

Compressive Gate Set Tomography


Raphael Brieger^{1,*}, Ingo Roth^{2,3} and Martin Kliesch^{1,4,†}

¹*Institute for Theoretical Physics, Heinrich Heine University Düsseldorf, Germany*

²*Quantum Research Centre, Technology Innovation Institute, Abu Dhabi, UAE*

³*Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, Germany*

⁴*Institute for Quantum-Inspired and Quantum Optimization, Hamburg University of Technology, Germany*

 (Received 25 May 2022; revised 10 January 2023; accepted 25 January 2023; published 10 March 2023)

Flexible characterization techniques that provide a detailed picture of the experimental imperfections under realistic assumptions are crucial to gain actionable advice in the development of quantum computers. Gate set tomography self-consistently extracts a complete tomographic description of the implementation of an entire set of quantum gates, as well as the initial state and measurement, from experimental data. It has become a standard tool for this task but comes with high requirements on the number of sequences and their design, making it already experimentally challenging for only two qubits. In this work, we show that low-rank approximations of gate sets can be obtained from significantly fewer gate sequences and that it is sufficient to draw them at random. This coherent noise characterization however still contains the crucial information for improving the implementation. To this end, we formulate the data processing problem of gate set tomography as a rank-constrained tensor completion problem. We provide an algorithm to solve this problem while respecting the usual positivity and normalization constraints of quantum mechanics. For this purpose, we combine methods from Riemannian optimization and machine learning and develop a saddle-free second-order geometrical optimization method on the complex Stiefel manifold. Besides the reduction in sequences, we numerically demonstrate that the algorithm does not rely on structured gate sets or an elaborate circuit design to robustly perform gate set tomography. Therefore, it is more flexible than traditional approaches. We also demonstrate how coherent errors in shadow estimation protocols can be mitigated using estimates from gate set tomography.

DOI: [10.1103/PRXQuantum.4.010325](https://doi.org/10.1103/PRXQuantum.4.010325)

I. INTRODUCTION

The precise characterization of digital quantum devices is crucial for several reasons: (i) to obtain “actionable advice” on how imperfections on their implementation can be reduced, e.g., by experimental control, (ii) to tailor applications to unavoidable device errors so that their effect can be mitigated, and (iii) to benchmark the devices for the comparison of different physical platforms and implementations. There is already a wide variety of protocols to characterize components of a digital quantum computing device with a trade-off between the information gained about the system and the associated resource requirements and assumptions of the scheme [1,2].

One particularly important requirement for practical characterization protocols for quantum gates is their robustness against errors in the state preparation and measurement (SPAM). There are two general approaches that SPAM-robustly characterize the implementation of entire *gate sets* of a quantum computer. On the low complexity side there is randomized benchmarking (RB) [3–5] and variants thereof [6] that typically aim at determining a single measure of quality for an experiment, though with the exception of RB tomography protocols [7–10]. Yet for the targeted improvement of individual quantum operations, protocols that provide more detailed information beyond mere benchmarking are crucial.

This is the motivation of self-consistent gate set tomography (GST) [11–18]. GST estimates virtually all parameters describing a noisy implementation of a quantum computing device simultaneously from the measurements of many gate sequences. This comprises tomographic estimates for all channels implementing the gate set elements, the initial state(s), and the measurement(s). The full tomographic information can then be used to compute arbitrary error measures for verification and to provide advice on

*brieger@hhu.de

†martin.kliesch@tuhh.de

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

error mitigation and device calibration [15,19–23]. Concomitant with the massive amount of inferred information and minimal assumptions, these protocols come with enormous resource requirements in terms of the necessary number of measurement rounds and the time and storage consumption of the classical postprocessing. Standard GST, as described by, e.g., Nielsen *et al.* [18], uses many carefully designed gate sequences in the experiment and a sophisticated and challenging data processing pipeline in post. To arrive at physically interpretable estimates, i.e., completely positive and trace preserving (CPT) maps, additional postprocessing is required. The massive amount of specific data consumed by standard GST already limits its practical applicability for two-qubit gate sets. The focus of Nielsen *et al.* [18] and their implementation pyGSTi is on so-called *long sequence GST*, a method to improve an initial GST estimate by using gate sequences in which a building block is repeated many times. The resulting error amplification of the building block is then used to significantly improve the accuracy of the GST estimate at the cost of larger measurement effort. In our work we focus on *short sequences* and the problem of finding an initial estimate without assuming any prior knowledge on the gate set and minimal experimental requirements.

The most important diagnostic information for a quantum computing device is often already contained in a low-rank approximation of the processes, states, and measurements. Coherent errors are typically those that can be corrected by experimental control and are of interest for refining calibration models. The strength of incoherent noise on the other hand is arguably well captured by average error measures as provided by RB outputs. Moreover, current fault tolerance thresholds often rely on worst-case error measures for which no good direct estimation technique exists [24–27] and coherent errors in particular hinder their indirect inference from average error measures [28–30]. For standard state and process tomography, it was realized that low-rank assumptions can crucially reduce the sample complexity, the required number of measurements, and the postprocessing complexity [9,31–40] as well as improve the stability against imperfections in the measurements [41] using compressed sensing techniques [42–44].

In this work, we take a fresh look at the data processing problem of GST from a compressed sensing perspective and regard it as a highly structured tensor completion problem. We develop a reconstruction method, called mGST, that exploits the geometric structure of CPT maps with low Kraus ranks. In numerical simulations we demonstrate that our structure-exploiting mGST approach (i) allows for maximal flexibility in the design of gate sequences, so that standard GST gate sequences and random sequences work equally well, and (ii) obtains low-rank approximations of the implemented gate set from a significantly reduced number of sequences and samples. This allows us

to successfully perform GST with gate sets and sequences that are not amenable to the standard GST implementation pyGSTi [17,45]. As one example, while the sequence design of pyGSTi uses at least 907 specific sequences to reconstruct a two-qubit gate set, we numerically demonstrate low-rank reconstruction from 200 random sequences of maximal length $\ell = 7$ with runtimes of less than an hour on a standard desktop computer. Thus, compressive GST significantly lowers the experimental resource requirements for maybe the most prominent use cases of GST, making it a tool that can be more easily and routinely applied. At the same time, for the default gate sets and sequences from the standard GST implementation, the novel algorithm matches state-of-the-art results. The runtime and storage requirements of mGST still scale exponentially in the number of qubits, as does the amount of parameters of the gate set it identifies. This limits the feasibility of the classical postprocessing of compressive GST to gate sets acting on only a few qubits without further assumptions. Nonetheless, we demonstrate that coherent errors and depolarizing noise of a three-qubit gate set can be completely characterized, from as little as 128 sequences of length $\ell = 7$ on desktop hardware in a few hours.

To give a novel example of how information about coherent errors can be used, we simulate a ten-qubit system and perform GST on neighboring two-qubit pairs. The resulting gate set estimates then allow us to calibrate the postprocessing step of the shadow estimation protocol [46], which is widely used for the sample efficient estimation of observables. More concretely, we find in simulations with moderate coherent errors that shadow estimates of the ground-state energy of a ten-qubit Heisenberg Hamiltonian are heavily biased when knowledge of the noisy gate implementation is limited. Information from GST on two-qubit pairs allows us to reduce this bias by about an order of magnitude.

Our mGST reconstruction method relies on manifold optimization over complex Stiefel manifolds [47–53] in order to include the low-rank CPT constraints. Such constraints emerge in several optimization problems [54–57] with applications in machine learning [52,58], quantum chemistry [59], signal processing in wireless communication [60,61], and more recently in the quantum information literature; see e.g., Refs. [62–64]. In order to deal with the nonconvex optimization landscape, we adopt a second-order saddle-free Newton method [65] in this setting. This involves the derivation of an analytic expression for geodesics, as well as an expression for the Riemannian Hessian in the respective product manifolds. Another important motivation for phrasing GST as a randomly subsampled tensor completion problem is to bring it closer to potential analytical recovery guarantees common for related tensor completion problems [66–73], opening up a new research direction.

Finally, being able to perform GST from random sequences enables one to use the same type of data for different, increasingly refined characterization tasks from filtered RB [6], cross-entropy benchmarking [74], and RB tomography [7–9] to GST. Unifying these approaches, random gate sequences can be regarded as the “classical shadow” of a gate set from which many properties can be estimated efficiently [75]. Compressive gate set tomography provides more detailed diagnostic information and only requires to further increase the amount of data without changing the experimental instructions.

With randomized linear GST [76] and fast Bayesian tomography [77], related alternatives to tackle the GST data processing problem have been proposed. Here, the gates are assumed to be well approximated by an *a priori* known unitary followed by a noise channel that is either linearized around the identity [76] or around a prior noise estimate [77]. This allows for treatment of the outcome probabilities as approximately linear functions. The resulting scheme already works for random sequence data but comes at the expense of much stronger assumptions compared to the compressed sensing approach taken with mGST.

The rest of the paper is structured into three parts. In the following section, we formalize the data processing problem of GST as a constrained reconstruction problem. In Sec. III, we formulate the data processing problem as a geometric optimization task and derive the mGST algorithm. In Sec. IV we demonstrate the performance of the novel algorithm in numerical simulations and compare our results with the standard GST processing pipeline of pyGSTi.

II. THE DATA PROCESSING PROBLEM OF GATE SET TOMOGRAPHY

In GST a quantum computing device is modeled as follows. The device is initialized with a state $\rho \in \mathcal{S} := \{\sigma \in \mathcal{H} : \sigma \geq 0, \text{Tr}[\sigma] = 1\}$ on a finite-dimensional Hilbert space $\mathcal{H} = \mathbb{C}^d$. Subsequently, a sequence of noisy operations from a fixed gate set $(\mathcal{G}_i)_{i \in [n]}$ can be applied, where we use the notation $[n] := \{1, 2, 3, \dots, n\}$. The noisy operations $\mathcal{G}_i : L(\mathcal{H}) \rightarrow L(\mathcal{H})$ are CPT maps on $L(\mathcal{H})$, the set of linear operators on \mathcal{H} . We define $[n]_\ell^* := \bigcup_{k=0}^{\ell} [n]^k$ such that $\mathbf{i} \in [n]_\ell^*$ defines a gate sequences with length of at most ℓ and associated CPT map $\mathcal{G}_\mathbf{i} := \mathcal{G}_{i_\ell} \circ \dots \circ \mathcal{G}_{i_1}$, the concatenation of the gates in sequence \mathbf{i} . In the end, a measurement is performed described by a positive operator-valued measure (POVM) with elements $(E_j)_{j \in [n_E]}$, satisfying $\sum_j E_j = \mathbb{1}$ and $0 \leq E_j \leq \mathbb{1}$ for all $j \in [n_E]$. The full description of the noisy quantum computing device is thus given by the triple

$$\mathcal{X} = ((E_j)_{j \in [n_E]}, (\mathcal{G}_i)_{i \in [n]}, \rho) \quad (1)$$

of a quantum state, a physical gate set, and a POVM. Let $I \subseteq [n]_\ell^*$ be the set of accessible gate sequences with $n_{\text{seq}} := |I|$ denoting the number of sequences. The probability of measuring outcome j upon applying a gate sequence $\mathbf{i} \in I$ is

$$p_{j|\mathbf{i}}(\mathcal{X}) = \text{Tr}[E_j \mathcal{G}_\mathbf{i}(\rho)]. \quad (2)$$

By $(\mathbf{p}_\mathbf{i}(\mathcal{X}))_j := p_{j|\mathbf{i}}(\mathcal{X})$ we denote the corresponding vector and, moreover, often omit the argument \mathcal{X} . While being a fairly general description, this gate set model relies on the following assumptions.

- (i) The physical system needs to be well characterized by a Hilbert space of fixed dimension.
- (ii) The system parameters need to be time independent over different experiments.
- (iii) A gate’s action is independent of the gates applied before and after (Markovianity).

There exist multiple descriptions of quantum computing devices within the gate set model that yield the same measurement probabilities on all sequences. Below, we provide a more detailed description of this freedom in terms of *gauge transformations*. These are linear transformations that, when simultaneously applied to all gates, input state, and POVM elements, leave the measurement statistics (2) invariant.

The task of GST is to infer the device’s full description (1) from measured data. To this end, one estimates the output probabilities for a set of different sequences $I \subset [n]_\ell^*$ by repeatedly performing the measurements of the corresponding sequences. Thus, we can state the *data processing problem of GST* as follows.

Problem (GST data processing): *Let \mathcal{X} be a gate set and $I \subset [n]_\ell^*$ a set of sequences. Given empirical estimates $\{y_{j|\mathbf{i}}\}_{\mathbf{i} \in I, j \in [n_E]}$ of $\{p_{j|\mathbf{i}}(\mathcal{X})\}_{\mathbf{i} \in I, j \in [n_E]}$, find the device description $\mathcal{X} = ((E_j)_{j \in [n_E]}, (\mathcal{G}_i)_{i \in [n]}, \rho)$ up to the gauge freedom.*

Note that GST aims at solving an *identification problem*. That is, for sufficient data, find the unique device description of the device compatible with the data. In particular, the input data $\{\hat{p}_{j|\mathbf{i}}\}_{\mathbf{i} \in I, j}$ is required to uniquely single out the device description. This is related but distinct from the corresponding *learning task* to find a description that generalizes on unseen data.

A. Compressive gate set description

At the heart of our approach is to capture this data processing problem as a highly structured *tensor completion problem*. The structure allows us to reduce the required size and structural assumption of set I , in order to determine \mathcal{X} . It is instructive to visualize the problem with

tensor network diagrams. The gate set can be viewed as a tensor of five indices and the action of gate i on the initial state ρ can be visualized as

$$\mathcal{G}_i(\rho) = \begin{array}{c} i \\ \text{---} \text{---} \text{---} \text{---} \text{---} \\ \boxed{\mathcal{G}} \text{---} \boxed{\rho} \\ \text{---} \text{---} \text{---} \text{---} \end{array}, \quad (3)$$

where each leg represents an open index and the joining of legs represents summation of the corresponding indices; see, e.g., Chapter 5.1 of Ref. [78] for more information on the tensor network notation. Neglecting the finite statistics in estimating the probabilities, the GST data processing problem can thus be rephrased as the problem of completing the translation-invariant *matrix product state* (MPS) [79–81] or *tensor train* [82]

$$p_{j|i} = \text{Tr}[E_j \mathcal{G}_{i_l} \circ \dots \circ \mathcal{G}_{i_2} \circ \mathcal{G}_{i_1}(\rho)]$$

$$= \begin{array}{c} j \\ \text{---} \text{---} \text{---} \text{---} \text{---} \\ \boxed{E} \text{---} \boxed{\mathcal{G}} \text{---} \dots \text{---} \boxed{\mathcal{G}} \text{---} \boxed{\mathcal{G}} \text{---} \boxed{\rho} \\ \text{---} \text{---} \text{---} \text{---} \end{array} \quad (4)$$

from access to a couple of its entries. By the following assumptions one can introduce more structure. First, we assume that the elements of the device description \mathcal{X} satisfy the physicality assumptions regarding normalization and positivity. Second, the assumption that they have low-rank approximations yields additional compressibility of \mathcal{X} .

