Experiment Design for Minimax Fidelity Estimation

Technical Manuscript

Jacob Spainhour
Department of Applied Mathematics
University of Colorado Boulder
Boulder, CO, USA
jacob.spainhour@colorado.edu

Akshay Seshadri
Department of Physics
University of Colorado Boulder
Boulder, CO, USA
akshay.seshadri@colorado.edu

Stephen Becker
Department of Applied Mathematics
University of Colorado Boulder
Boulder, CO, USA
stephen.becker@colorado.edu

Abstract

Technological limitations require that additional resources be spent to verify the output of a quantum device. We consider this verification through the lens of fidelity estimation, in which measurements of the quantum state directly inform how "close" a constructed state is to the intended target. This is in contrast to tomography schemes that first reconstruct the complete state, as these often require a greater number of measurements to obtain a reasonably accurate estimate. To be experimentally viable, a central goal of any approach is to accurately estimate the fidelity from as few observations, and types of observations, as possible. We present a technique that designs an experimental measurement protocol of a known target state, finding one that minimizes the width of a nearly optimal minimax confidence interval around the true value of the fidelity. Importantly, the nature of the underlying fidelity estimation scheme means that this design procedure is robust to the availability of measurements, and can be designed prior to the collection of any observations.

I Introduction

The verification and validation of existing quantum technologies is a critical component of the current Noisy Intermediate Scale Quantum era. Even when there is advance knowledge of properties of an experimentally acquired quantum state, there are a number of reasons that the true state might deviate from expectation. In this work, we are interested in the case where the state is created with a pure target state in mind, but the construction process is subject to noise and other experimentally induced errors [Cer+20]. In particular, the methods discussed are indifferent to the origins of these errors and the type of device hardware more generally.

Our objective is to design a system for accurately measuring the *fidelity* between an experimentally produced quantum state σ and the original target ρ. A common approach to this problem is full scale tomography, in which the complete state is reconstructed through a classical post-measurement analysis of observables, such as Maximum Likelihood Estimation (MLE) [Gro+10]. Once this state is constructed, it allows for analysis of many observables of interest, including the fidelity. However, this comes at an often unreasonable experimental cost, requiring a large number of measurements needed to obtain a reasonably accurate reconstruction, particularly when the fidelity is the only value of interest [Blu10]. Instead, it is preferable to have a method of directly measuring the fidelity without appealing to a full reconstruction. This has led to the Direct Fidelity Estimation (DFE) technique, in which many fewer measurements of a pre-determined experimental protocol can be used to rigorously compute a confidence interval on the true fidelity without needing to perform the full tomography [FL11]. A recent alternative to DFE is proposed by [Ses+21b], describing a theoretical framework for fidelity estimation that is more flexible to a wider range of experimental protocols. This technique produces both a single estimator of fidelity (itself a function of arbitrary observation data), as well as a rigorous, near-optimal bound on the error in the worst case of σ. This defines a near-optimal minimax confidence interval on the true fidelity, and we refer to this process for a single experimental protocol as the optimal minimax method.

Within the framework of the optimal minimax method for fidelity estimation, our work considers the width of this confidence interval, the *risk*, as a function of the experimental protocol. Our principle contribution is to design a measurement scheme that minimizes this risk by more efficiently allocating experimental resources, measured through the copies of the experimental state used for observation. This means we perform the minimization within a fixed budget for the total number measurements, recognizing that acquiring copies of a state and performing measurements is expensive. Because the optimal minimax method minimizes the error in the worst case, it constructs the relevant statistical estimator of the fidelity prior to execution of the experimental protocol. As a result, this means the entire optimization process does not need access to any observed data. This is very valuable from a practical perspective, as this means that the same optimal measurement scheme can be used across multiple experimental configurations, provided the intended target state remains the same. Ultimately, the proposed method has the potential to greatly reduce the cost of verifying the creation of quantum states while simultaneously improving the accuracy of estimates for the fidelity.

