of (27).

In order to find the critical point $\mathbf{s}^\#$ in question we define the Lagrangian function

$$L(\mathbf{s}) = f(\mathbf{s}) + \lambda g(\mathbf{s}), \tag{28}$$

where we have—as already announced—ignored the inequality constraint $\langle \mathbf{s}, \mathbf{s} \rangle \leq 1$. As a consequence, $\lambda \in \mathbb{R}$ denotes the single Lagrangian multiplier associated with the remaining normalization constraint. The necessary condition for an optimal solution of (27) then reads

$$\mathbf{r} - \frac{p_\rho \mathbf{s}}{\sqrt{1 - \langle \mathbf{s}, \mathbf{s} \rangle}} + \lambda \mathbf{1} = \mathbf{0}. \tag{29}$$

Taking the inner product of this vector-identity with the ‘all-ones’ vector $\mathbf{1}$ results in

$$0 = \langle \mathbf{1}, \mathbf{0} \rangle = \langle \mathbf{1}, \mathbf{r} \rangle - \frac{p_\rho \langle \mathbf{1}, \mathbf{s} \rangle}{\sqrt{1 - \langle \mathbf{s}, \mathbf{s} \rangle}} + \lambda \langle \mathbf{1}, \mathbf{1} \rangle = 1 - \frac{p_\rho}{\sqrt{1 - \langle \mathbf{s}, \mathbf{s} \rangle}} + d\lambda, \tag{30}$$

where we have used $\langle \mathbf{1}, \mathbf{r} \rangle = \sum_{i=1}^d \hat{r}_i = \text{Tr}(\hat{\rho}) = 1$ and the normalization constraint, which likewise assures $\langle \mathbf{1}, \mathbf{s} \rangle = 1$. This equation allows us to replace $\sqrt{1 - \langle \mathbf{s}, \mathbf{s} \rangle}$ by $\frac{p_\rho}{1 + d\lambda}$ and reinserting this into (29) results in the equivalent vector equation

$$\mathbf{r} - (1 + d\lambda)\mathbf{s} + \lambda \mathbf{1} = \mathbf{0}. \tag{31}$$

This can be readily inverted to yield

$$\mathbf{s} = \frac{1}{1 + d\lambda}(\mathbf{r} + \lambda \mathbf{1}). \tag{32}$$

In order to determine the value of λ , we revisit (30) which in combination with (32) demands

$$\begin{aligned} p_\rho^2 &= (1 + d\lambda)^2(1 - \langle \mathbf{s}, \mathbf{s} \rangle) = (1 + d\lambda)^2 - \langle \mathbf{r}, \mathbf{r} \rangle - 2\lambda \langle \mathbf{1}, \mathbf{r} \rangle - \lambda^2 \langle \mathbf{1}, \mathbf{1} \rangle \\ &= d(d - 1)\lambda^2 + 2(d - 1)\lambda + 1 - \text{Tr}(\hat{\rho}^2), \end{aligned} \tag{33}$$

where we have once more used $\langle \mathbf{1}, \mathbf{r} \rangle = 1$ as well as $\langle \mathbf{r}, \mathbf{r} \rangle = \sum_{i=1}^d \hat{r}_i^2 = \text{Tr}(\hat{\rho}^2)$. This results in the quadratic equation

$$\lambda^2 + \frac{2}{d}\lambda - \frac{1}{d(d - 1)}(p_\rho^2 + \text{Tr}(\hat{\rho}^2) - 1), \tag{34}$$

for λ whose two possible solutions correspond to

$$\lambda_\pm = -\frac{1}{d} \left(1 \mp \sqrt{\frac{d(p_\rho^2 + \text{Tr}(\hat{\rho}^2)) - 1}{d - 1}} \right). \tag{35}$$

Note that the argument of the square-root is non-negative, because the purity $\text{Tr}(\hat{\rho}^2)$ of any quantum state is lower-bounded by $1/d$. Also, the second solution λ_- is vacuous, since it leads to an immediate contradiction. Indeed, it follows by inspection that $\lambda_- < -1/d$ holds. Together with (30) this implies the contradictory relation

$$\sqrt{1 - \langle \mathbf{s}, \mathbf{s} \rangle} = \frac{p_\rho}{1 + d\lambda_-} < 0, \quad (36)$$

because p_ρ is positive by assumption.