In more detail, physically implementable gate sets are completely positive and trace preserving if and only if they admit a Kraus decomposition [83], i.e., the i th gate implementation can be written as

$$\mathcal{G}_i(\rho) = \sum_{l=1}^{r_K} \mathcal{K}_{il} \rho \mathcal{K}_{il}^\dagger \quad (5)$$

for each $i \in [n]$, where r_K is the (maximum) Kraus rank of the CPT maps $\{\mathcal{G}_i\}$. We use the notation \mathcal{K} to denote the tensor containing all Kraus operators of all gates and \mathcal{K}_i to denote the tensor containing the Kraus operators for gate i . In terms of tensor network diagrams the decomposition is represented as

$$\begin{array}{c} i \\ \text{---} \text{---} \text{---} \text{---} \\ \boxed{\mathcal{G}} \\ \text{---} \text{---} \text{---} \end{array} = \begin{array}{c} \text{---} \text{---} \text{---} \text{---} \\ \boxed{\mathcal{K}} \\ \text{---} \text{---} \text{---} \text{---} \\ \boxed{\mathcal{K}^*} \\ \text{---} \text{---} \text{---} \end{array}. \quad (6)$$

Moreover, the trace preservation constraints

$$\sum_{l=1}^{r_K} \mathcal{K}_{il}^\dagger \mathcal{K}_{il} = \mathbb{1} \quad \text{for all } i \quad (7)$$

require \mathcal{K}_i viewed as a matrix in $\mathbb{C}^{r_K d \times d}$ to be an isometry, i.e.,

$$\begin{array}{c} \boxed{\mathcal{K}_i} \\ \text{---} \text{---} \text{---} \text{---} \\ \boxed{\mathcal{K}_i^*} \\ \text{---} \text{---} \text{---} \end{array} = \begin{array}{c} \text{---} \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \end{array}. \quad (8)$$

Constraints on a low Kraus rank r_K can be naturally enforced in this parametrization by reducing the row dimension of \mathcal{K}_i . The initial state and POVM elements are constrained to be positive matrices, which we hence parameterize as

$$E_j = A_j^\dagger A_j, \quad \rho = B B^\dagger \quad (9)$$

with $A_j \in \mathbb{C}^{r_E \times d}$ and $B \in \mathbb{C}^{d \times r_\rho}$, where r_E and r_ρ are the matrix ranks. For the matrices A_j to form a valid POVM, they have to satisfy a similar condition to the Kraus operators,

$$\sum_{j=1}^{r_E} A_j^\dagger A_j = \mathbb{1}, \quad (10)$$

while the initial state is of unit trace if

$$\|B\|_F = 1. \quad (11)$$

With the physicality constraints incorporated, measurement outcome probabilities are given in terms of tensor network diagrams as

$$p_{j|i} = \text{Tr}[E_j \mathcal{G}_{i_l} \circ \dots \circ \mathcal{G}_{i_2} \circ \mathcal{G}_{i_1}(\rho)]$$

$$= \begin{array}{c} j \\ \text{---} \text{---} \text{---} \text{---} \text{---} \\ \boxed{E} \text{---} \boxed{\mathcal{G}} \text{---} \dots \text{---} \boxed{\mathcal{G}} \text{---} \boxed{\mathcal{G}} \text{---} \boxed{\rho} \\ \text{---} \text{---} \text{---} \text{---} \end{array}$$

$$= \begin{array}{c} \text{---} \text{---} \text{---} \text{---} \text{---} \\ \boxed{A} \text{---} \boxed{\mathcal{K}} \\ \text{---} \text{---} \text{---} \text{---} \text{---} \\ \boxed{A^*} \text{---} \boxed{\mathcal{K}^*} \\ \text{---} \text{---} \text{---} \text{---} \end{array} \dots \begin{array}{c} \text{---} \text{---} \text{---} \text{---} \\ \boxed{\mathcal{K}} \text{---} \boxed{\mathcal{K}} \\ \text{---} \text{---} \text{---} \text{---} \\ \boxed{\mathcal{K}^*} \text{---} \boxed{\mathcal{K}^*} \\ \text{---} \text{---} \text{---} \end{array} \begin{array}{c} \text{---} \text{---} \text{---} \text{---} \\ \boxed{B} \\ \text{---} \text{---} \text{---} \text{---} \\ \boxed{B^*} \\ \text{---} \text{---} \text{---} \end{array}. \quad (12)$$

Thus, we arrive at a compressive device description $\mathcal{X}_c = (A, \mathcal{K}, B)$ that considerably reduces the amount of parameters compared to the triple $\mathcal{X} = ((E_j)_{j \in [n_E]}, (\mathcal{G}_i)_{i \in [n]}, \rho)$ when choosing small dimensions r_ρ , r_K , and r_E . Correspondingly, we can adapt the GST data processing task to demand only a compressive device description.

Problem (Compressive GST data processing): *Let \mathcal{X} be a gate set and $I \subset [n]_I^*$ a set of sequences. Given empirical estimates $\{y_{j|i}\}_{i \in I, j \in [n_E]}$ of $\{p_{j|i}(\mathcal{X})\}_{i \in I, j \in n_E}$ and ranks r_ρ , r_K , and r_E , find the compressive device description $\mathcal{X}_c = (A, \mathcal{K}, B)$ of dimensions r_ρ , r_K , and r_E , respectively, so that the normalization constraints (7), (10), and (11) are satisfied.*

As before, the set of sequences needs to be large enough so that this identification problem is well defined. Again, a desired compressive device description \mathcal{X}_c can only be determined up to gauge freedom. Note that, for the identification problem to be well defined, it is not required that the true gate set \mathcal{X} that generated the data is of low rank itself. As one usually aims to implement unit rank states, unitary gates, and basis measurements, i.e., for $r_\rho = r_K = r_E = 1$, it can be expected however that a compressive device description \mathcal{X}_c is also often a good approximation to the true gate set. Moreover, coherent errors are arguably the most relevant, since they give actionable advice on error mitigation and are complementary to the incoherent error measures provided by randomized benchmarking experiments. By choosing the ranks in \mathcal{X}_c , a problem-specific decision can be made that balances the information gained with the computational and sample complexity of the model reconstruction.

Since the $p_{j|i}(\mathcal{X})$ are high degree polynomials in the gate set parameters, the compressive GST data processing problem is different from compressed sensing for standard state and process tomography, where the map from the model parameters to the outcome probabilities is linear.

Next, we discuss another unique problem of GST, the gauge freedom, more explicitly and introduce relevant error measures for gates sets \mathcal{X} .

B. Gauge freedom and gate set metrics

So far we have not made explicit what “finding a device description” actually means. What is well studied in the GST and RB literature [6,14,16,84–86] is that, without additional prior assumptions, there is a freedom in representing a device in the gate set model. In particular, this freedom needs to be considered when defining a metric for gate sets [84] with respect to which we want to recover the device description.

Gauge freedom refers to the following observation. The observable measurement probabilities $p_{j|i}$ of the form (4)

are invariant under the transformation

$$\rho \mapsto \mathcal{T}^{-1}(\rho), \quad (13)$$

$$\mathcal{G}_i \mapsto \mathcal{T}^{-1} \circ \mathcal{G}_i \circ \mathcal{T} \quad \text{for all } i, \quad (14)$$

$$E_j \mapsto \mathcal{T}^\dagger(E_j) \quad \text{for all } j, \quad (15)$$

for any invertible super operator $\mathcal{T} : \mathcal{L}(\mathcal{H}) \rightarrow \mathcal{L}(\mathcal{H})$, where \mathcal{T}^\dagger denotes the adjoint of \mathcal{T} with respect to the Hilbert-Schmidt inner product. This invariance is also the well-known *gauge freedom of MPS* [87].

If the gate set is universal and the initial state is pure then the gauge transformation \mathcal{T} has to be either a unitary or antiunitary channel. This statement can be seen as follows.

In our case the \mathcal{G}_i are constrained to be CPT. Hence, $\mathcal{G} \mapsto \mathcal{T}^{-1}(\mathcal{G})\mathcal{T}$ has to map CPT maps to CPT maps. Similarly, $\mathcal{T}^{-1}\rho$ has to be a density operator and $\{\mathcal{T}^\dagger(E_j)\}_j$ a valid POVM.

A more explicit condition on \mathcal{T} can be obtained by considering gauge action on entire sequences. For all sequences \mathbf{i} , we have

$$p_{j|i} = \text{Tr}[E_j \mathcal{T} \circ \mathcal{T}^{-1} \circ \mathcal{G}_i(\rho)],$$

where $\mathcal{G}_i(\rho)$ is a positive operator if the gates \mathcal{G}_i are CPT. Now $\mathcal{T}^{-1}\mathcal{G}_i(\rho)$ has to be positive as well for all sequences \mathbf{i} . Thus, if the gate set is universal, the map \mathcal{T}^{-1} has to be positive and trace preserving for all states. An analogous statement can be made for \mathcal{T}^\dagger , by considering that $\mathcal{T}^\dagger \mathcal{G}_i^\dagger(E_j)$ has to be positive definite for all POVM elements E_j . This implies that \mathcal{T} has to be a positive map as well.

It has been shown that any positive invertible map \mathcal{T} with a positive inverse can be written either as $\mathcal{T}(\rho) = PU\rho U^\dagger P^\dagger$ or $\mathcal{T}(\rho) = PU\rho^T U^\dagger P^\dagger$ for $U \in \text{GL}(d, \mathbb{C})$ and $P \geq 0$ [88, Theorem 2]. The condition that \mathcal{T} needs to be trace preserving then yields $P = \mathbb{1}$, as can be seen from the Kraus decomposition. Hence, \mathcal{T} is indeed either a unitary or antiunitary channel.

We note that the map $\rho \mapsto U\rho^T U^\dagger$ is positive but not completely positive. However, it has the property that $\mathcal{T}^{-1}\mathcal{G}_i\mathcal{T}$ is CPT whenever \mathcal{G}_i is CPT. This can be seen by observing that the Choi matrix of $\mathcal{T}^{-1}\mathcal{G}_i\mathcal{T}$ is given by $(U^* \otimes U)\text{Choi}(\mathcal{G}_i)^T(U^T \otimes U^\dagger)$, which is positive definite for \mathcal{G}_i being CPT.

However, actual gate set implementations are noisy and hence not universal in the sense that they cannot prepare any pure state. Therefore, in practice, the gauge freedom can be larger [18]. For instance, if all gate implementations are given by unital channels then an additional freedom exists: depolarizing noise can be commuted through the circuit. Therefore, it can be distributed arbitrarily among the initial state, gates, and measurement. Meaningful distance measures for gate sets should have the same gauge freedom as the GST data. The problem of finding a gauge

invariant distance has been studied by Lin *et al.* [84]. For individual gate sequences, any measure that compares only the ideal and observed outcome probabilities is naturally gauge invariant. The authors thus proposed using the total variation error, a natural error measure to compare probability distributions, for individual gate sequences. Let

$$p_{j|i}(E_j, \mathcal{G}_i, \rho) := \text{Tr}[E_j \mathcal{G}_i(\rho)] \quad (16)$$

denote the probabilities of measuring the j th output of the POVM with elements E_j after applying the sequence \mathbf{i} of gates in \mathcal{G}_i to the state ρ . The *total variation error* for sequence \mathbf{i} between two gate sets $\hat{\mathcal{X}} = \{\{\hat{E}_j\}, \hat{\mathcal{G}}, \hat{\rho}\}$ and $\mathcal{X} = ((E_j), \mathcal{G}, \rho)$ is defined as

$$\delta d_i(\hat{\mathcal{X}}, \mathcal{X}) := \frac{1}{2} \sum_j |\text{Tr}[\hat{E}_j^\dagger \hat{\mathcal{G}}_i(\hat{\rho})] - \text{Tr}[E_j^\dagger \mathcal{G}_i(\rho)]|. \quad (17)$$

The *mean variation error* (MVE) is defined as [84]

$$\text{MVE}_I(\hat{\mathcal{X}}, \mathcal{X}) := \mathbb{E}_{\mathbf{i} \sim I}[\delta d_i(\hat{\mathcal{X}}, \mathcal{X})] \quad (18)$$

with respect to a set of sequences I , where $\mathbf{i} \sim I$ means that \mathbf{i} is drawn uniformly from I . Often, we omit the subscript I in the following. The MVE corresponds to taking the natural worst-case error measure over the measurement outcomes (the total variation distance) and averaging it over the available gate sequences. Often I is chosen as the set of all gate sequences up to some length ℓ . Then the expectation value (18) contains a sum over exponentially many terms. However, since they are all non-negative, the expectation value can be estimated efficiently via Monte Carlo sampling [84].

A closely related error measure is the *mean squared error* (MSE)

$$\mathcal{L}_I(\hat{\mathcal{X}}, \mathcal{X}) := \mathbb{E}_{\mathbf{i} \sim I} \sum_{j \in [n_E]} (\text{Tr}[\hat{E}_j^\dagger \hat{\mathcal{G}}_i(\hat{\rho})] - \text{Tr}[E_j^\dagger \mathcal{G}_i(\rho)])^2, \quad (19)$$

which averages the squared deviation over all sequences and POVM elements.

III. GST DATA PROCESSING VIA RIEMANNIAN OPTIMIZATION

In the previous section, we defined the compressive GST data processing problem and introduced metrics for the quality of reconstruction. We now turn to devising a concrete algorithm for the data processing problem. To this end, we formulate the reconstruction problem as a constraint optimization problem of a loss function for the data fitting. A natural candidate for the loss function is the MVE restricted to the set of measured sequences. As a proxy, we instead minimize the MSE that depends smoothly on the

gate set and is therefore more suitable for local optimization. In terms of the compressive device description, the MSE (19) can be written as

$$\mathcal{L}_I(A, \mathcal{K}, B|\mathbf{y}) := \frac{1}{|I|} \sum_{\mathbf{i} \in I} \sum_j (p_{j|i}(A, \mathcal{K}, B) - y_{j|i})^2, \quad (20)$$

where $y_{j|i}$ is the empirical estimate of $\text{Tr}[E_j \mathcal{G}_i(\rho)]$. Correspondingly, the compressive GST data processing problem can be cast as the constraint optimization problem

$$\underset{A, \mathcal{K}, B}{\text{minimize}} \quad \mathcal{L}_I(A, \mathcal{K}, B|\mathbf{y}) \quad (21a)$$

$$\text{subject to} \quad \sum_{l=1}^{r_K} \mathcal{K}_{il}^\dagger \mathcal{K}_{il} = \mathbb{1} \quad \text{for all } i \in [n], \quad (21b)$$

$$\sum_{j=1}^{r_E} A_j^\dagger A_j = \mathbb{1}, \quad (21c)$$

$$\|B\|_F = 1. \quad (21d)$$

The constraints restrict the objective variables to embedded matrix manifolds. Therefore, algorithms for the optimization problem can be derived by generalizing standard optimization algorithms for functions on the Euclidean space to the geometric structure of these manifolds.