II Methods Consider a quantum system with a d-dimensional Hilbert space over $\mathbb{C}$. We identify each state in this system with a density matrix, a $d \times d$ positive semidefinite matrix with unit trace, and we denote the set of such density matrices as $\mathcal{X}$. Given a pure target state ρ and the experimentally constructed state σ, the fidelity between the two is given by the function $F(\sigma; \rho) = \text{Tr}(\sigma\rho)$, a linear function of σ. Because this fidelity is dependent on the unknown state, we can only estimate it based on a set of observations from a particular measurement protocol. Each of the L measurement settings of the protocol is described by a positive operator-valued measure (POVM) $\{E_1^{(\ell)}, \dots E_{N_\ell}^{(\ell)}\}$, such that the ℓ^{th} POVM can have one of N_ℓ possible outcomes $k \in \{1, \dots, N_\ell\}$. Each measurement, or *shot*, is repeated S_ℓ times to produce independently and identically distributed outcomes $\{o_1^{(\ell)}, \dots, o_{S_\ell}^{(\ell)}\}$. For a given state σ, the probabilities of each outcome are modeled by Born's rule with $p_\sigma^{(\ell)}(k) = \text{Tr}(E_k^{(\ell)}\sigma)$. In the following examples, we consider only the set of $4^d - 1$ Pauli measurements with $N_\ell = 2^d$ outcomes each, but the general framework is readily applicable to other types of measurements.

Given this information about the measurement protocol, as well as a confidence level of $1 - \delta$, the optimal minimax estimator $\widehat{F}$ is constructed through the optimization techniques described in [Ses+21a]. This estimator is an affine function of the observed outcomes from each measurement, mapping them to an estimate of the true fidelity F. The error in this estimate is defined within the supporting statistics literature [JN09] as the risk $\widehat{\mathcal{R}}$, essentially the width of the confidence interval around the estimate that contains F with a probability of $1 - \delta$. As is the case for the estimator $\widehat{F}$, this risk $\widehat{\mathcal{R}}$ can be computed prior to taking any measurements of σ. The risk can also be computed independently of the estimator, being the optimal value of the saddle point optimization problem

$$\widehat{\mathcal{R}} = \inf_{\alpha > 0} \left\{ \alpha \ln(2/\delta) + \max_{\chi_1, \chi_2 \in \mathcal{X}} \left[\frac{1}{2}\langle \rho, \chi_1 \rangle - \frac{1}{2}\langle \rho, \chi_2 \rangle + \alpha \ln(\text{AffH}(A(\chi_1), A(\chi_2))) \right] \right\}, \tag{1}$$

where $\text{AffH}(A(\chi_1), A(\chi_2))$ is the Hellinger affinity, a jointly log-concave function defined by

$$\ln\left(\text{AffH}(A(\chi_1), A(\chi_2))\right) = \sum_{\ell=1}^{L} \frac{S_\ell}{2} \ln \left[\sum_{k=1}^{N_\ell} p_{\chi_1}^{(\ell)}(k) p_{\chi_2}^{(\ell)}(k) \right]. \tag{2}$$

Written in this way, Equation 1 can be solved reasonably efficiently for a given experimental protocol using standard convex optimization software such as `cvxpy` (See [Ses+21b], [Ses+21a] for additional details). In fact, it is this feature of the optimal minimax method that permits optimization over the measurement protocol itself.

Our goal is to decide how to allocate a fixed budget of S_{tot} observations, measured in shots, over a given fixed set of L available measurement settings. The criterion we choose is the risk from Equation 1, which is a function of the vector of shot counts $\vec{S} = (S_1, \dots, S_L)^T$, which we minimize subject to the budget constraints:

$$\begin{aligned} &\text{minimize} && \widehat{\mathcal{R}}(\vec{S}), \\ &\text{subject to} && \sum_{\ell=1}^{L} S_\ell \leq S_{\text{tot}}. \end{aligned} \tag{3}$$

In minimizing the risk, we first note that although the S_ℓ are in principle integer-valued, with each representing discrete measurements of the experimental state, the objective function $\widehat{\mathcal{R}}$ allows for a real-valued treatment of these arguments. This means that during minimization of the risk we can allow each S_ℓ to take non-integer values, and simply round them at the end. This does not guarantee a global minimizer, but finding a true global minimizer is an intractable combinatorial problem. Instead, we simply search the space of possible measurements for a *better* protocol than conventional alternatives. In particular, we can verify that we improve on the initial choice of $\vec{S}$.