Consequently we are left with one meaningful value λ_+ for the Lagrangian multiplier and inserting it into (32) yields the unique critical solution

$$\mathbf{s}^\# = \frac{1}{d} \mathbf{1} + \sqrt{\frac{d-1}{d(p_\rho^2 + \text{Tr}(\hat{\rho}^2)) - 1}} \left(\mathbf{r} - \frac{1}{d} \mathbf{1} \right). \quad (37)$$

Recall that throughout this proof we are exploiting a one-to-one correspondence between vectors $\mathbf{s} = (s_1, \dots, s_n)^T \in \mathbb{R}^d$ and hermitian $d \times d$ -matrices $\sigma = \sum_{i=1}^n s_i |b_i\rangle\langle b_i|$ that commute with $\hat{\rho}$. Consequently, the critical vector $\mathbf{s}^\#$ corresponds to the critical matrix presented in (26).

Plugging the critical point $\mathbf{s}^\#$ into the objective function $f(\mathbf{s})$ furthermore yields the corresponding critical function value

$$\begin{aligned} f(\mathbf{s}^\#) &= \langle \mathbf{r}, \mathbf{s}^\# \rangle + p_\rho \sqrt{1 - \langle \mathbf{s}^\#, \mathbf{s}^\# \rangle} = \frac{\langle \mathbf{r}, \mathbf{r} \rangle + \lambda_+ \langle \mathbf{1}, \mathbf{r} \rangle}{1 + d\lambda_+} + \frac{p_\rho^2}{1 + d\lambda_+} \\ &= \frac{d(p_\rho^2 + \text{Tr}(\hat{\rho}^2)) - 1 + 1 + d\lambda_+}{d(1 + d\lambda_+)}, \\ &= \frac{1}{d} \left(1 + \frac{d(p_\rho^2 + \text{Tr}(\hat{\rho}^2)) - 1}{1 + d\lambda_+} \right) = \frac{1}{d} \left(1 + \sqrt{d-1} \sqrt{d(p_\rho^2 + \text{Tr}(\hat{\rho}^2)) - 1} \right), \end{aligned} \quad (38)$$

where we have once more replaced $\sqrt{1 - \langle \mathbf{s}^\#, \mathbf{s}^\# \rangle}$ by $\frac{p_\rho}{(1 + d\lambda_+)}$ and combined that with the fact that

$$(1 + d\lambda_+) = \sqrt{\frac{d(p_\rho^2 + \text{Tr}(\hat{\rho}^2)) - 1}{d-1}} \text{ holds.}$$

With such a unique critical point $\mathbf{s}^\#$ at hand, we are now ready to show that it strictly obeys the inequality constraint $\langle \mathbf{s}^\#, \mathbf{s}^\# \rangle$ we have ignored so far. By employing the same equalities we have used in the previous paragraph, we can readily establish such a claim:

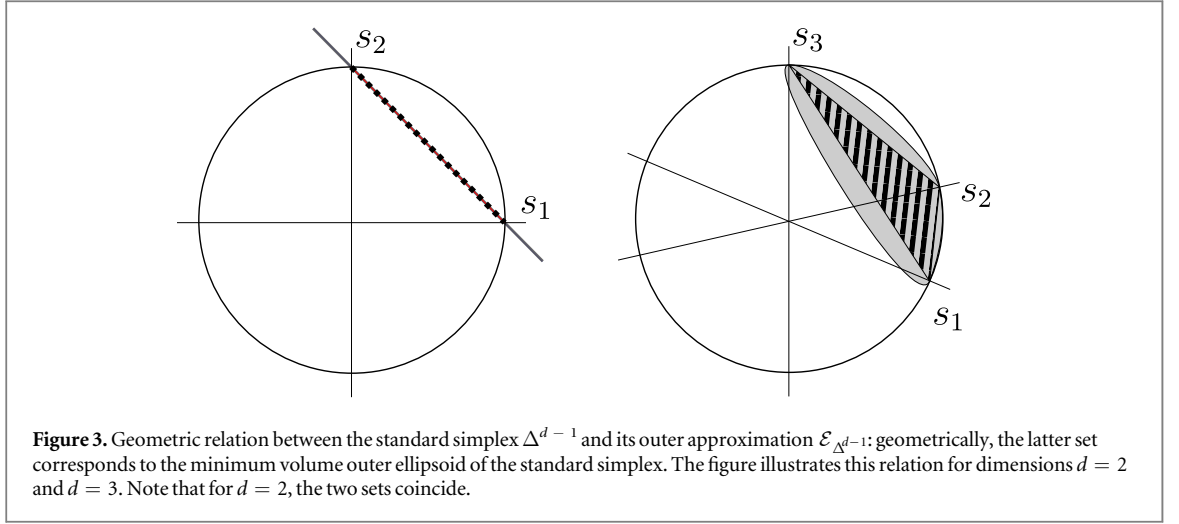
$$\langle \mathbf{s}^\#, \mathbf{s}^\# \rangle = 1 - (1 - \langle \mathbf{s}^\#, \mathbf{s}^\# \rangle) = 1 - \frac{p_\rho^2}{(1 + d\lambda_+)^2} < 1. \quad (39)$$