A. The complex Stiefel manifold

In order to formulate our main reconstruction algorithm, we need to understand the matrix manifold that encompasses the physicality constraints mentioned in Sec. II A. We start by summarizing the elementary properties of these manifolds, to then derive a parametrization of geodesics and the Riemannian Hessian, thereby extending what was previously done for their real counterparts in Ref. [48]. For a comprehensive introduction to optimization on matrix manifolds, we refer the reader to the book by Absil *et al.* [89].

Let $(K_l)_{l \in [r]}$ be the Kraus operators of a fixed gate. By stacking them along their row dimension to a new matrix $K \in \mathbb{C}^{dr \times d}$, we can write the CPT constraint as $K^\dagger K = \mathbb{1}$. In the following we set $D = dr$. The set

$$\text{St}(D, d) := \{K \in \mathbb{C}^{D \times d} : K^\dagger K = \mathbb{1}_d\} \quad (22)$$

is called the $D \times d$ *complex Stiefel manifold*. This manifold is the set of isometries of the Euclidean space and contains the special cases of the sphere $\text{St}(D, 1)$ and the unitary matrices $U(D) = \text{St}(D, D)$. We regard it here as a submanifold of $\mathbb{C}^{dr \times d}$.

The tangent space of $\text{St}(D, d)$ at K is given by

$$T_K \text{St}(D, d) = \{\Delta \in \mathbb{C}^{D \times d} : K^\dagger \Delta = -\Delta^\dagger K\}. \quad (23)$$

The canonical inner product of $\Delta_1, \Delta_2 \in T_K \text{St}(D, d)$ can be defined as

$$\langle \Delta_1, \Delta_2 \rangle_K = \text{Re}\{\text{Tr}(\Delta_1^\dagger \Gamma \Delta_2)\} \quad (24)$$

with $\Gamma = \mathbb{1} - \frac{1}{2}KK^\dagger$. Another choice is the standard Hilbert-Schmidt inner product of the embedding matrix space. However, the advantage of the canonical inner product is that it weights all degrees of freedom on the tangent space equally. The Stiefel $\text{St}(D, d)$ together with the metric given by Eq. (24) is a Riemannian manifold. The *normal space* is defined by

$$\begin{aligned} N_K \text{St}(D, d) &= \{\Delta_\perp \in \mathbb{C}^{D \times d} : \langle \Delta, \Delta_\perp \rangle_K \\ &= 0 \text{ for all } \Delta \in T_K \text{St}(D, d)\}. \end{aligned}$$

The projector onto the normal space at position K is given by

$$P_N(X) = K(K^\dagger X + X^\dagger K)/2 \quad (25)$$

for $X \in \mathbb{C}^{d \times d}$ and we can write the projector onto the tangent space at K as

$$P_T(X) = X - P_N(X). \quad (26)$$

We wish to optimize the MSE over $\text{St}(D, d)$. In analogy to the optimization over $U(n)$ in Ref. [47], we move along geodesics, which are the locally length minimizing curves. In Appendix A 1 we show that, within $\text{St}(D, d)$, a geodesic starting at $K_{t=0} \equiv K$ and going in the direction $\Delta \in T_K \text{St}(D, d)$ can be written as

$$K_t(K, \Delta) = (K, Q) \exp \left[t \begin{pmatrix} K^\dagger \Delta & -R^\dagger \\ R & 0 \end{pmatrix} \right] \begin{pmatrix} \mathbb{1} \\ 0 \end{pmatrix} \quad (27)$$

with Q, R given by the QR decomposition of $(\mathbb{1} - KK^\dagger)\Delta$. Note that $\dot{K}_t|_{t=0} = \Delta$. Often simpler curves that just satisfy $K_0 = K$ and $\dot{K}_t|_{t=0} = \Delta$ are used instead of the geodesic in order to save computation time [89]. However, computing the exponential of the $2d$ -dimensional matrix in Eq. (27) provides no bottleneck in our scenario as the inversion of the $2nd^2r_K$ -dimensional Hessian is more costly (see Sec. IV F).

In order to identify the Riemannian gradient and Hessian, we generalize results from the real case [48] to the complex case. Then we use the second-order Taylor approximation of the objective function, which will be given below in terms of the MSE (19) along geodesics (see Appendix A 2). The same treatment can be applied to the

POVM given by the matrices A_j from decomposition (9), where we define A as the matrix obtained from stacking the A_j along their row dimension. The physicality constraint on A is then equivalent to $A \in \text{St}(dn_E, r_E)$ with n_E being the number of POVM elements and r_E their maximal rank. Finally, the constraint $\|B\|_F^2 = \text{vec}(B)^\dagger \text{vec}(B) = 1$ on the initial state (9) can also be captured by the Stiefel manifold via the requirement that $\text{vec}(B) \in \text{St}(dr_\rho, 1)$, where $\text{vec}(B) \in \mathbb{C}^{dr_\rho}$ is the vectorization of $B \in \mathbb{C}^{d \times r_\rho}$.

B. The mGST estimation algorithm

With a better understanding of the underlying manifold structure we can now formulate a concrete optimization approach to tackle the estimation problem (21). The least-squares cost function (19) is a polynomial of order at most the sequence length squared in the parameters of \mathcal{G} , with a highly degenerate global minimum due to the gauge freedom. In analogy to the alternating minimization techniques that are successful for matrix product state completion [66,90,91] we alternate between updates on A, \mathcal{K} , and B . Each update would naively be done via a local optimization approach such as gradient descent. However, we observe that following the gradient direction on the respective manifolds is problematic around saddle points, which are frequently encountered in our optimization problem. In principle, the gradient direction points away from saddle points, yet the norm of the gradient can be arbitrarily small. There are different approaches in the literature to deal with this problem. For instance, information about the curvature can be included [65] or, if a saddle point is encountered, random update directions can be chosen to escape the area of vanishing gradient [50,92]. We find that the so-called saddle-free Newton (SFN) method [65] yields considerably better results than first-order methods. There the update direction is given by $-|H|^{-1}g$ with H being the Hessian, g the gradient, and the absolute value $|H|$ defined by spectral calculus. An instructive way to see why this leads to a speedup is to write the Hessian H as $H = \sum_i \lambda_i |v_i\rangle\langle v_i|$, where v_i is the eigenvector to eigenvalue λ_i . The update direction of the SFN method then reads $-|H|^{-1}g = -\sum_i |\lambda_i|^{-1} |v_i\rangle\langle v_i|g$. Since the vectors $|v_i\rangle$ form a basis, this can be interpreted as a rescaling of $|g\rangle$ by $|\lambda_i|^{-1}$ in the directions $|v_i\rangle$. As with the standard Newton method, this leads to a large rescaling if the curvature in a particular direction is small, resulting in large steps even close to the saddle point. Taking the absolute value of the eigenvalues then ensures that saddle points are repulsive. For numerical stability, it is beneficial to introduce a damping term that offsets the eigenvalues of H that are very close to zero before the inversion.

Algorithm 1 describes a single step of the damped saddle-free Newton method with damping parameter λ and is a generalization of the original SFN method [65] to manifolds.

input: Curve parametrization $Y_t(Y_0, \Delta)$, objective function $\mathcal{L}_I(Y_t)$, damping parameter λ

- 1 Compute the gradient \mathbf{G} and Hessian H of $\mathcal{L}_I(Y_t)$ at Y_0 .
- 2 Determine the update direction

$$\begin{pmatrix} \Delta \\ \Delta^* \end{pmatrix} = (|H| + \lambda \mathbb{1})^{-1} \begin{pmatrix} \mathbf{G}^* \\ \mathbf{G} \end{pmatrix}.$$
- 3 Determine the step size $\tau = \underset{t}{\operatorname{argmin}} \mathcal{L}_I(Y_t(Y_0, \Delta))$.

return $Y_\tau(Y_0, \Delta)$

Algorithm 1. SFN update.

Algorithm 1 is formulated in a way that is compatible with an update in Euclidean space as well as an update on the Stiefel manifold. In Euclidean space we update along the curve $Y_t(\cdot, \cdot) : \mathbb{C}^{D \times d} \times \mathbb{C}^{D \times d} \rightarrow \mathbb{C}^{D \times d}$ with $Y_t(Y_0, \Delta) = Y_0 + t \Delta$ for an update direction Δ . On the Stiefel manifold we have $Y_t(\cdot, \cdot) : \operatorname{St}(D, d) \times \mathfrak{T} \rightarrow \operatorname{St}(D, d)$ with the curve given by geodesic (27) and \mathfrak{T} being the tangent bundle on $\operatorname{St}(D, d)$. The step size is determined by locally optimizing over the parameter t using standard gradient-free optimizers. We derive an expression for the Hessian on the Stiefel manifold in Appendix A.2. In Appendix A.3, we also provide a detailed discussion and expressions for the optimization in complex Euclidean space.

Algorithm 2 describes the main mGST routine. It can be run with different choices of smooth objective functions, and we use MSE (19) by default. In our numerics we often find that optimizing the loglikelihood function after the MSE can improve estimates (see Appendix A.4 for a discussion).

The algorithm alternates updates on A , \mathcal{K} , and B . Updates are performed using Algorithm 1 on the tangent spaces of the respective Stiefel manifolds. In order to

input: Data $\{\mathbf{y}_{j|i}\}_{i \in I, j \in [n_E]}$, batch size κ , Kraus rank r_K , initialization $(A^0, \mathcal{K}^0, B^0)$, stopping criterion

- 1 $i \leftarrow 0$
- 2 **repeat**
- 3 Select batch $J \subset I$ of size $|J| = \kappa$ at random
- 4 $A^{i+1} \leftarrow$ update A^i with objective $\mathcal{L}_J(\cdot, K^i, B^i; \mathbf{y})$ along geodesic on $\operatorname{St}(dn_E, d)$
- 5 $\mathcal{K}^{i+1} \leftarrow$ update \mathcal{K}^i with objective $\mathcal{L}_J(A^{i+1}, \cdot, B^i; \mathbf{y})$ along geodesic on $\operatorname{St}(r_K d, d)^{\times n}$
- 6 $B^{i+1} \leftarrow$ update B^i with objective $\mathcal{L}_J(A^{i+1}, \mathcal{K}^{i+1}, \cdot; \mathbf{y})$ along geodesic on $\operatorname{St}(d^2, 1)$
- 7 $i \leftarrow i + 1$
- 8 **until** stopping criterion is met at $i = i_*$;
- 9 **return** $(A^{i_*}, \mathcal{K}^{i_*}, B^{i_*})$

Algorithm 2. mGST.

achieve good convergence, we run the optimization with mGST in two consecutive steps: we start from a random initialization and perform a coarse grained optimization with a small batch size κ , i.e., only using κ many random gate sequences from I for each update step. The batching of data results in lower computation time for the derivatives and adds a factor of randomness to the optimization, which avoids getting stuck at suboptimal points to a certain degree. We terminate the first optimization loop when the objective function $\mathcal{L}_I(A^i, \mathcal{K}^i, B^i | \mathbf{y})$ is smaller than an *early stopping value* δ , which is obtained from the data as follows.

For a number of m samples per sequence, the outcome probabilities of each sequence for the true gate set are given by

$$y_{j|i} = k_{j|i}/m, \quad (28)$$

where $k_{j|i}$ is the number of times outcome j is measured upon applying gate sequence \mathbf{i} . Because of Born's rule, $k_{j|i}$ is distributed according to the multinomial distribution $M(m, (p_{1|i}, \dots, p_{n_E|i}))$ with probabilities $\{p_{j|i}\}_j$ and m trials. We estimate the expectation value of the objective function from the values $y_{j|i}$. This provides us with a rough estimate for how low the objective function value can become, given the sample counts $k_{j|i}$. Then we set the early stopping value to be twice that estimate,

$$\delta := 2 \mathbb{E}_{\tilde{k}_{j|i} \sim M(m, (y_{j|i}))} \frac{1}{|I|} \sum_{\mathbf{i} \in I} \sum_j (y_{\mathbf{i}}^j - \tilde{k}_{j|i}/m)^2. \quad (29)$$

Hence, we require the objective function on the full data set to be close to its expectation value for the measured probabilities $y_{j|i}$ obtained from m samples.

While computationally inexpensive the minibatch stochastic optimization does not converge to an optimal point on the full data set I . In a second optimization loop, we initialize the mGST algorithm with the result from the first run and use all the data for the updates. Formally, we choose the batch size $\kappa = |I|$ and thereby make the random batch selection obsolete. We perform these more costly update steps until the change in objective function reaches a desired relative precision ϵ ,

$$\mathcal{L}_I(A^i, \mathcal{K}^i, B^i | \mathbf{y}) - \mathcal{L}_I(A^{i-1}, \mathcal{K}^{i-1}, B^{i-1} | \mathbf{y}) \leq \delta \epsilon, \quad (30)$$

or a maximal number of iterations is exceeded.

The first optimization run is initialized with a random gate set parameterized by A^0, \mathcal{K}^0 , and B^0 (see Sec. II A). For the random initialization, we make use of the *Gaussian unitary ensemble* (GUE). A matrix H belongs to the GUE if $H = (M + M^\dagger)/2$, where M is a *complex Gaussian matrix*, i.e., the real and imaginary parts of each M_{ij} are independently drawn from $\mathcal{N}(0, 1)$, the normal distribution with zero mean and unit variance. In this case we

write $H \sim \text{GUE}$. For A^0 and each gate in \mathcal{K}^0 , we take the first d columns of e^{iH} with $H \sim \text{GUE}$ to obtain a random isometry \mathcal{K}^0 . For B_0 , we take a complex Gaussian matrix and normalize it such that $\text{Tr}[B^{0\dagger}B^0] = 1$.

Importantly, due to the nature of nonconvex optimization, several initializations can be needed to converge to a satisfactory minimum.