Even without the integer constraints, the function $\widehat{\mathcal{R}}$ is generally nonconvex, making the nested optimization problem in Equation 3 difficult. We solve this outer level optimization using the `scipy.optimize` library, utilizing the constrained trust-region methods within since they are appropriate for multidimensional nonconvex problems. We further improve the performance of these methods by providing an analytic gradient of $\widehat{\mathcal{R}}$, which is

$$\nabla_{\vec{S}} \widehat{\mathcal{R}}_\ell(\vec{S}) = \alpha^* \ln \left(\sum_{k=1}^{N_\ell} \sqrt{p_{\chi_1^*}^{(\ell)}(k) p_{\chi_2^*}^{(\ell)}(k)} \right),$$

where α^*, χ_1^*, and χ_2^* are the optimal values of the saddle point problem in Equation 1 at the input vector $\vec{S}$. Importantly, if the objective function has already been evaluated at this input vector, then constructing the gradient comes at essentially no additional cost since α^*, χ_1^*, and χ_2^* are already known. Through this optimization procedure, we find an optimal experimental protocol for a given target state ρ.

GHZ State ρ. $S_{\text{tot}} = 320$

Measurement	Uniform Shots	DFE-Derived Shots	Optimized Shots
XXX	11	40	53
XYY	11	40	53
YXY	11	40	53
YYX	11	40	53
ZZZ	11	160	106
Each Other Measurement	11	0	0
Risk $\widehat{R}$:	0.2237	0.1191	0.1136

W State ρ. $S_{\text{tot}} = 320$

Measurement	Uniform Shots	DFE-Derived Shots	Optimized Shots
XXX	11	0	41
XXZ	11	36	31
XZX	11	36	31
YYY	11	0	41
YYZ	11	36	31
YZY	11	36	31
ZXX	11	36	31
ZYY	11	36	31
ZZZ	11	107	53
Each Other Measurement	11	0	0
Risk $\widehat{R}$:	0.2269	0.1751	0.1456

TABLE I: Our optimized measurement scheme provides a more efficient utilization of available experimental resources, providing a lower value for the risk using the same volume of measurements.

III Results and Discussion Given a set of available measurement settings, our method for optimal experiment design is capable of producing tighter confidence intervals than conventional alternatives. As a naive approach, we uniformly distribute the budget among all available non-trivial measurements. To compare to a more sophisticated technique, we also consider a practical variant of the well-studied DFE method presented in [FL11] in which a subset of Pauli observables $\{W_i\}_{i=1}^{4^d}$ (excluding the identity) are selected at random according to the weighting function $Pr(i) = \text{Tr}(\rho W_i)^2/2^d$. To make an appropriate comparison, we consider the common approach of designing an experimental protocol that distributes the measurement budget among settings according to their relative weighting in the importance sampling rule. Then, instead of computing the risk according to the theory of DFE, we construct the optimal minimax confidence interval using this DFE-derived protocol. This provides a uniform comparison between all three experimental protocols, establishing the same type of worst case bound for each method. This approach is quite different from the intended usage of DFE, which is inherently probabilistic with respect to which measurements are taken as opposed to establishing a fixed measurement protocol. See [Zha+21] for an example of a similar usage of a fixed DFE-derived protocol used for fidelity estimation, where the importance weighting function determines which measurements are taken.

This comparison to DFE illustrates another advantage of the optimal minimax method for fidelity estimation. The theory for DFE associates each Pauli operator with a "coarse," 2-outcome POVM for each measurement setting, projecting the state onto the subspace of each operator associated with the eigenvalues ± 1. In contrast, by using the optimal minimax method to construct our estimator, we are able to fully utilize a set of "fine" POVMs for each measurement setting, which produce 2^d outcomes each. This is information which is commonly available experimentally, but often discarded due to a lack of appropriate theory. From these observations, that we also note that using an optimized scheme of this kind becomes more advantageous over the DFE-derived protocol as the problem scale increases. This is because as the size of the system increases, the fine POVM becomes more expressive relative to the coarser alternative. Furthermore, the use of a fine POVM for each measurement naturally induces sparsity in the number of different of measurements compared to the classical DFE technique. For example, utilizing this full collection of outcomes in the calculation of the risk allows us to completely eliminate redundancies between corresponding I and Z components of the Pauli operator. This is in stark contrast to the standard implementation of the DFE method, which uses coarser POVMs that are not invariant to this exchange.