The strict inequality on the right follows from the fact that $p_\rho > 0$ holds by assumption. This indeed establishes, that $\mathbf{s}^\#$ is also a critical point of the optimization problem (27). Since this optimization corresponds to maximizing a concave function over a convex set, the unique critical point $\mathbf{s}^\#$ must correspond to the unique maximum of (27). \square

4.3. Detailed proofs of corollary 1 and corollary 2

We conclude the proof section with providing detailed proofs of the remaining statements, namely that section 3 reproduces the main result in [6] for the particular case of a single qubit, i.e. $d = 2$ (corollary 1) that the bounds presented in theorem 3 are strictly better than the ones outlined in theorem 2 (corollary 2).

Proof of corollary 1. We start this section by pointing out that in the particular case of dimension $d = 2$, the two relaxations we have employed in the previous subsection are not relaxations at all. Indeed, for dimension two, fidelity and super-fidelity coincide, and moreover the sets $\{(y_1, y_2)^T \in \mathbb{R}^2: y_1 + y_2 = 1, y_1, y_2 \geq 0\}$ and $\{(y_1, y_2)^T \in \mathbb{R}^2: y_1 + y_2 = 1, y_1^2 + y_2^2 \leq 1\}$ coincide (this one-to-one correspondence is illustrated in figure 3 below). These low-dimensional equivalences assure that all the relaxations employed in the derivation of theorem 3 are actually tight. Consequently, in this particular low-dimensional case, we solve the actual problem of interest.



For deducing the claimed statement from this fact, we consider equation (2.9) in [6]:

$$F = \frac{1}{2} \left(1 + \sum_{\chi} \|\mathbf{V}_{\chi}\|_2 \right). \quad (40)$$

Here χ simply means the data generated via the measurement. The vector \mathbf{V}_{χ} is defined as follows:

$$\mathbf{V}_{\chi} = \mathbb{E}_{\rho} [\mathbf{r} \Pr(\chi | \rho)], \quad (41)$$

where \mathbf{r} is related to the usual Bloch vector $\mathbf{r} = (x, y, z)$ via

$$\mathbf{r} = \left(\sqrt{1 - \|\mathbf{r}\|_2^2}, \mathbf{r} \right). \quad (42)$$

We point out that this F is not the same average fidelity we have considered but the following quantity (which corresponds to our equation (4) above):

$$F = \max_{\sigma} \mathbb{E}_{\rho} \left[\mathbb{E}_{\chi | \rho} [F(\rho, \sigma(\chi))] \right]. \quad (43)$$

Note however that, by employing Bayes' rule, this is equal to

$$F = \max_{\sigma} \mathbb{E}_{\chi} \left[\mathbb{E}_{\rho | \chi} [F(\rho, \sigma(\chi))] \right], \quad (44)$$

and thus maximizing the posterior average fidelity is equivalent to maximizing the total average fidelity. Our bound applies directly to the former but trivially extends to the latter.

Thus, to establish corollary 1, we need to extract the posterior average fidelity from the expressions above. First, using Bayes' rule, we calculate

$$\mathbf{V}_{\chi} = \Pr(\chi) \mathbb{E}_{\rho | \chi} [\mathbf{r}]. \quad (45)$$

Using the fact that $\|\mathbf{r}\|_2^2 = 2 \operatorname{Tr}(\rho^2) - 1$ and

$$\operatorname{Tr}(\mathbb{E}_{\rho | \chi} [\rho]^2) = \frac{1}{2} \left(1 + \|\mathbb{E}_{\rho | \chi} [\mathbf{r}]\|_2^2 \right), \quad (46)$$

we find

$$\|\mathbf{V}_{\chi}\|_2^2 = \Pr(\chi)^2 \left(2 \mathbb{E}_{\rho | \chi} \left[\sqrt{1 - \operatorname{Tr}(\rho^2)} \right]^2 + 2 \operatorname{Tr}(\mathbb{E}_{\rho | \chi} [\rho]^2) - 1 \right). \quad (47)$$