IV. NUMERICAL ANALYSIS

In this section, we evaluate the performance of mGST in different scenarios in numerical simulations. In particular, we compare its performance to the state-of-the-art implementation for gate set tomography, pyGSTi [18], in the regimes where both methods can be applied.

For pyGSTi to be applicable, one has to use structured gate sequences inspired by standard quantum process tomography. In Sec. IV A we evaluate the performance of mGST and pyGSTi on minimal measurement sequences and different models to find that mGST benefits from flexibility in the sequence design and a fully general model parametrization. Section IV B numerically validates the expected inverse square-root scaling of the reconstruction error with the number of measurement samples per sequence for different noise regimes. Section IV C numerically determines the required number of random sequences to accurately reconstruct simple and random gate set models with mGST for different Kraus ranks. In Sec. IV D we follow up with a numerical demonstration of unitary noise characterization for a three-qubit gate set using *a priori* knowledge in the initialization. Finally, in Secs. IV E and IV F we discuss the choice of initialization and hyperparameters, as well as the runtime of mGST.

For a model of n gates reconstructed from m measurements of sequence length ℓ , we validate the performance of mGST by computing MVE (18) over all possible n^ℓ sequences, or 10^4 random sequences of length ℓ if $n^\ell > 10^4$. Usually, $m \ll \min(n^\ell, 10^4)$ and the MVE can be thought of as a generalization error on the predicted output probabilities of the gate set estimate. The gate sets studied in this section all use the same target initial state $|0\rangle\langle 0|$ and computational basis measurement, although with different levels of noise applied to them. For instance, we often use global depolarizing noise, which acts on a quantum state ρ as $\rho \mapsto (1-p)\rho + p\mathbb{1}/d$. For the numerics presented here, we use a maximum of 100 reinitializations (if not stated otherwise). A discussion of the required number of initializations is given in Sec. IV E. A PYTHON implementation of mGST and a short tutorial can be found on GitHub [93].

A. Gate set and measurement structure

We compare mGST and pyGSTi for the minimal number of sequences doable with each method and for gate

sets of different conditionings, without using the compression capabilities of mGST yet. We find that mGST is more flexible in the sequence design and model parametrization, while generating estimators with lower mean variation errors in several regimes.

The traditional strategy for GST, akin to standard quantum process tomography, is to generate a frame for $L(\mathcal{H})$, measure each gate in that frame and generate an estimate for each gate by applying the pseudoinverse of the measurement operator.

This is particularly important for the first reconstruction step in pyGSTi where the sequences that generate the frame are called *fiducials*. The strategy of pyGSTi is to obtain an initial estimate via the pseudoinverse, followed up by local optimization of a particular cost function [18]. In contrast, we perform mGST using random initializations and thereby do not rely on designated fiducial sequences.

Figure 1 compares mGST to pyGSTi focusing on the regime of very few gate sequences, showing what is needed in terms of measurement effort to obtain low mean variation errors for different gate sets. In order to test pyGSTi in the regime of low sequence counts, we replace the five standard fiducial sequences with two or three fiducial sequences drawn uniformly at random, thereby reducing the total sequence number from 92 to 18 or 39 sequences. Since mGST is compatible with any sequence design, we use between 10 and 58 random sequences for mGST to explore the low sequence count region.

The first gate set we study is the so-called *XYI model*, the standard single-qubit example in the pyGSTi package [45]. The XYI model consists of the identity gate, a $\pi/2$ X rotation and a $\pi/2$ Y rotation on the Bloch sphere, with initial state $|0\rangle\langle 0|$ and measurement in the computational basis. Results for the ideal XYI model can be seen on the left in Figure 1, with mGST and pyGSTi performing identically for 92 sequences. Comparing the results for 18 sequences we find that mGST does not converge on more than 50% of trials, which reflects in the median MVE being above 10^{-1} , while pyGSTi achieves lower median errors. Comparing the 38 and 39 sequence medians however, we find that mGST yields lower error models than pyGSTi.

The subsequent models analyzed successively deviate from the simple unitary XYI model and highlight the versatility of our manifold approach. Since the full CPT parametrization used in our optimization (21) is agnostic to any special gate set properties, we expect it to perform well for all possible CPT maps as gate implementations. For instance, for random and specific non-Markovian channels. pyGSTi on the other hand uses a parametrization of Lindblad type and is therefore based on a more limited model space.

To illustrate this comparison, we perturb the XYI model by adding amplitude damping noise to each gate. The

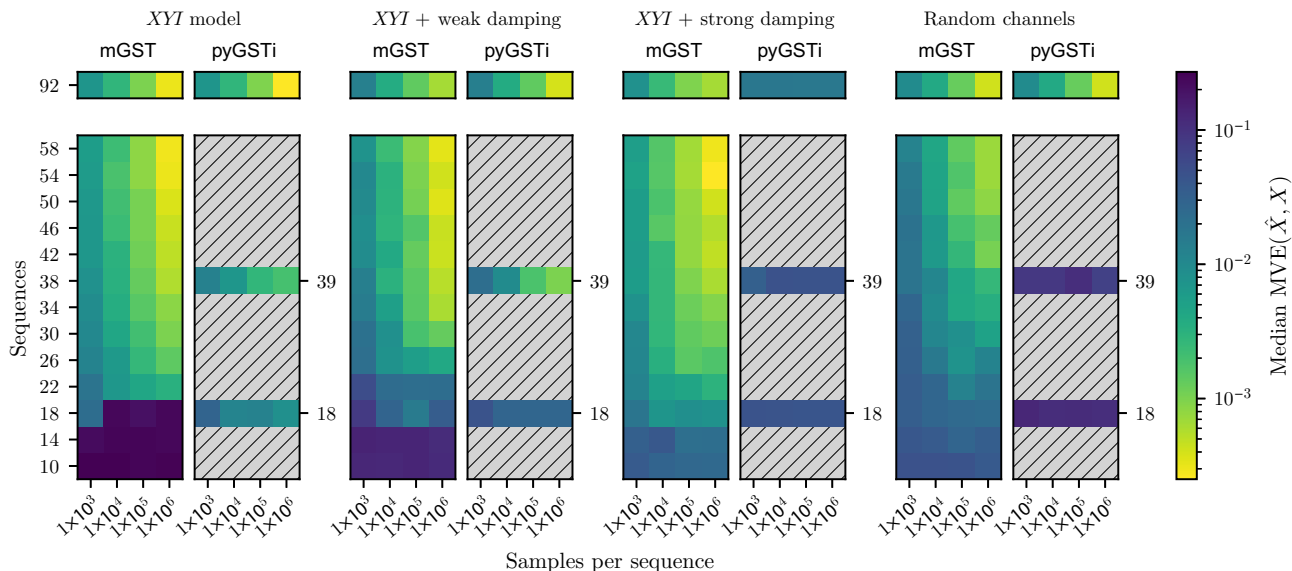


FIG. 1. MVE comparison between mGST (with loglikelihood cost function and $r_K = 4$) and pyGSTi, showing the dependence on the number of sequences, the number of samples per sequence, and the gate set on a single qubit. The number of sequences used by mGST in the range 10–58 are drawn uniformly at random, while the sequences for pyGSTi need to follow the pyGSTi fiducial design and are limited to the fixed sequence counts (18 and 39). We choose independently drawn random fiducials for each instance. The 92 sequences used by both mGST and pyGSTi are taken from the standard pyGSTi sequence design for the XYI model. All sequences are of length $\ell = 7$. The MVE depicted in each square is the median result for ten instances, each with random statistical measurement noise and a random sequence drawn from the uniform distribution. In the random channel scenario, a new random channel is used for each instance. The XYI model is a simple unitary model used in the GST literature and the weak damping model consists of amplitude damping noise on each gate with $\Gamma = 0.94$, while the strong damping model uses $\Gamma = -0.6$. A complete description of the models used can be found in the main text. Each model has additional depolarizing noise of strength $p = 0.01$ on the initial state.

amplitude damping channel can be written in terms of the Kraus operators

$$K_1 = \begin{pmatrix} 1 & 0 \\ 0 & \Gamma \end{pmatrix} \quad \text{and} \quad K_2 = \begin{pmatrix} 0 & \sqrt{1 - |\Gamma|^2} \\ 0 & 0 \end{pmatrix},$$

which arise, e.g., from the Jaynes-Cummings model of a qubit system interacting with a quantized bosonic field [94]. How well mGST and pyGSTi perform on a model with $\Gamma = 0.94$ can be seen in the center left block of Fig. 1 (XYI + weak damping). We find generally similar performance, with mGST being a bit more accurate on 38 sequences, and a bit less accurate on 92 sequences with 10^6 samples.

Increasing the interaction time between qubit and environment leads to memory effects and strong non-Markovianity of the amplitude damping channel at $\Gamma = -0.6$. This scenario is shown in the center right plot of Fig. 1, and we see that while the accuracy of mGST is the same as before, the model parametrization of pyGSTi cannot fit the model with MVEs below 10^{-2} , independent of the sequence or sample count.

For the last comparison (rightmost block in Fig. 1), we look at the performance for random full Kraus rank channels. Each channel is constructed by drawing a Haar random $d^3 \times d^3$ unitary and then taking its first d columns.

The resulting $d^3 \times d$ matrix is an isometry and therefore constitutes a valid set of Kraus operators. Note that this construction is different from the previous construction of random channels via the Gaussian unitary ensemble. The results show that mGST can reconstruct these models from low sequence counts, while pyGSTi does not yield good estimators. Using the standard sequence design of 92 sequences, mGST and pyGSTi have identical accuracy again, suggesting that random channels are typically well within the model space of pyGSTi after all. These demonstrations show that mGST is indeed flexible in the sequence design with state-of-the-art performance for arbitrary gate set implementations.

B. Number of samples per sequence

The probability associated with every sequence is estimated from a finite number of samples. Here, we study the resulting effect on the reconstruction accuracy as measured by the MVE more closely.

For a high number m of samples per sequence, each probability y_i^j in the objective function is estimated with an error of order $1/\sqrt{m}$. Therefore, we expect the MVE to also decrease as $1/\sqrt{m}$ if the algorithm converges to the global minimum. This scaling was observed to hold for pyGSTi [18]. In order to be able to compare the scaling




FIG. 2. Reconstruction of the XYI gate set for different levels of depolarizing noise with strength p and unitary noise with strength γ on each gate. The unitary noise is given by $e^{i\gamma H}$ with $H \sim \text{GUE}$. Additional depolarizing noise with $p = 0.01$ is applied before measurement. The mGST algorithm is run on the loglikelihood cost function with $r_K = 4$ (max), which is the same as for pyGSTi. As gate sequences, we again used the standard pyGSTi fiducial sequences, with the number of measurements per sequence between 10^2 and 10^8 . The lines connect data points of which each is the median over ten runs. For each run, a new random over-rotation is drawn and new measurements are simulated. The measurement sequences are the 92 sequences provided by the pyGSTi software, with a maximum sequence length of $\ell \leq 7$.

of mGST directly to that of pyGSTi, we use a standard pyGSTi setting: the gate set is the XYI model (with $\pi/2$ rotations) and the gate sequences are the standard pyGSTi sequences for this model with a maximum sequence length of $\ell = 7$.

We add noise to the gate set by varying the amount of depolarizing noise with strength p on each gate and also over-rotating each gate by a random unitary. The random unitaries are given by $e^{i\gamma H}$ with $H \sim \text{GUE}$. In particular, this means that H can be bounded on average as follows. We can write $H = (M_1 + M_1^T + i(M_2 - M_2^T))/2$ with M_i being independent Gaussian matrices. Next, we use Gordon’s theorem for Gaussian matrices (see, e.g., Theorem 5.32 of Ref. [95]), which tells us that $\mathbb{E}\|M_1\|[\infty] \leq 2\sqrt{d}$. The relevant magnitude of the random generator H is then in expectation upper bounded as

$$\mathbb{E}\|H\|_\infty \leq 2\mathbb{E}\|M_1\|_\infty \leq 4\sqrt{d}. \quad (31)$$

State preparation and measurement are assumed to be noise-free in this setup; however, for a fixed sequence length, the depolarizing noise per gate is equivalent to a global depolarizing channel applied before measurement, since it commutes with the unitary gates. Figure 2 depicts the resulting MVE scaling of the reconstruction where data are generated using different numbers of samples per sequence m .

We observe that mGST follows the expected scaling in m , matching the scaling of pyGSTi for different levels of unitary and depolarizing noise.

C. Number of sequences

Arguably the most challenging experimental requirement of GST is the number of measurement settings (sequences) that are required for a successful gate reconstruction. One of the main motivations of compressive GST is to employ structure constraints, i.e., to reduce the number of degrees of freedom of the reconstruction problem, in order to reduce the required number of measurements. Instead of reconstructing arbitrary quantum channels we aim at reconstructing low-rank approximations of the gate set elements. In addition, we expect that by using the mGST algorithm, compressive recovery is already possible from a “few” randomly selected sequences. We here numerically demonstrate that this is indeed the case.

The top row of Fig. 3 shows the median performance in MVE against the number of randomly chosen sequences for different Kraus ranks. On the left are the results for the single-qubit XYI model as defined in Sec. IV A. On the right are the results for the XYICNOT gate set that is based on the identity, controlled-NOT (CNOT) gate, and Pauli- X and Pauli- Y rotations on each qubit individually, with rotation angle $\pi/2$.

We observe a phase transition in the MVE that indicates a minimal number of sequences that are required for the successful reconstructions of the gate sets. As expected, constraining the reconstruction to a lower Kraus rank indeed reduces the amount of required sequences in the reconstruction in most cases.