To demonstrate our experiment design approach fully, we consider an optimized protocol on two 3-qubit examples, one with a GHZ stabilizer state and the other a W state. In both cases we consider the risk associated with a confidence level of 0.95 and use a budget of $S_{\text{tot}} = 320$ shots. We compile these results in Table I, showing the improvement to our risk estimate using the optimized protocol over the DFE-derived protocol. Note that although the inherent simplicity of the GHZ leads to an only slightly improved risk, our use of the set of fine POVMs nonetheless causes both techniques to use dramatically fewer types of measurements than would otherwise be assumed by the coarse POVM DFE method. The complex nature of the W state more clearly demonstrates the advantages of the proposed method, achieving a narrower confidence interval with the same number of shots, while simultaneously revealing additional measurements to take that may not be suggested by any underlying theory.

The above numerical experiment shows improvement in terms of the minimax risk, which is a *bound* on the estimation error when applied to measured outcomes. In line with this, we can perform numerical simulations of

observed measurements and use them to evaluate the true, asymmetric confidence interval, thereby verifying that the estimator produced by the optimized minimax procedure improves along with the risk. These simulations are performed on a pure and *totally random* 3-qubit state to which varying levels of depolarizing noise are added. As before, we consider the risk associated with a confidence level of 0.95. For each level of depolarizing noise, we simulate 5000 sets of measured data, using a budget of $S_{\text{tot}} = 1250$ shots for each instance. This budget was selected to ensure that the risk is less than 0.10 in the optimized case. Being able to work with arbitrary pure quantum states is a unique advantage of the optimal minimax method, as many other techniques that involve quantum state verification or machine learning require a specific structure for the target state.

For each level of depolarizing noise in Figure 1, we plot a distribution of measured fidelities along with the true fidelity with the target state. For each experimental protocol (DFE-derived and optimized) we plot the optimal minimax confidence interval, as well as an empirically constructed confidence interval that is selected to have minimum width while capturing 95% of the estimates. We make particular note of the fact that because the optimal minimax risk is minimized for a given target state with no reference to observation data, the width of the minimax confidence interval is equal across the noise levels. Nevertheless, we observe in these histograms that the optimized experimental protocol produces not just improvements over the alternative in terms of risk, but also in the estimate of the fidelity upon being evaluated with observation data, with estimates more concentrated around the true value.