Plugging this back into (40), we have

$$F = \frac{1}{2} \left(1 + \sum_{\chi} \Pr(\chi) \sqrt{2 \mathbb{E}_{\rho | \chi} \left[\sqrt{1 - \operatorname{Tr}(\rho^2)} \right]^2 + 2 \operatorname{Tr}(\mathbb{E}_{\rho | \chi} [\rho]^2) - 1} \right), \quad (48)$$

$$= \frac{1}{2} \left(1 + \mathbb{E}_{\chi} \left[\sqrt{2 \mathbb{E}_{\rho | \chi} \left[\sqrt{1 - \operatorname{Tr}(\rho^2)} \right]^2 + 2 \operatorname{Tr}(\mathbb{E}_{\rho | \chi} [\rho]^2) - 1} \right] \right), \quad (49)$$

$$= \mathbb{E}_\chi \left[\frac{1}{2} \left(1 + \sqrt{2\mathbb{E}_{\rho|\chi} \left[\sqrt{1 - \text{Tr}(\rho^2)} \right]^2 + 2 \text{Tr} \left(\mathbb{E}_{\rho|\chi} [\rho^2] \right) - 1} \right) \right]. \quad (50)$$

Thus, implied by the results of [6], the maximum posterior average fidelity (dropping the χ for parallelism) is

$$\max_\sigma \mathbb{E}_\rho [F(\rho, \sigma)] = \frac{1}{2} \left(1 + \sqrt{2 \left(\mathbb{E}_\rho \left[\sqrt{1 - \text{Tr}(\rho^2)} \right]^2 + \text{Tr} \left(\mathbb{E}_\rho [\rho^2] \right) \right) - 1} \right). \quad (51)$$

This coincides with our main result (12) for dimension $d = 2$. □

Proof of corollary 2. For notational simplicity, let us introduce the short-hand notation

$$s_\rho := \text{Tr} \left(\mathbb{E}_\rho [\rho^2] \right) - \text{Tr} \left(\mathbb{E}_\rho [\rho]^2 \right), \quad (52)$$

such that the bound presented in theorem 2 simply reads $\max_{\sigma \in \mathcal{S}} \mathbb{E}_\rho [F(\rho, \sigma)] \leq 1 - \frac{s_\rho}{4}$. Note furthermore that $0 \leq s_\rho \leq 1$ holds. As already mentioned, the lower bound follows from invoking Jensen’s inequality, while the upper bound is a simple consequence of the fact that the purity of any state is at most one. A vanishing s_ρ would correspond to a trivial Fuchs–van de Graaf bound of one which is the first case instance covered by corollary 2. Therefore we can from now on safely assume that $s_\rho > 0$ holds. Under this assumption we prove the second claim by starting with the bound presented in theorem 3 and upper-bounding it via a chain of inequalities which will ultimately lead to the bound presented in theorem 2. Indeed, pick any dimension d and an arbitrary distribution $d \rho$ over states. Then Jensen’s inequality assures

$$\mathbb{E}_\rho \left[\sqrt{1 - \text{Tr}(\rho^2)} \right]^2 \leq 1 - \mathbb{E}_\rho \left[\text{Tr}(\rho^2) \right], \quad (53)$$

and the right-hand side of expression (12) in theorem 3 can be upper-bounded by

$$\frac{1}{d} + \frac{\sqrt{d-1}}{d} \sqrt{d-1-ds_\rho}, \quad (54)$$

because the square root function is monotonically increasing on the positive reals. Adding and subtracting s_ρ in the last square root and once more invoking monotonicity allows us to continue via

$$\frac{1}{d} + \frac{\sqrt{d-1}}{d} \sqrt{(d-1)(1-s_\rho) - s_\rho} < \frac{1}{d} + \frac{d-1}{d} \sqrt{1-s_\rho}, \quad (55)$$

where we have used $s_\rho > 0$ in the last line to obtain strict inequality. Since the square root is a concave function, the inequality $\sqrt{1-s_\rho} \leq 1 - \frac{1}{2}s_\rho$ is valid for any $s_\rho \leq 1$ and consequently

$$\frac{1}{d} + \frac{d-1}{d} \sqrt{1-s_\rho} \leq 1 - \frac{d-1}{2d} s_\rho, \quad (56)$$

is true. Finally, we use the simple fact that $\frac{d-1}{d} \geq \frac{1}{2}$ holds for any $d \geq 2$ to arrive at $1 - \frac{1}{4}s_\rho$ which is just the Fuchs–van de Graaf bound. Since a strict inequality sign connects the expressions in (55), the claimed strict majorization follows. □