An intriguing exception is the $r_K = 1$ reconstruction of the two-qubit gate set that exhibits the worst reconstruction performance compared to higher rank constraints. We suspect that this is due to the optimization problem being more dependent on the initialization for $r_K = 1$. In more general settings, it has been observed that the optimization over matrix product states with fixed Kraus rank can be unstable and using rank-adaptive optimization techniques yield much better performance [96,97]. This motivates us to use a slightly higher rank in the optimization than the expected rank of an effective approximation of the gate set. In accordance with this intuition, we find that it is beneficial to constrain the optimization to $r_K = 2$ in order to achieve an accurate unit-rank approximation. The same effect is also observed in the single-qubit example when taking a detailed look at the number of required initializations (see Sec. IV E), yet less pronounced. In the bottom row of Fig. 3 we show the recovery rates for random unitary models, with the reconstruction now using a fixed Kraus rank of $r_K = 2$. Note that there are three sources of randomness present in the data: first the Haar-random unitary gates,

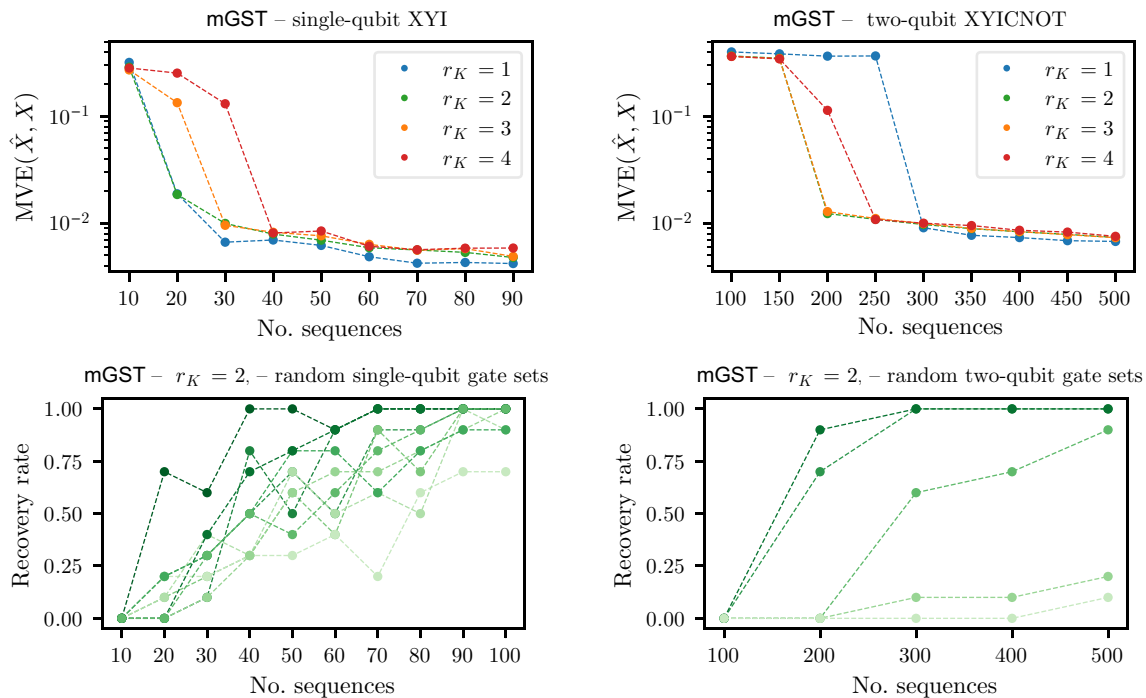


FIG. 3. Median error of mGST run on the least-squares cost function, plotted over the number of sequences for a single-qubit model (top left) and a two-qubit model (top right). Each data point is the median over the results from ten different random sequences. Measurement data for the XYI model—and the XYICNOT gate set—are taken from a noisy version with depolarizing noise of strength $p = 0.001$ on each gate, depolarizing noise with strength $p = 0.01$ before measurement, as well as independent random unitary rotations e^{iyH} with $\gamma = 0.001$ and $H \sim \text{GUE}$ on each gate. In the bottom left plot the recovery rates for the reconstruction of different models of three Haar random unitaries are shown. For each gate set, the average over ten draws of random sequences is shown. A gate set is classified as recovered if the MVE falls below 0.03. The bottom right plot depicts the recovery rate for random two-qubit gate sets of the form $\mathcal{G} = \{\mathbb{1}_4, \mathbb{1}_2 \otimes U_1, \mathbb{1}_2 \otimes U_2, U_1 \otimes \mathbb{1}_2, U_2 \otimes \mathbb{1}_2, U_{12}\}$, where U_1 and U_2 are Haar random single-qubit unitaries and U_{12} is a Haar random two-qubit unitary. Additionally, each single-qubit and two-qubit gate contains depolarizing noise of strength $p = 0.001$ and depolarizing noise of strength 0.01 is applied before measurement. The recovery rate is averaged over ten random sequence draws. For all gate sets, the sequences are drawn uniformly at random with sequence length $\ell = 7$ and $m = 1000$ samples per sequence. The maximum numbers of initializations are 80, 33, 17, and 10 for Kraus ranks 1, 2, 3, 4, respectively. They are chosen such that the maximal computation time is equal among different ranks.

then the random drawing of gate sequences, and finally the random initialization of the algorithm. Each shade of green corresponds to one random gate set, and the recovery rate tells us how many of the ten random sequence sets lead to a successful reconstruction, given a budget of 33 initializations. We find that the random single-qubit gate sets all have similar recovery rates, with a successful reconstruction possible from $n_{\text{seq}} = 20$ to $n_{\text{seq}} = 30$ sequences, and a high rate of recovery at $n_{\text{seq}} = 100$ sequences. In the two-qubit case a different picture emerges, where two of the random gate sets show a high recovery rate at $n_{\text{seq}} = 200$ sequences (akin to the XYICNOT model), while the least favorable random gate set is only recoverable at $n_{\text{seq}} = 500$ sequences. This shows that random gate sets for two qubits can have very different conditioning.

1. Sequence number comparison to pyGSTi

Comparing the number of random sequences needed for mGST and the number of sequences for pyGSTi is

not straightforward. The standard pyGSTi data processing pipelines crucially rely on specific, fixed sequence construction. For this reason, pyGSTi cannot be applied to the type of data that we use here. We can however compare the number of random sequences with the number of deterministic sequences that the standard implementation of pyGSTi uses. For the single-qubit XYI model, the minimal number of sequences given in the pyGSTi implementation is $n_{\text{seq}} = 92$. This is significantly larger than the number of random sequences at which the phase transition of mGST in Fig. 3 appears. However, the $n_{\text{seq}} = 92$ sequences are overcomplete by design, and we find that pyGSTi can also reconstruct the XYI model with $n_{\text{seq}} = 48$ sequences. Yet we do not find the same sequence design to be successful for three Haar-random single-qubit gates, indicating that the choice of sequences is well tailored to the XYI model. The reduction in sequences becomes more pronounced for the two-qubit gate set studied in the top right of Fig. 3. For this gate set, the minimal number of gate sequences that pyGSTi uses is

$n_{\text{seq}} = 907$, which is significantly larger than what mGST needs.

D. Characterizing unitary errors using prior knowledge

In the previous section we demonstrated compressive gate set tomography for one- and two-qubit gate sets using agnostic random initializations. A major obstacle in going beyond reconstructing entire two-qubit gate sets even compressively on desktop hardware is that besides runtime and storage, the number of required random initializations until proper convergence also grows in principle with the number of qubits. This is due to longer sequences being required for tomographic completeness, leading to a higher-order polynomial in the cost function. This situation can be remedied by using prior knowledge, such as the target gate set, for the initialization. In this case, a gate set in the vicinity of the initial point will be found that is in better agreement with the data. In a conceivable experimental scenario the gates are more or less known due to the physical setup and previous benchmarking rounds, but further calibration requires information about present coherent errors. The use of prior knowledge can be seen as a situational tool to further reduce runtime when applicable, but for general purpose verification and characterization, no initial point is to be assumed for compressive GST with the mGST algorithm.

To showcase the characterization of unitary errors using prior knowledge, we take a three-qubit gate set that is the direct generalization of the previous two-qubit XYICNOT model, by adding the local X and Y rotations as individual gates to the third qubit and adding a CNOT gate between qubits 2 and 3. We then apply a global random rotation to each gate individually, as well as depolarizing noise on each gate. From random sequences of fixed sequence length we can then, in theory, fit the noisy model perfectly via an $r_K = 1$ approximation, as the depolarizing channels commute with the unitary gates and can be pulled into initial state or measurement. In Fig. 4 we see that mGST is indeed able to precisely reconstruct the rotated gates, as shown by the average diamond norm error. We chose a comparatively high number of 10^5 samples per sequence to showcase that high accuracy can indeed be realized using this method: for instance, only 256 sequences are enough to achieve an average diamond norm distance of around 0.007 between the reconstructed unitary gates and the true unitary gates, which include over-rotations. The fact that these over-rotations are modeled as being global on all three qubits suggests that we can efficiently characterize unitary crosstalk as well, by capturing the effects of single- and two-qubit gates on their neighbors within a three-qubit region.

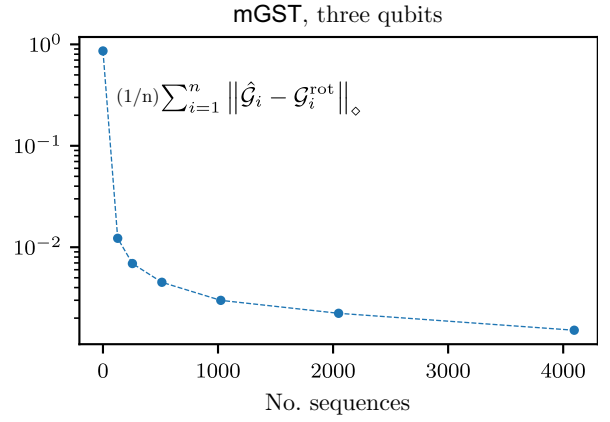


FIG. 4. Average diamond distance between three-qubit rotated target gates $\mathcal{G}_i^{\text{rot}}$ and their unitary ($r_K = 1$) mGST estimators $\hat{\mathcal{G}}_i$ as a function of the number of sequences. The 0-sequence data mark the average diamond distance between initialization \mathcal{G}_i and $\mathcal{G}_i^{\text{rot}}$. The rotated gates $\mathcal{G}_i^{\text{rot}}$ are related to their counterparts \mathcal{G}_i by independent random (global) over-rotations on each gate, given by $\exp(i\gamma H)$ with $H \sim \text{GUE}$ and $\gamma = 0.05$, leading to $\mathbb{E}\|\gamma H\|_\infty \leq 2\sqrt{2}/5$. The data are simulated from the gates $\mathcal{G}_i^{\text{rot}}$ with depolarizing noise of strength $p = 0.01$ on each gate and the measurements are taken from random sequences of length $\ell = 7$, with $m = 10^5$ samples per sequence.

E. Implementation details and calibration

We now provide more details on the simulations, the criteria for successful recovery, and the required number of initializations. To simulate measurements on a gate sequence \mathbf{i} , we first compute the outcome probabilities $p_{j|\mathbf{i}}$ from Eq. (16) of the POVM elements according to the model gate set in question. Afterwards we draw m samples from the multinomial distribution $M(m, (p_{1|\mathbf{i}}, \dots, p_{n_E|\mathbf{i}}))$, where $\sum_j p_{j|\mathbf{i}} = 1$. Let $k_{j|\mathbf{i}}$ be the number of times outcome j occurred for sequence \mathbf{i} . Then Algorithm 2 optimizes the objective function (20) on the estimated probabilities $y_{j|\mathbf{i}} = k_{j|\mathbf{i}}/m$.

For the single-qubit examples, the batch size $\kappa = 50$ is chosen, while for the two-qubit example, we use $\kappa = 120$. The choice of batch size determines the number of values summed over in Eq. (20). Therefore, the computation time of the objective function and its derivatives scales linearly in κ , making a small batch size favorable. However, it cannot be set too small, otherwise the update directions become highly erratic, and no convergence is reached. A general rule of thumb is to set the batch size close to the number of free parameters in the model. Another hyperparameter is the damping value λ for the saddle-free Newton method described in Algorithm 1. We find that a fixed value of $\lambda = 10^{-3}$ leads to the best results across the models tested.

Judging whether mGST recovers a gate set by looking at the attained objective function value can only be done if the set of measured sequences is informationally complete.




FIG. 5. Least-squares objective function versus MVE at the end of the mGST optimization, with the gray bar indicating the range of stopping values below which a run is considered successful. The plot illustrates that results with a large MVE also have a large objective function if enough sequences are measured. The experiment is the same as in the top left of Fig. 3 for $r_K = 4$. The color of each point indicates the number of measured sequences on which mGST is run.

Then there is a unique (up to gauge) global minimum in the least-squares minimization problem and the minimum corresponds to the true gate set in the limit of infinitely many samples per sequence.

In Fig. 5 we take a look at the correlation between the final least-squares objective function value $\mathcal{L}(\hat{\mathcal{X}}, \mathcal{y})$ and the mean variation error $\text{MVE}(\hat{\mathcal{X}}, \mathcal{X})$. We see that, for a low number of sequences (10–20), a low objective function value does not imply a low MVE, yet for higher numbers of sequences, an objective function value below 10^{-3} implies an MVE around 10^{-2} . For sufficiently many sequences, the gray line indicating our success criterion clearly separates two clusters of points, meaning that no intermediate quality fits are found in our model space. In this sequence regime either the algorithm converges to a fit as good as the sample count allows, or it does not converge at all. Therefore, restarting the algorithm when an initialization turns out to be bad yields practically optimal results. A thorough analysis of the probability of obtaining an informationally complete set of random sequences is left for future work.

To give an intuition on how many initializations are required for mGST to converge, we can take a look at data from a modified XYI model with gates $\{\mathbb{1}, e^{i\alpha\sigma_y/2}, e^{i\alpha\sigma_x/2}\}$, simulating more difficult gate set conditioning. Figure 6 shows histograms for the number of reinitializations needed for convergence. The data combine the results for α between $\pi/18$ and $\pi/2$, with depolarizing noise of strength $p = 0.001$ on each gate and $p = 0.01$ on the initial state, as well as a maximum of 100 reinitializations.

We observe that in 48.4% of all cases convergence is reached on the first attempt, and in 90.4% of cases four or fewer reinitializations are required. The histogram




FIG. 6. Relative frequencies of reinitialization counts for mGST run on modified XYI models (see the main text for details).

indicates that the chance of needing multiple reinitializations rapidly decays and that rank-1 optimization is more sensitive to bad initializations compared to rank-4 optimization.

F. Runtime and scaling

In order to assess the runtime scaling of Algorithm 2 in the problem parameters we identify the two most time-consuming steps as the computation of the second derivative terms in H and the diagonalization of H for the SFN update on the gates.

Recall that r_K denotes the Kraus rank of the gate estimators, n the number of gates in the gate set, ℓ the number of gates per sequence, and κ the number of sequences per batch. The computation of the second derivative terms scales as $O(\kappa \ell^3 d^6)$, while the eigendecomposition of H scales as $O((2nr_K d^2)^3)$. A computationally less expensive variant is to not optimize over all variables of the full gate tensor at once, but over the individual gates one after another. The complexity of the eigendecompositions then reduces to $O(n(2r_K d^2)^3)$, which is beneficial for large gate sets. However, this approach also leads to slower convergence, and we choose to optimize over the full tensor by default.

Table I contains runtimes for the system sizes studied in our numerical experiments. We find that reconstruction of single-qubit gate sets can be achieved within seconds and low-rank two-qubit gate sets can be reconstructed within minutes on a single modern 32-core CPU. The runtimes for the gradient descent method on the two-qubit example are very fast up to high ranks; however, we generally find that the default Newton method is more reliable across different gate set and sequence scenarios. If a good initialization is known, the runtime is reduced drastically. For instance, the three-qubit reconstruction done with prior knowledge about the gates in Sec. IV D took 5 h and 30 min to complete.