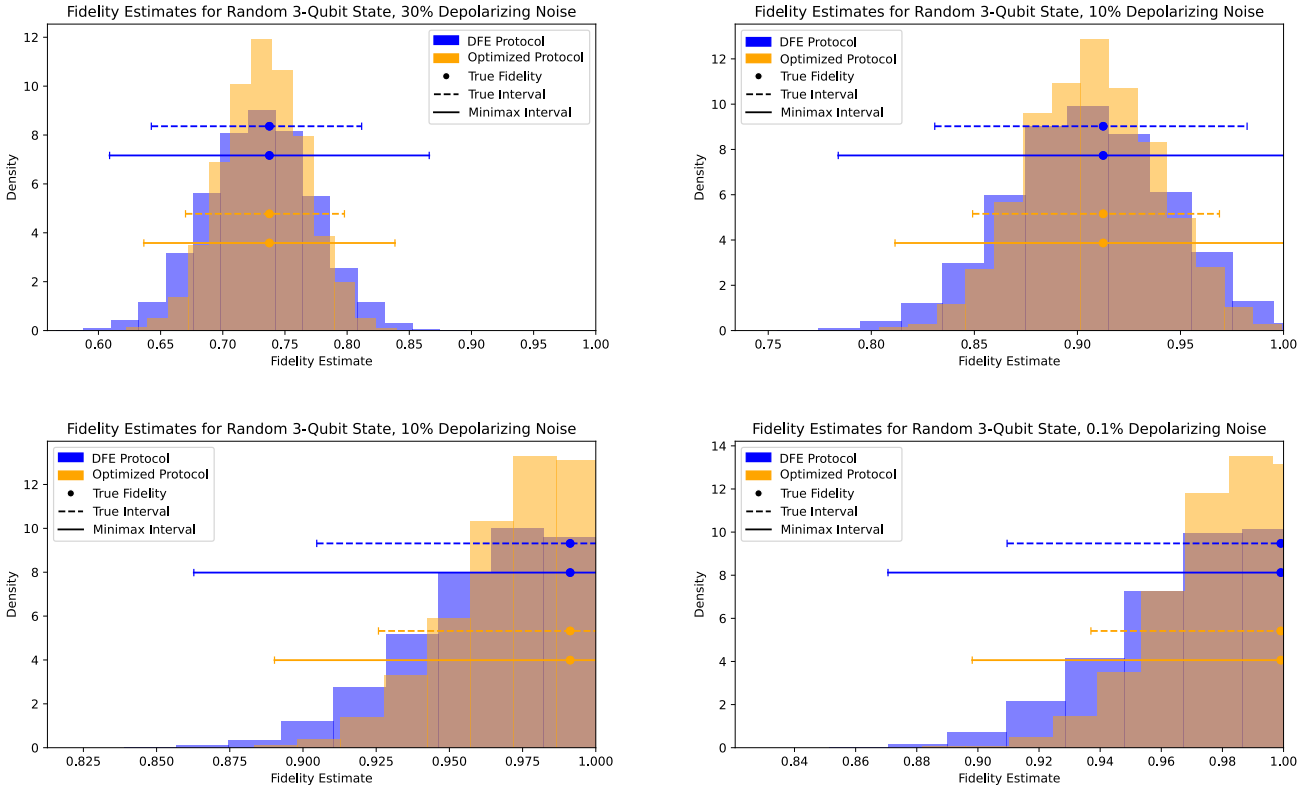

Fig. 1: For varying levels of depolarizing noise, over two types of experimental protocols (our new proposal in orange), we observe the relationship between the risk defined by the optimal minimax method and the empirically derived confidence interval derived simulated observations applied to each estimator. We see that in both cases, the optimized experimental protocol (orange) shows clear improvements across all noise levels.

IV Conclusions

The optimal minimax method is able to provide rigorous error bounds on the estimate of fidelity for a given measurement protocol. Overall, we have demonstrated that numerical optimization over the different available measurement protocols is an effective method of intelligently allocating experimental resources. We see that this approach not only reduces the minimax risk in quantum fidelity estimation, but also leads directly to improvements in the fidelity estimates themselves. The results of this work indicate that future advancements in the optimization of minimax bounds have the potential to greatly impact the field. As one example, measurement sparsity is hugely important in an experimental context, and should be prioritized even beyond the inherent sparsity obtained

from using fine POVM measurement settings. This is simply because switching measurements is a particularly expensive and time consuming part of the verification process. With the understanding that even small changes in the measurement protocol can *severely* improve the resulting fidelity estimator, modifications to the minimization problem in Equation 3 to directly encourage sparsity in the vector of shot counts have the potential to greatly improve efficiency in verifying modern quantum devices.

References

[Blu10] Robin Blume-Kohout. "Optimal, reliable estimation of quantum states". In: *New Journal of Physics* 12.4 (Apr. 2010), p. 043034. DOI: 10.1088/1367-2630/12/4/043034. URL: https://dx.doi.org/10.1088/1367-2630/12/4/043034.

[Cer+20] Marco Cerezo et al. "Variational Quantum Fidelity Estimation". In: *Quantum* 4 (Mar. 2020), p. 248. ISSN: 2521-327X. DOI: 10.22331/q-2020-03-26-248. URL: http://dx.doi.org/10.22331/q-2020-03-26-248.

[FL11] Steven T. Flammia and Yi-Kai Liu. "Direct Fidelity Estimation from Few Pauli Measurements". In: *Physical Review Letters* 106.23 (June 2011). DOI: 10.1103/physrevlett.106.230501. URL: https://doi.org/10.1103%2Fphysrevlett.106.230501.

[Gro+10] David Gross et al. "Quantum State Tomography via Compressed Sensing". In: *Phys. Rev. Lett.* 105 (15 Oct. 2010), p. 150401. DOI: 10.1103/PhysRevLett.105.150401. URL: https://link.aps.org/doi/10.1103/PhysRevLett.105.150401.

[JN09] Anatoli B Juditsky and Arkadi S Nemirovski. "Nonparametric estimation by convex programming". In: (2009).

[Ses+21a] Akshay Seshadri et al. *Theory of versatile fidelity estimation with confidence*. 2021. arXiv: 2112.07947 [quant-ph].

[Ses+21b] Akshay Seshadri et al. *Versatile fidelity estimation with confidence*. 2021. arXiv: 2112.07925 [quant-ph].

[Zha+21] Xiaoqian Zhang et al. "Direct Fidelity Estimation of Quantum States Using Machine Learning". In: *Phys. Rev. Lett.* 127 (13 Sept. 2021), p. 130503. DOI: 10.1103/PhysRevLett.127.130503. URL: https://link.aps.org/doi/10.1103/PhysRevLett.127.130503.