5. Geometric interpretation of the relaxation leading to equation (25)

Recall that in order to arrive at theorem 3, we have replaced the feasible set

$$\Delta^{d-1} = \left\{ \mathbf{s} \in \mathbb{R}^d: \langle \mathbf{1}, \mathbf{s} \rangle = 1, \mathbf{s} \geq \mathbf{0} \right\}, \quad (57)$$

of the optimization problem (11) by

$$\mathcal{E}_{\Delta^{d-1}} = \left\{ \mathbf{s} \in \mathbb{R}^d: \langle \mathbf{1}, \mathbf{s} \rangle = 1, \langle \mathbf{s}, \mathbf{s} \rangle \leq 1 \right\}, \quad (58)$$

which is a convex outer approximation of Δ^{d-1} . This follows from the basic fact that $x^2 \leq x$ holds for any $x \in [0, 1]$. Since the vector components s_i of any $\mathbf{s} \in \Delta^{d-1}$ have to obey $s_i \in [0, 1]$, we can readily conclude

$$\langle \mathbf{s}, \mathbf{s} \rangle = \sum_{i=1}^d s_i^2 \leq \sum_{i=1}^d s_i = 1. \quad (59)$$

Note that the converse is true if and only if $d = 1, 2$ —a fact which we have exploited in proving corollary 1.

Geometrically, the former set corresponds to the standard simplex in \mathbb{R}^d . In this section we prove that the latter one is in fact the minimum volume covering ellipsoid of the standard simplex which furthermore corresponds to a $(d - 1)$ -dimensional Euclidean ball. For dimensions two and three this situation is illustrated in figure 3.

Proposition 1. (Geometric nature of $\mathcal{E}_{\Delta^{d-1}}$) *The convex outer-approximation $\mathcal{E}_{\Delta^{d-1}}$ of the d -simplex corresponds to a $(d - 1)$ -dimensional Euclidean ball with radius $\sqrt{\frac{d-1}{d}}$ and center $\frac{1}{\sqrt{d}}\mathbf{1}$ which is contained in the $(d - 1)$ -dimensional hyperplane $\mathcal{H}_{1,1} := \{\mathbf{s} \in \mathbb{R}^d: \langle \mathbf{1}, \mathbf{s} \rangle = 1\}$.*

Proof. By definition, the set $\mathcal{E}_{\Delta^{d-1}}$ corresponds to the intersection of the Euclidean unit ball $\mathcal{B}_1(\mathbf{0}) = \{\mathbf{s} \in \mathbb{R}^d: \langle \mathbf{s}, \mathbf{s} \rangle \leq 1\}$ and the hyperplane $\mathcal{H}_{1,1}$. This assures $\mathcal{E}_{\Delta^{d-1}} \subseteq \mathcal{H}_{1,1}$ by construction.

One way to establish that $\mathcal{E}_{\Delta^{d-1}}$ is furthermore itself an Euclidean ball, is using ‘generalized cylindrical coordinates’ for the Euclidean unit ball $\mathcal{B}^d(\mathbf{0}, 1)$: Such coordinates use the fact that $\mathcal{B}^d(\mathbf{0}, 1)$ is equivalent to the union of a family of $(d - 1)$ -dimensional unit balls. More concretely: let $\mathbf{z} \in \mathbb{R}^d$ be an arbitrary unit vector and let $\zeta \in \mathbb{R}$ denote a parameter. For each value of this parameter, we define the hyperplane $\tilde{\mathcal{H}}_{\mathbf{z},\zeta} = \{\mathbf{s} \in \mathbb{R}^d: \langle \mathbf{z}, \mathbf{s} \rangle = \zeta\}$ which in particular contains the vector $\zeta\mathbf{z}$ by construction. Furthermore, let $\tilde{\mathcal{B}}^{d-1}(\mathbf{z}, \zeta) \subset \tilde{\mathcal{H}}_{\mathbf{z},\zeta}$ be the $(d - 1)$ -dimensional Euclidean ball with radius $\sqrt{1 - \zeta^2}$ and center $\zeta\mathbf{z}$ that is contained in the hyperplane $\tilde{\mathcal{H}}_{\mathbf{z},\zeta}$. Clearly each element in such a union of sets is contained in the d -ball, and letting ζ range from -1 to 1 covers the entire d -ball. In order to see this, decompose any $\mathbf{s} \in \mathcal{B}^d(\mathbf{0}, 1)$ as $\mathbf{s} = \langle \mathbf{s}, \mathbf{z} \rangle \mathbf{z} + \mathbf{z}^\perp$ such that $\langle \mathbf{z}^\perp, \mathbf{z} \rangle = 0$ and set $\zeta = \langle \mathbf{s}, \mathbf{z} \rangle$. Pythagoras’ theorem then assures $\|\mathbf{z}^\perp\|_2 \leq \sqrt{1 - \zeta^2}$ and consequently $\mathbf{s} \in \tilde{\mathcal{B}}^{d-1}(\mathbf{z}, \zeta)$.