TABLE I. Average runtimes for one, two, and three qubits with selected Kraus ranks on a modern 32-core CPU. The average runtime is calculated as the average time until the first successful reconstruction with a random initialization is achieved. The one-qubit gate set is the same as in Fig. 2, the two-qubit gate set the same as in Fig. 3, and the three-qubit gate set is the same as in Fig. 4 but without using any prior knowledge. For the gradient descent method in the three-qubit scenario, no convergence is achieved within the maximum iteration limit. The average runtimes for pyGSTi are determined on the same models and with the same sequences.

mGST: default (Newton)				
	$r_K = 1$	$r_K = 4$	$r_K = 8$	$r_K = 16$
One qubit	4 s	14 s	/	/
Two qubits	8 min 8 s	39 min 50 s	2 h 26 min 10 s	9 h 1 min 43 s
Three qubits	3 d 2 h	/	/	/
mGST: gradient descent				
	$r_K = 1$	$r_K = 4$	$r_K = 8$	$r_K = 16$
One qubit	26 s	17 s	/	/
Two qubits	5 min 40 s	6 min 0 s	5 min 51 s	4 min 47 s
Three qubits	No convergence	/	/	/
pyGSTi				
One qubit	5 s			
Two qubits	1 h 6 min 48 s			
Three qubits	No result after 4 d 14 h			

V. NOISE-MITIGATION FOR SHADOW ESTIMATION WITH COMPRESSIVE GST

Gate set estimates provide a detailed picture of the imperfections that can inform prioritization and further experimental efforts [15,19,21–23]. Coherent error estimates can often be directly corrected for by adjusting or optimizing the control. We expect that the more economically accessible compressive estimates from mGST can be used in place of traditional GST estimates in the above applications, in particular when complemented with RB estimates of incoherent noise effects. However, it is beyond the scope of this work to demonstrate mGST in a full engineering cycle of a quantum computing device.

To still showcase the value of GST estimates, we sketch another novel application that can be demonstrated without simulating a whole engineering cycle. Generally speaking, noise characterizations can be used to mitigate noise-induced biases in other quantum characterization protocols by adapting the classical postprocessing [98,99]. Given a device that can repeatedly prepare a quantum state, a fundamental characterization task is to estimate the expectation values of observables from measurements. In particular, an informationally complete measurement allows one to estimate arbitrary observables from the same data in the postprocessing. Such an informationally complete measurement can be implemented on a quantum computing device by measuring in sufficiently many random bases, the prototypical example being measurements in random Pauli bases. Huang *et al.* [46] showed how to derive optimal guarantees with exponential confidence for estimating multiple observables simultaneously using a

median-of-means estimator and explicit bounds on the variance for random basis measurement that constitute unitary 3-designs. They introduced the term “classical shadow” to refer to the elements of a dual frame of the informationally complete POVM corresponding to observed samples; see Appendix A 5 for a brief summary. Importantly, one can often arrive at high-precision estimates of observables long before one has measured all the informationally complete bases in multiqubit systems.

In practice, however, implementing a random bases measurement, say, by applying a unitary rotation followed by a computational bases measurement will suffer from noise from the gates and readout. This has motivated the development of robust variants of shadow estimation that either make use of simple depolarizing noise models of known strength [100] or perform a separate RB-style experiment that estimates the depolarizing noise strength induced by a gate-independent channel acting between the rotation and the measurement [101]. Using GST estimates provides a complimentary, flexible approach to mitigate even highly gate-dependent noise with finite correlations in shadow estimation.

We demonstrate how GST estimates on two-qubit pairs can be used to calculate noise-robust classical shadow estimators in postprocessing. Our robust estimation scheme consists of two distinct stages, each consisting of multiple steps. (I) The *calibration stage* consists of three phases. (i) The local channels implementing each combination of two local gates are reconstructed with mGST. (ii) The gauge of these gate estimates is matched to the gauge in which the ideal gates and the observables are given. For this step, we use the gauge optimization provided by the pyGSTi

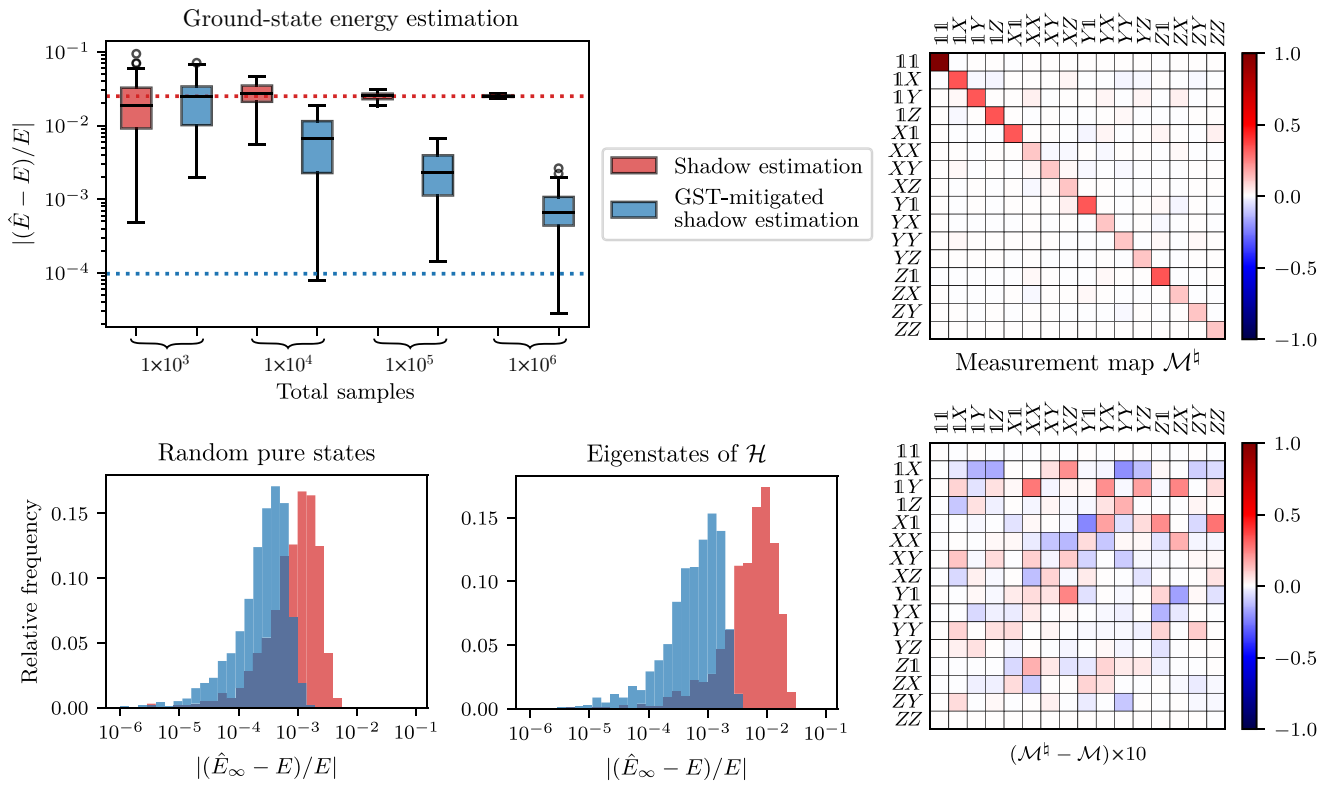


FIG. 7. Energy estimation for the ten-qubit Hamiltonian $\mathcal{H} = \frac{1}{2} \sum_{j=1}^{10} (\sigma_x^j \sigma_x^{j+1} + \sigma_y^j \sigma_y^{j+1} + \sigma_z^j \sigma_z^{j+1} - \sigma_z^j)$ with periodic boundary conditions. The top left plot shows the sample dependence of the relative accuracy $|(\hat{E} - E)/E|$ for estimating the ground-state energy. The colored blocks extend from the first quartile to the third quartile around the median (black line) of the data (50 repetitions per sample value). The whiskers extend from the 5th to the 95th percentile and the dashed lines indicate the infinite sample expectation values. In the bottom left and bottom middle plots two histograms are shown that compare the theoretical infinite sample energy estimates \hat{E}_∞ (biases) for 1000 random pure states and for all 1024 eigenstates of the Hamiltonian, respectively. All simulations are done with noisy Clifford gates, whose average gate fidelity to their ideal counterparts is at 0.99 ± 10^{-3} . In the top right plot the Pauli transfer matrix of a two-qubit effective measurement map \mathcal{M}^b under this noise model is shown. The bottom right plot displays the difference between \mathcal{M}^b and its noise-free counterpart \mathcal{M} . Gate estimates for GST-mitigated shadow estimation are produced with mGST ($r_K = 2$), using 400 random sequences equally distributed among sequence lengths $\{6, 7, 8, 9\}$ with 10^4 samples per sequence. The noise in the simulations is given by two-qubit random unitary noise $e^{i\gamma K}$ with $K \sim \text{GUE}$. The error parameter is $\gamma = 0.14$ on H and on HS , leading to the aforementioned average gate fidelities of about 0.99. For the bottom left histogram, random pure states are generated as $U|0\rangle$, with U drawn according to the Haar measure.

package [45]. (iii) From the gauge-optimized channel estimates, we numerically calculate the *effective measurement map* when implementing random Pauli measurements with the characterized noisy gate set.

(II) After calibration, the second stage is a *shadow estimation* protocol that consists of two separate phases: (i) the *data acquisition* by repeatedly measuring the unknown state of the quantum device in a randomly selected Pauli basis; (ii) the *classical postprocessing* where estimators of the observables are calculated using the data. We use the inverse of the effective measurement map from the calibration stage to calculate the empirical estimators. We give a more detailed description of the individual steps of the procedure in Appendix A 5.

Using an empirically estimated effective measurement map instead of the ideal theoretical result is the essential

modification compared to standard shadow estimation. In this way, we also “invert” the effect of the noise on our estimator. The right column of Fig. 7 shows an effective measurement map implemented with imperfect gates.

As a proof of concept, we choose the following simple but practically relevant setup. Random local Pauli measurements are implemented by native measurements in the computational basis after rotating with a Hadamard gate H (if the Pauli- X basis is to be measured) or a phase gate S followed by a Hadamard gate (for measurement in the Pauli- Y basis). Since throughout the protocol the S gate only turns up before application of the Hadamard gate, we treat the sequence HS as a single gate.

We assume that the dominant noise associated with the single-qubit rotations of each local gate in the experiment stays confined to two neighboring qubits. This assumption

makes both the gate set estimation and the postprocessing of the shadow estimation highly scalable.

Figure 7 shows the results of our scheme in simulations of the energy estimation of a Heisenberg Hamiltonian on a ten-qubit system. We observe that using the estimated effective measurement map instead of the ideal theoretical one significantly reduces the relative error $|(\hat{E} - E)/E|$ between the estimated energy \hat{E} and the true energy E of a given state. There are two contributions to the relative error in a shadow estimation protocol: first, the statistical fluctuation from the randomness of both the Pauli basis selection and the single-shot measurements; second, the systematic bias introduced in the postprocessing due to imperfect implementations of the measurements. The histograms in Fig. 7 show the infinite measurement limit of the relative error and thus directly reflect the bias. We observe that, for a fixed noise model, the magnitude of the bias depends heavily on the selected initial state, with relative errors being distributed over 2 orders of magnitude. When comparing the most likely errors between standard shadow estimation and GST-mitigated shadow estimation, we find that using GST data leads to a reduction in relative error by half an order of magnitude for random pure states and an order of magnitude for eigenstates of the Hamiltonian. The simulation of the protocol in the top left of Fig. 7 includes statistical fluctuations and shows how estimates spread for different sample counts when the ground-state energy is estimated. We find that, from 10^4 samples on, the GST-mitigated protocol yields a significantly more accurate estimate.

VI. CONCLUSION AND OUTLOOK

We have revisited the data processing task of GST from a compressed sensing perspective, regarding it as a highly structured and constrained tensor completion problem. In this formulation, we can naturally require the reconstructed gate set to be physical and, moreover, of low rank. Compressive gate set tomography thus aims at extracting considerably fewer parameters of the gate set. At the same time we have argued that the low-rank approximation to the implementation of a gate set contains the most valuable information about experimental imperfections for the practitioner.

The set of Kraus operators of a low-rank gate can be regarded as isometries that make up the complex Stiefel manifold. This observation has motivated the solution of the compressive GST data processing problem via geometrical optimization on the respective product manifolds. We have devised the optimization algorithm mGST that performs an adapted saddle-free Newton method on the manifold. To this end, we have derived the Riemannian Newton equation, Hessian equation, and geodesic curves.

In numerical experiments we have studied the performance of the mGST algorithm. We have compared it to pyGSTi, the state-of-the-art approach to the GST data processing problem, in settings where both algorithms can be applied and using full rank mGST estimates. We have found that in these settings mGST matches the performance of pyGSTi, while offering a larger model space, more flexibility in the sequence design, and allowing for low-rank assumptions. Moreover, we have numerically demonstrated that making use of the low-rank constraints significantly reduces the required number of measured sequences and the runtime of the reconstruction algorithm for a standard single- and two-qubit model. Importantly, we have found that we can successfully reconstruct generic unitary channels and depolarizing noise of one- and two-qubit gate sets from *random* gate sequences. This reduces the demands of GST both for experiments and classical postprocessing: the data that compressive GST requires are virtually identical with the experimental data produced by randomized benchmarking experiments. The classical postprocessing of mGST for a low-rank reconstruction of two-qubit gate sets takes only minutes even on desktop hardware, compared to over an hour with pyGSTi. We expect that this speedup and the low number of sequences required can lift two-qubit GST from being a protocol that is unpractical in many situations to one that is routinely applied, thus enabling it to be used in the engineering cycle for the design and calibration of gate sets.

Furthermore, compressive GST makes it feasible to perform self-consistent tomography on three-qubit systems. Making use of often available prior knowledge about an initialization can further reduce the computing time. We demonstrated this by performing tomography of unitary errors for three-qubit gate sets, using only a small number of random gate sequences, and with the postprocessing still running on desktop hardware in a few hours. We expect that even going slightly beyond three qubits is feasible by simply using more computing power. We leave it to future work to further tweak the numerical implementation in order to improve the scalability of the classical postprocessing. We also expect that progressively longer sequences can be added at the end of our optimization method, much in the same fashion as in pyGSTi, in order to further improve reconstruction accuracy of gate set estimates.