The structure of the particular problem at hand suggests to fix $\mathbf{z} = \frac{1}{\sqrt{d}}\mathbf{1}$. Indeed, such a particular choice of \mathbf{z} assures equality of the hyperplane $\mathcal{H}_{1,1}$ which contains $\mathcal{E}_{\Delta^{d-1}}$ and the hyperplane $\tilde{\mathcal{H}}_{\frac{1}{\sqrt{d}}\mathbf{1}, \frac{1}{\sqrt{d}}}$, we have just introduced. Consequently, the ‘cylindrical representation’ of the Euclidean unit ball assures that the intersection $\mathcal{E}_{\Delta^{d-1}} = \mathcal{B}_1(\mathbf{0}) \cap \mathcal{H}_{1,1}$ corresponds to the $(d - 1)$ -ball $\tilde{\mathcal{B}}^{d-1}\left(\frac{1}{\sqrt{d}}\mathbf{1}, \frac{1}{\sqrt{d}}\right)$ associated with the hyperplane $\tilde{\mathcal{H}}_{\frac{1}{\sqrt{d}}\mathbf{1}, \frac{1}{\sqrt{d}}}$ and a parameter value $\zeta = \frac{1}{\sqrt{d}}$. By definition, this ball has center $\frac{1}{d}\mathbf{1}$ and radius $\sqrt{1 - \zeta^2} = \sqrt{\frac{d-1}{d}}$ which completes the proof. □

The next statement establishes that our choice of replacing the original feasible set Δ^{d-1} in the proof of theorem 3 by the larger convex set $\mathcal{E}_{\Delta^{d-1}}$ is in a precise sense the tightest possible elliptic relaxation of the original optimization problem.

Proposition 2. *The set $\mathcal{E}_{\Delta^{d-1}}$ is the unique minimal volume covering ellipsoid of the standard simplex Δ^{d-1} .*

The proof exploits the following standard result about Löwner–John ellipsoids that is originally due to John. However, here we make use of a slightly more general version presented in [31].

Theorem 4. *Theorem 2.1 in [31]. Let $K \subset \mathbb{R}^d$ be a convex body and let K be contained in the Euclidean unit ball $\mathcal{B}^d(\mathbf{0})$. Then the following statements are equivalent:*

1. $\mathcal{B}^d(\mathbf{0})$ is the unique minimum volume ellipsoid containing K .
2. There exist contact points $\mathbf{u}_1, \dots, \mathbf{u}_m$ lying both in the boundary of K and $\mathcal{B}^d(\mathbf{0})$, and positive numbers $\lambda_1, \dots, \lambda_m$, $m \geq d$, such that

$$\sum_{i=1}^m \lambda_i \mathbf{u}_i = \mathbf{0} \quad \text{and} \quad \sum_{i=1}^m \lambda_i |\mathbf{u}_i\rangle\langle \mathbf{u}_i| = \mathbf{1}. \tag{60}$$

Proof. In proposition 1 we have established that the set $\mathcal{E}_{\Delta^{d-1}}$ corresponds to a $(d - 1)$ -ball with radius $\sqrt{\frac{d-1}{d}}$ and center $\frac{1}{d}\mathbf{1}$ that (like the standard simplex) is contained in the hyperplane $\mathcal{H}_{1,1}$. A quick calculation reveals that all vertices of the standard simplex Δ^{d-1} —which are just the standard basis vectors $\mathbf{e}_1, \dots, \mathbf{e}_d$ —have Euclidean distance $\sqrt{\frac{d-1}{d}}$ to the ball’s center. Consequently they are contained in the boundary of the ball $\mathcal{E}_{\Delta^{d-1}}$

and we have found sufficiently many contact points for applying theorem 4. Since volume is translationally invariant we can furthermore shift the coordinate's origin into the point $\frac{1}{d}\mathbf{1}$ (which is the center of the ball $\mathcal{E}_{\Delta^{d-1}}$). This has the advantage that the affine space $\mathcal{H}_{1,1}$ containing both Δ^{d-1} and $\mathcal{E}_{\Delta^{d-1}}$ turns into $\mathcal{H}_{1,0}$ which is a linear subspace. Note that with respect to the (translated) standard basis, the orthogonal projection onto this subspace is given by

$$P = \mathbf{1} - \frac{1}{d} |\mathbf{1}\rangle\langle\mathbf{1}|.$$

With respect to this new coordinate system, the d contact points (vertices of the simplex) amount to $\tilde{\mathbf{e}}_i = \mathbf{e}_i - \frac{1}{d}\mathbf{1}$. Choosing unit weights $\lambda_i = 1$ for all $m = d$ contact points $\mathbf{u}_i = \tilde{\mathbf{e}}_i$ and calculating

$$\sum_{i=1}^m \lambda_i \mathbf{u}_i = \sum_{i=1}^n \tilde{\mathbf{e}}_i = \sum_{i=1}^n \left(\mathbf{e}_i - \frac{1}{d}\mathbf{1} \right) = \mathbf{0} \quad (61)$$

reveals that the first condition for theorem 4 is fulfilled. A similar calculation reveals

$$\sum_{i=1}^m \lambda_i |\mathbf{u}_i\rangle\langle\mathbf{u}_i| = \mathbf{1} - \frac{1}{d} |\mathbf{1}\rangle\langle\mathbf{1}|.$$

This, however equals just the projector P onto the subspace $\mathcal{H}_{1,0}$ which contains the entire $(d - 1)$ -dimensional problem of interest. Restricted to its range, a projector corresponds to the identity which establishes the second condition for theorem 4. Since this statement is invariant under re-scaling, we can also apply it here, where the radius of the $(d - 1)$ -dimensional surrounding Euclidean ball is not one but $\sqrt{\frac{d-1}{d}}$. \square

6. Conclusion

In this work we have derived upper bounds on the average fidelity of any estimator with no restrictions on the dimension or the distribution being averaged over. Furthermore, we have shown a sharp distinction in the optimization problems of maximizing average fidelity between measures supported only on pure states and those with full support. In the former case, we have provided the exact optimal estimator, while in both cases we argued based on numerical evidence that the mean estimator is a good proxy for the optimal solution.

Interestingly, we found that the analytical bound (12) (which is based on super-fidelity [20]) is strictly tighter than a corresponding one obtained using the well known, and often used, Fuchs–van de Graaf inequalities [19].

These results have obvious applications to practical Bayesian quantum tomography [13], since the bound can be computed *online*—that is, it is only a property of the current distribution under consideration. But we also expect our bound to be of interest in other theoretical work on tomography, where a benchmark is needed to make statements about absolute average performance of some candidate protocol.

Acknowledgments

We thank Robin Blume-Kohout, Steve Flammia and Patrick Hayden for fruitful discussions. CF was supported by National Science Foundation grant number PHY-1212445, the Canadian Government through the NSERC PDF program, the IARPA MQCO program, the ARC via EQuS project number CE11001013, and by the US Army Research Office grant numbers W911NF-14-1-0098 and W911NF-14-1-0103. RK was supported by scholarship funds from the State Graduate Funding Program of Baden-Württemberg, the Excellence Initiative of the German Federal and State Governments (Grant ZUK 43), the ARO under contracts, W911NF-14-1-0098 and W911NF-14-1-0133 (Quantum Characterization, Verification, and Validation), the Freiburg Research Innovation Fund, and the DFG (GRO 4334-1-1). Furthermore, the authors would like to thank the *Mathematisches Forschungsinstitut Oberwolfach* and the organizers of the Oberwolfach workshop *New Horizons in Statistical Decision Theory* (September 2014, Workshop ID: 1437) where work on this project commenced.