To demonstrate the use of compressive GST for error mitigation, we have introduced one novel application where low-rank mGST reconstructions are used to alleviate the effect of coherent errors in classical shadow estimation. The protocol uses a set of gates to implement basis changes before the measurement. We have demonstrated that with tomographic information on these gates through two-qubit compressive GST, more accurate ground state energy estimates are obtained in practically

relevant regimes. This constitutes just one example where the full information of a low-rank GST estimate is used to correct errors, and we expect that the reduced runtime requirements of mGST enable frequent use of GST for error diagnosis and mitigation.

Finally, besides making GST more applicable and flexible in practice, our reformulation is motivated by bringing it closer to theoretical recovery guarantees quantifying a required and sufficient number of random sequences for accurate reconstruction. Regarding the data processing of GST as a translation-invariant matrix-product-state completion problem makes it more amenable to prove techniques from compressed sensing. For example, establishing local convergence guarantees for mGST would allow one to quantify the assumptions on the experimental implementation that justify certain initialization of the algorithm. We hope that our work can serve as a foundation and inspiration in the quest of establishing mathematically rigorous guarantees for GST.

ACKNOWLEDGMENTS

We thank Lennart Bittel for helpful discussions and hints regarding optimization methods and their implementation and Markus Heinrich for a discussion on inverses of positive maps on operator spaces. We are particularly thankful to Kenneth Rudinger and Robin Blume-Kohout for valuable feedback on a previous version of the manuscript and help with finding an error in our code. The work of R.B. and M.K. has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) within the Emmy Noether program (Grant No. 441423094) and by the German Federal Ministry of Education and Research (BMBF) within the funding program “quantum technologies—from basic research to market” via the joint project MIQRO (Grant No. 13N15522). I.R. acknowledges funding from the BMBF (DAQC) and the Einstein Foundation (Einstein Research Unit).

APPENDIX A

In this appendix, we provide the mathematical details required for the saddle-free Newton method within the Riemannian optimization framework; see Appendices A 1–A 3. Moreover, we compare the dependence of the mGST MVE on the choice of objective function (mean squared error versus maximum likelihood) in Appendix A 4.

1. Geodesics on the Stiefel manifold

Edelman *et al.* [48] derived the geodesic on the real Stiefel manifold by solving the respective geodesic equation. We now show that the simple generalization given in Eq. (27) is indeed the correct geodesic in the complex case. For a curve $K_t \equiv K(t)$, the general geodesic

equation is [89, Chapter 5.4, Proposition 5.3.2]

$$P_{T(K_t)}(\ddot{K}_t + C_{K_t}(\dot{K}_t, \dot{K}_t)) = 0, \quad (\text{A1})$$

where the Christoffel symbol C_{K_t} depends on the chosen metric. Here, we use the canonical metric

$$\langle \Delta_1, \Delta_2 \rangle_K = \text{Re}\{\text{Tr}(\Delta_1^\dagger \Gamma \Delta_2)\} =: g(\Delta_1, \Delta_2) \quad (\text{A2})$$

with $\Gamma = \mathbb{1} - \frac{1}{2}K_t K_t^\dagger$. Using the Einstein summation convention, the Christoffel symbol at K can be computed as

$$(C_{K_t}^k)_{ij} = \frac{1}{2}g_{kl}^{-1} \left(\frac{\partial g_{lj}}{\partial K_{ti}} + \frac{\partial g_{li}}{\partial K_{tj}} - \frac{\partial g_{ij}}{\partial K_{tl}} + \text{c.c.} \right), \quad (\text{A3})$$

where $C_{K_t}^k$ is the k th component of the Christoffel symbol at K_t with respect to a basis $\{E_k, E_k^*\}_{k \in [Dd]}$ on the ambient space $\mathbb{C}^{D \times d}$.

Lemma 1: *The geodesic equation on the complex Stiefel manifold $\text{St}(D, d)$ equipped with the canonical metric for the curve $K_t : \mathbb{R} \rightarrow \text{St}(D, d)$ is given by*

$$P_{T(K_t)}(\ddot{K}_t + \dot{K}_t \dot{K}_t^\dagger K_t - K_t \dot{K}_t^\dagger \dot{K}_t - \dot{K}_t K_t^\dagger \dot{K}_t) = 0. \quad (\text{A4})$$

Proof. By noting that $\Gamma^{-1} = \mathbb{1} + K_t K_t^\dagger$ we can determine the function $g^{-1}(\Delta_1, \Delta_2)$ via the condition $g(g^{-1}(\Delta_1, \cdot), \Delta_2) = \text{Tr}[\Delta_1^\dagger \Delta_2 + \Delta_2 \Delta_1^\dagger]$, meaning that the inverse g^{-1} would recover the standard symmetric inner product on $T_K \text{St}(D, d)$. One can quickly verify that $g^{-1}(\Delta_1, \cdot) = \Gamma^{-1} \Delta_1$ satisfies this condition.

We determine the derivatives of g needed for the Christoffel symbol by explicitly writing out g as

$$g(\Delta_1, \Delta_2) = \text{Tr} \left[\Delta_1^\dagger \left(\mathbb{1} - \frac{1}{2}K_t K_t^\dagger \right) \Delta_2 + \Delta_2^\dagger \left(\mathbb{1} - \frac{1}{2}K_t K_t^\dagger \right) \Delta_1 \right], \quad (\text{A5})$$

from where we can find the derivatives by K and K^* as

$$\frac{\partial g_{ij}}{\partial K_{tl}} = \frac{\partial g}{\partial K_{tl}}(E_i, E_j) = -\frac{1}{2} \text{Tr}[E_i^\dagger E_l K_t^\dagger E_j + E_j^\dagger E_l K_t^\dagger E_i], \quad (\text{A6})$$

$$\begin{aligned} \frac{\partial g}{\partial K_{tl}^*}(E_i, E_j) &= \left(\frac{\partial g}{\partial K_{tl}}(E_i, E_j) \right)^* \\ &= -\frac{1}{2} \text{Tr}[E_i^\dagger K_t E_l^\dagger E_j + E_j^\dagger K_t E_l^\dagger E_i]. \end{aligned} \quad (\text{A7})$$

With these derivatives we can calculate

$$\begin{aligned}
C_{K_t}^k(\dot{K}_t, \dot{K}_t) &= (C_{K_t}^k)_{ij}(\dot{K}_t)_i(\dot{K}_t)_j \\
&= \frac{1}{2} g_{kl}^{-1} \left(\frac{\partial g_{lj}}{\partial K_{ii}} + \frac{\partial g_{li}}{\partial K_{jj}} - \frac{\partial g_{ij}}{\partial K_{ll}} + \text{c.c.} \right) \dot{K}_{ii} \dot{K}_{jj} \\
&= \frac{1}{2} \left(\frac{\partial g}{\partial K_{ii}}(g^{-1}(E_k, \cdot), \dot{K}_t) \dot{K}_{ii} + \frac{\partial g}{\partial K_{jj}}(g^{-1}(E_k, \cdot), \dot{K}_t) \dot{K}_{jj} - \frac{\partial g}{\partial K_{ll}}(\dot{K}_t, \dot{K}_t)(g^{-1}(E_k, \cdot))_l + \text{c.c.} \right) \\
&= -\frac{1}{2} \text{Re}\{\text{Tr}[2E_k^\dagger \Gamma^{-1} \dot{K}_t K_t^\dagger \dot{K}_t + 2\dot{K}_t^\dagger \dot{K}_t K_t^\dagger \Gamma^{-1} E_k - 2\dot{K}_t^\dagger \Gamma^{-1} E_k K_t^\dagger \dot{K}_t]\} \\
&= -\text{Re}\{\text{Tr}[\dot{K}_t K_t \dot{K}_t^\dagger \Gamma^{-1} E_k + \dot{K}_t^\dagger \dot{K}_t K_t^\dagger \Gamma^{-1} E_k - K_t^\dagger \dot{K}_t \dot{K}_t^\dagger \Gamma^{-1} E_k]\} \\
&= \langle (K_t^\dagger \dot{K}_t \dot{K}_t^\dagger - \dot{K}_t^\dagger \dot{K}_t K_t^\dagger - \dot{K}_t^\dagger K_t \dot{K}_t^\dagger)^\dagger, E_k \rangle \\
&= \langle \dot{K}_t \dot{K}_t^\dagger K_t - K_t \dot{K}_t^\dagger \dot{K}_t - \dot{K}_t K_t^\dagger \dot{K}_t, E_k \rangle,
\end{aligned} \tag{A8}$$

where we have used the facts that $(\Gamma^{-1})^\dagger = \Gamma^{-1}$ and $\text{Re Tr}[X] = \text{Re Tr}[X^\dagger]$. We now first write out the geodesic equation (A1) on the ambient space,

$$\begin{aligned}
\langle \ddot{K}_t, E_k \rangle + C_{K_t}^k(\dot{K}_t, \dot{K}_t) &= \langle \ddot{K}_t, E_k \rangle + \langle \dot{K}_t \dot{K}_t^\dagger K_t - K_t \dot{K}_t^\dagger \dot{K}_t - \dot{K}_t K_t^\dagger \dot{K}_t, E_k \rangle \\
&= 0 \quad \text{for all } E_k,
\end{aligned} \tag{A9}$$

which is equivalent to

$$\ddot{K}_t + \dot{K}_t \dot{K}_t^\dagger K_t - K_t \dot{K}_t^\dagger \dot{K}_t - \dot{K}_t K_t^\dagger \dot{K}_t = 0. \tag{A10}$$

This completes the proof. ■

To arrive at the geodesic equation (27), it remains to project the above equation onto the tangent space. Indeed, with the explicit form of the geodesic equation from Lemma 1 we can show that the immediate generalization from the geodesic in the real case [48] gives a valid geodesic for the complex case.

Lemma 2: *The curve given by*

$$K_t = (K, Q) \exp \left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \begin{pmatrix} \mathbb{1} \\ 0 \end{pmatrix} \tag{A11}$$

is a geodesic on $\text{St}(D, d)$, determined through the initial conditions $K_{t=0} = K$ and $\dot{K}_{t=0} = \Delta$, with Q, R given by the QR decomposition of $(\mathbb{1} - KK^\dagger)\Delta$ and $A = K^\dagger \Delta$.

Proof. We recall that the ambient space splits into the tangent space and its orthogonal complement, the normal space. Therefore, the condition that the projection of the left-hand side onto the tangent space in Eq. (A4) vanishes is equivalent to demanding that it lies solely in the normal space. If it is in the normal space, K_t^\dagger applied from the left will yield a Hermitian matrix. We now show that this is indeed the case. For that, we first need to determine the first and second derivatives of K_t :

$$\begin{aligned}
\dot{K}_t &= (K, Q) \exp \left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \begin{pmatrix} \mathbb{1} \\ 0 \end{pmatrix} \\
&= \underbrace{K_t A}_{\dot{K}_1} + \underbrace{(K, Q) \exp \left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \begin{pmatrix} 0 \\ \mathbb{1} \end{pmatrix} R}_{\dot{K}_2},
\end{aligned} \tag{A12}$$

$$\begin{aligned}
 \ddot{K}_t &= (K, Q) \exp \left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix}^2 \begin{pmatrix} \mathbb{1} \\ 0 \end{pmatrix} \\
 &= \underbrace{K(t)(A^2 - R^\dagger R)}_{\ddot{K}_1} + \underbrace{(K, Q) \exp \left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \begin{pmatrix} 0 \\ \mathbb{1} \end{pmatrix}}_{\ddot{K}_2} RA.
 \end{aligned} \tag{A13}$$

We immediately see that $K_t^\dagger \ddot{K}_1 = A^2 - R^\dagger R$, which is Hermitian, as A is skew Hermitian. We will now show that $K_t^\dagger \ddot{K}_2 = 0$ starting with

$$\begin{aligned}
 K_t^\dagger \ddot{K}_2 &= (\mathbb{1}, 0) \exp \left(-t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \begin{pmatrix} K^\dagger \\ Q^\dagger \end{pmatrix} (K, Q) \exp \left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \begin{pmatrix} 0 \\ \mathbb{1} \end{pmatrix} RA \\
 &= (\mathbb{1}, 0) \exp \left(-t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \begin{pmatrix} \mathbb{1} & K^\dagger Q \\ Q^\dagger K & \mathbb{1} \end{pmatrix} \exp \left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \begin{pmatrix} 0 \\ \mathbb{1} \end{pmatrix} RA \\
 &= (\mathbb{1}, 0) \exp \left(-t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \left(\begin{pmatrix} \mathbb{1} & 0 \\ 0 & \mathbb{1} \end{pmatrix} + \begin{pmatrix} 0 & K^\dagger Q \\ Q^\dagger K & 0 \end{pmatrix} \right) \\
 &\quad \times \exp \left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \begin{pmatrix} 0 \\ \mathbb{1} \end{pmatrix} RA \\
 &= (\mathbb{1}, 0) \exp \left(-t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \begin{pmatrix} 0 & K^\dagger Q \\ Q^\dagger K & 0 \end{pmatrix} \exp \left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \begin{pmatrix} 0 \\ \mathbb{1} \end{pmatrix} RA.
 \end{aligned} \tag{A14}$$

To simplify the last expression, we set

$$\begin{pmatrix} U_{00} & U_{01} \\ U_{10} & U_{11} \end{pmatrix} = \exp \left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \tag{A15}$$

and obtain $K_t^\dagger \ddot{K}_2 = (U_{00}^\dagger K^\dagger Q U_{11} + U_{10}^\dagger Q^\dagger K U_{01}) RA$. From the series representation of the matrix exponential we gather that $U_{11} = \mathbb{1} + R \cdot X$ for some matrix X . Moreover, $U_{10} = R\tilde{X}$ and $U_{01}^\dagger = \tilde{X}^\dagger R^\dagger$ for some \tilde{X} , leading to

$$K_t^\dagger \ddot{K}_2 = (U_{00}^\dagger K^\dagger Q + U_{00}^\dagger K^\dagger Q R X + \tilde{X}^\dagger R^\dagger Q^\dagger K U_{01}) RA = 0, \tag{A16}$$

since $K^\dagger QR = K^\dagger (\mathbb{1} - KK^\dagger) \Delta = 0$.