References

- [1] Nielsen M A and Chuang I L 2010 *Quantum Computation and Quantum Information* (Cambridge: Cambridge University Press)
- [2] Wootters W K 1981 Statistical distance and Hilbert space *Phys. Rev. D* **23** 357
- [3] Jozsa R 1994 Fidelity for mixed quantum state *J. Mod. Opt.* **41** 2315
- [4] Emerson J, Weinstein Y S, Lloyd S and Cory D G 2002 Fidelity decay as an efficient indicator of quantum chaos *Phys. Rev. Lett.* **89** 284102
- [5] Khaneja N, Reiss T, Kehlet C, S-Herbruggen T and Glaser S J 2005 Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms *J. Magn. Reson.* **172** 296

- [6] Bagan E, Ballester M A, Gill R D, Monras A and Muñoz-Tapia R 2006 Optimal full estimation of qubit mixed states *Phys. Rev. A* **73** 032301
- [7] Emerson J, da Silva M P, Moussa O, Ryan C A, Laforest M, Baugh J, Cory D G and Laflamme R 2007 Symmetrized characterization of noisy quantum processes *Science* **317** 1893
- [8] Flammia S T and Liu Y-K 2011 Direct fidelity estimation from few Pauli measurements *Phys. Rev. Lett.* **106** 230501
- [9] Da-Silva M P, Landon-Cardinal O and Poulin D 2011 Practical characterization of quantum devices without tomography *Phys. Rev. Lett.* **107** 210404
- [10] Cook R L, Riofrio C A and Deutsch I H 2014 Single-shot quantum state estimation via a continuous measurement in the strong backaction regime *Phys. Rev. A* **90** 032113
- [11] Berger J 1985 *Statistical Decision Theory and Bayesian Analysis* (Berlin: Springer)
- [12] Hradil Z 1997 Quantum-state estimation *Phys. Rev. A* **55** R1561(R)
- [13] Blume-Kohout R 2010 Optimal, reliable estimation of quantum states *New J. Phys.* **12** 043034
- [14] Blume-Kohout R 2010 Hedged maximum likelihood quantum state estimation *Phys. Rev. Lett.* **105** 200504
- [15] Gross D, Liu Y-K, Flammia S T, Becker S and Eisert J 2010 Quantum state tomography via compressed sensing *Phys. Rev. Lett.* **105** 150401
- [16] Kueng R, Rauhut H and Terstiege U 2014 Low rank matrix recovery from rank-one measurements (arXiv:1410.6913)
- [17] Boyd S and Vandenberghe L 2004 *Convex Optimization* (Cambridge: Cambridge University Press)
- [18] Blume-Kohout R and Hayden P Accurate quantum state estimation via Keeping the experimentalist honest (arXiv:quant-ph/0603116)
- [19] Fuchs C A and Graaf J V D 1999 Cryptographic distinguishability measures for quantum-mechanical states *IEEE Trans. Inf. Theory* **45** 1216
- [20] Miszczak J A, Puchała Z, Horodecki P, Uhlmann A and Życzkowski K 2009 Sub- and super-fidelity as bounds for quantum fidelity *Quantum Inf. Comput.* **9** 103
- [21] Horn R A and Johnson C R 2013 *Matrix Analysis* 2nd edn (Cambridge: Cambridge University Press)
- [22] Chase B A and Geremia J M 2009 Single-shot parameter estimation via continuous quantum measurement *Phys. Rev. A* **79** 022314
- [23] Granade C E, Ferrie C, Wiebe N and Cory D G 2012 Robust online Hamiltonian learning *New J. Phys.* **14** 103013
- [24] Wiebe N, Granade C, Ferrie C and Cory D G 2014 Hamiltonian learning and certification using quantum resources *Phys. Rev. Lett.* **112** 190501
- [25] Huszar F and Houlby N M T 2012 Adaptive Bayesian quantum tomography *Phys. Rev. A* **85** 052120
- [26] Ferrie C 2014 High posterior density ellipsoids of quantum states *New J. Phys.* **16** 023006
- [27] Ferrie C 2014 Quantum model averaging *New J. Phys.* **16** 093035
- [28] Granade C and Ferrie C 2012 *QInfer: Library for Statistical Inference in Quantum Information* (<http://github.com/csferrie/python-qinfer>)
- [29] Holevo A S 1982 *Probabilistic and Statistical Aspects of Quantum Theory* (Amsterdam: North-Holland)
- [30] Życzkowski K, Penson K A, Nechita Ion and Collins Benoit 2011 Generating random density matrices *J. Math. Phys.* **52** 062201
- [31] Henk M 2012 Löwner-John ellipsoids *Documenta Mathematica, Extra Volume: Optimization Stories* **95** 95–106