This shows that \ddot{K}_t lies in the normal space, leaving us with the terms in the geodesic equation (A4) that depend only on \dot{K}_t :

$$\begin{aligned}
 &K_t^\dagger (\dot{K}_t \dot{K}_t^\dagger K_t - K_t \dot{K}_t^\dagger \dot{K}_t - \dot{K}_t K_t^\dagger \dot{K}_t) \\
 &= K_t^\dagger \dot{K}_t \dot{K}_t^\dagger K_t - \dot{K}_t^\dagger \dot{K}_t - (K_t^\dagger \dot{K}_t)^2 \\
 &= (A + K_t^\dagger \dot{K}_2)(A + K_t^\dagger \dot{K}_2)^\dagger - (A^\dagger A + A^\dagger K_t^\dagger \dot{K}_2 + \dot{K}_2^\dagger K_t A + \dot{K}_2^\dagger \dot{K}_2) - (A + K_t^\dagger \dot{K}_2)^2 \\
 &= AA^\dagger - A^\dagger A - \dot{K}_2^\dagger \dot{K}_2 - A^2.
 \end{aligned}$$

The last line follows from $\dot{K}_2 A = \ddot{K}_2$ and our previous observation that $K_t^\dagger \ddot{K}_2 = 0$, which implies that $K_t^\dagger \dot{K}_2 = 0$ as well. The remaining term $\dot{K}_2^\dagger \dot{K}_2$ can be computed similarly to $K_t^\dagger \ddot{K}_2$ and we obtain

$$\begin{aligned}
\dot{K}_2^\dagger \dot{K}_2 &= R^\dagger(0, \mathbb{1}) \exp\left(-t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix}\right) \begin{pmatrix} \mathbb{1} & K^\dagger Q \\ Q^\dagger K & \mathbb{1} \end{pmatrix} \exp\left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix}\right) \begin{pmatrix} 0 \\ \mathbb{1} \end{pmatrix} R \\
&= R^\dagger(0, \mathbb{1}) \exp\left(-t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix}\right) \left(\begin{pmatrix} \mathbb{1} & 0 \\ 0 & \mathbb{1} \end{pmatrix} + \begin{pmatrix} 0 & K^\dagger Q \\ Q^\dagger K & 0 \end{pmatrix} \right) \\
&\quad \times \exp\left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix}\right) \begin{pmatrix} 0 \\ \mathbb{1} \end{pmatrix} R \\
&= R^\dagger R + R^\dagger \begin{pmatrix} 0 & \mathbb{1} \end{pmatrix} \exp\left(-t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix}\right) \begin{pmatrix} 0 & K^\dagger Q \\ Q^\dagger K & 0 \end{pmatrix} \\
&\quad \times \exp\left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix}\right) \begin{pmatrix} 0 \\ \mathbb{1} \end{pmatrix} R \\
&= R^\dagger R + R^\dagger (U_{11}^\dagger Q^\dagger K U_{01} + U_{01}^\dagger K^\dagger Q U_{11}) R \\
&= R^\dagger R + R^\dagger ((\mathbb{1} + X^\dagger R^\dagger) Q^\dagger K U_{01} + U_{01}^\dagger K^\dagger Q (\mathbb{1} + R X)) R \\
&= R^\dagger R,
\end{aligned} \tag{A17}$$

where we have again used the fact that $K^\dagger Q R = R^\dagger Q^\dagger K = 0$ in the last line.

We can now put all the terms obtained by multiplying Eq. (A10) with K_t^\dagger from the left together to obtain

$$\begin{aligned}
K_t^\dagger (\ddot{K}_t + \dot{K}_t \dot{K}_t^\dagger K_t - K_t \dot{K}_t^\dagger \dot{K}_t - \dot{K}_t K_t^\dagger \dot{K}_t) &= A^2 - R^\dagger R - A A^\dagger - A^\dagger A - R^\dagger R - A^2 \\
&= -2R^\dagger R - A^\dagger A - A A^\dagger.
\end{aligned} \tag{A18}$$

We see that these remaining terms are Hermitian and therefore the left-hand side of Eq. (A10) is in the normal space. ■

2. Complex Newton equation

In this section, we derive the Riemannian Hessian operator and solve the Hessian equation to obtain an update direction on the tangent space, which we can follow along the geodesic defined in Eq. (27). This can be done for each gate individually, or simultaneously over all gates, in which case we operate on the Cartesian product $\text{St}(D, d)^{\times n}$ of single Stiefel manifolds. We consider the latter case, whereby we obtain the single Stiefel Newton equation [Eq. (A27)] as a byproduct. The method is based on the real case [48]. See also Ref. [49] for a recent treatment of second-order optimization on the complex Stiefel manifold, where instead of following geodesics, each optimization step is done in Euclidean space followed by a projection onto the manifold.

First let us make a general observation that will be useful at several points.

Lemma 3 (Theorem 14 of Ref. [49]): *Let $f : T_K \text{St}(D, d) \rightarrow \mathbb{C}$ be a \mathbb{C} -linear function, and let $\langle \cdot, \cdot \rangle_K$ be the canonical metric on $\text{St}(D, d)$ as defined in Eq. (24). Then the solution to*

$$\text{Re}\{f(\Delta)\} = \langle X, \Delta \rangle_K \quad \text{for all } \Delta \in T_K \text{St}(D, d) \tag{A19}$$

is given by $X = F^* - K F^T K \in T_K \text{St}(D, d)$, where F is chosen such that $f(\Delta) = \text{Tr}(F^T \Delta)$.

Proof. It is straightforward to see that $X \in T_K \text{St}(D, d)$ by applying the projector onto the tangent space: $P_{T_K}(X) = F^* - K F^T K - \frac{1}{2} K (K^\dagger F^* + F^T K) + \frac{1}{2} K (F^T K + K^\dagger F^*) = F^* - K F^T K$.

To show that X solves Eq. (3), we use the fact that $K^\dagger \Delta$ is skew Hermitian, as well as the fact that $\text{Re Tr}[HS] = 0$ for any skew Hermitian matrix S and Hermitian matrix H . Substituting $X = F^* - KF^TK$ into Eq. (3) we obtain

$$\begin{aligned} \langle X, \Delta \rangle_K &= \text{Re Tr} [X^\dagger (\mathbb{1} - \frac{1}{2}KK^\dagger) \Delta] \\ &= \text{Re Tr} [(F^T - K^\dagger F^* K^\dagger) (\mathbb{1} - \frac{1}{2}KK^\dagger) \Delta] \\ &= \text{Re Tr} [(F^T - \frac{1}{2}(F^T K + K^\dagger F^*) K^\dagger) \Delta] \\ &= \text{Re Tr} [F^T \Delta] - \text{Re Tr} [\text{herm}(F^T K) K^\dagger \Delta] \\ &= \text{Re Tr} [F^T \Delta]. \end{aligned}$$

This completes the proof. ■

Our goal is to simultaneously update all gates along the geodesic $\bigoplus_{i=1}^n \mathcal{K}_i(t) \in \text{St}(D, d)^{\times n}$, with the single Stiefel geodesics $\mathcal{K}_i(t)$ being given by Eq. (27). We define the initial directions as $\Delta_i = \dot{\mathcal{K}}_i(0)$. The first step to identify the Riemannian gradient and Hessian is to compute the second-order Taylor series expansion of \mathcal{L} in t at $t = 0$. Using $(\mathcal{K}_i)_{lm} \equiv \mathcal{K}_{ilm}$ and Einstein notation, we find that

$$\begin{aligned} &\mathcal{L}(\mathcal{K}_1 \oplus \dots \oplus \mathcal{K}_n; \mathcal{K}_1^* \oplus \dots \oplus \mathcal{K}_n^*) \\ &= \mathcal{L}|_{t=0} + 2 \text{Re} \left\{ \frac{\partial \mathcal{L}}{\partial \mathcal{K}_{ilm}} \frac{\partial \mathcal{K}_{ilm}}{\partial t} \right\} \Big|_{t=0} \cdot t \\ &\quad + 2 \text{Re} \left\{ \frac{\partial^2 \mathcal{L}}{\partial \mathcal{K}_{jop} \partial \mathcal{K}_{ilm}} \frac{\partial \mathcal{K}_{jop}}{\partial t} \frac{\partial \mathcal{K}_{ilm}}{\partial t} + \frac{\partial^2 \mathcal{L}}{\partial \mathcal{K}_{jop}^* \partial \mathcal{K}_{ilm}} \frac{\partial \mathcal{K}_{jop}^*}{\partial t} \frac{\partial \mathcal{K}_{ilm}}{\partial t} + \frac{\partial \mathcal{L}}{\partial \mathcal{K}_{ilm}} \frac{\partial^2 \mathcal{K}_{ilm}}{\partial t^2} \right\} \Big|_{t=0} \\ &\quad \times t^2/2 + \mathcal{O}(t^3). \end{aligned} \tag{A20}$$

We have $\partial \mathcal{K}_{ilm} / \partial t|_{t=0} = (\Delta_i)_{lm}$ and define $(\mathcal{L}_{\mathcal{K}_i})_{lm} := \partial \mathcal{L} / \partial \mathcal{K}_{ilm}$, so that we can write

$$\frac{\partial \mathcal{L}}{\partial \mathcal{K}_{ilm}} \frac{\partial \mathcal{K}_{ilm}}{\partial t} = \text{Tr}(\mathcal{L}_{\mathcal{K}_i}^T \Delta_i) \quad \text{and} \quad \frac{\partial \mathcal{L}}{\partial \mathcal{K}_{ilm}} \frac{\partial \mathcal{K}_{ilm}}{\partial t} =: \mathcal{L}_{\mathcal{K}_i}[\Delta_i].$$

In a similar fashion we define

$$\mathcal{L}_{\mathcal{K}_j \mathcal{K}_i}[\Delta_j, \Delta_i] := \frac{\partial^2 \mathcal{L}}{\partial \mathcal{K}_{jop} \partial \mathcal{K}_{ilm}} \frac{\partial \mathcal{K}_{jop}}{\partial t} \frac{\partial \mathcal{K}_{ilm}}{\partial t},$$

where $\mathcal{L}_{\mathcal{K}_j \mathcal{K}_i}[\cdot, \cdot]$ is a bilinear function, which is symmetric per definition via the second derivative. For more details on how to compute these derivatives for the objective function used in the main text, see Appendix A 3.

Before determining the relevant terms for the update on $\text{St}(D, d)^{\times n}$ we first consider the gradient and Hessian, as well as the Newton equation for a single variable $K \in \text{St}(D, d)$, leaving all others constant.

The Riemannian gradient $G \in T_K \text{St}(D, d)$ can be identified from the first-order term in the Taylor expansion via its definition [49]

$$2 * \text{Re}\{\mathcal{L}_K[\Delta]\} = \langle G, \Delta \rangle_K \quad \text{for all } \Delta \in T_K \text{St}(D, d). \tag{A21}$$

The solution for G in the canonical metric (24) is given by

$$G = 2(\mathcal{L}_K^* - K \mathcal{L}_K^T K), \tag{A22}$$

as per Lemma 3.

Lemma 4: *The Riemannian Hessian $\text{Hess} : T_K \text{St}(D, d) \times T_K \text{St}(D, d) \rightarrow \mathbb{R}$ of a function $\mathcal{L} : \text{St}(D, d) \rightarrow \mathbb{R}$ with respect to the canonical metric on $\text{St}(D, d)$ is given by*

$$\begin{aligned} \text{Hess}(\Delta, \Omega) &= 2 \text{Re}\{\mathcal{L}_{KK}[\Delta, \Omega] + \mathcal{L}_{K^*K}[\Delta^*, \Omega]\} \\ &\quad + \text{Re}\{\text{Tr}[\mathcal{L}_K^T (\Delta K^\dagger \Omega + \Omega K^\dagger \Delta)] - \text{Tr}[\mathcal{L}_K^T K (\Delta^\dagger \Pi \Omega + \Omega^\dagger \Pi \Delta)]\}. \end{aligned} \tag{A23}$$

Proof. According to Proposition 5.5.5 of Ref. [89], we can compute $\text{Hess}(\Delta, \Omega)$ via

$$\text{Hess}(\Delta, \Omega) = \frac{1}{2} \frac{d^2}{dt^2} [\mathcal{L}(K(t(\Delta + \Omega))) - \mathcal{L}(K(t\Delta)) - \mathcal{L}(K(t\Omega))] \Big|_{t=0}, \quad (\text{A24})$$

where $K(t\Delta)$ satisfies $\dot{K}(t\Delta)|_{t=0} = \Delta$ (see Ref. [48] for a discussion of the real case). The individual terms in Eq. (A24) can be determined from our general Taylor approximation in Eq. (A20), i.e., with $i = j = 1$, if we take $K = K_1$. The term

$$\frac{\partial \mathcal{L}}{\partial K} \left[\frac{\partial^2 K(t\Delta)}{\partial t^2} \right] \Big|_{t=0}$$

contains second derivatives of the geodesic given in Lemma 2, which we write out next. We have

$$\begin{aligned} \ddot{K}(t) &= (K, Q) \exp \left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix}^2 \begin{pmatrix} \mathbb{1} \\ 0 \end{pmatrix} \\ &= (K, Q) \exp \left(t \begin{pmatrix} A & -R^\dagger \\ R & 0 \end{pmatrix} \right) \begin{pmatrix} A^2 - R^\dagger R \\ RA \end{pmatrix}. \end{aligned}$$

It follows using $QR = (\mathbb{1} - KK^\dagger)\Delta =: \Pi\Delta$ and $A = K^\dagger\Delta$ from the definition of the geodesic that

$$\begin{aligned} \ddot{K}(0) &= K(A^2 - R^\dagger R) + QRA \\ &= K(A^2 - R^\dagger Q^\dagger QR) + \Pi\Delta K^\dagger\Delta \\ &= K(K^\dagger\Delta K^\dagger\Delta - \Delta^\dagger \Pi^\dagger \Pi\Delta) + \Pi\Delta K^\dagger\Delta \\ &= K(K^\dagger\Delta K^\dagger\Delta - \Delta^\dagger \Pi\Delta) + \Delta K^\dagger\Delta - KK^\dagger\Delta K^\dagger\Delta \\ &= \Delta K^\dagger\Delta - K\Delta^\dagger \Pi\Delta. \end{aligned}$$

Putting the terms together, we arrive at

$$\begin{aligned} \frac{d^2}{dt^2} \mathcal{L}(K(t\Delta)) \Big|_{t=0} &= 2 \text{Re}\{\mathcal{L}_{KK}[\Delta, \Delta] + \mathcal{L}_{K^*K}[\Delta^*, \Delta]\} \\ &\quad + \text{Re}\{\text{Tr}[\mathcal{L}_K^T(\Delta K^\dagger\Delta - K\Delta^\dagger \Pi\Delta)]\}. \end{aligned} \quad (\text{A25})$$

The terms involving \mathcal{L}_{KK} and \mathcal{L}_{K^*K} satisfy $\mathcal{L}_{KK}[\Delta, \Omega] = \mathcal{L}_{KK}[\Omega, \Delta]$ and $\mathcal{L}_{K^*K}[\Delta^*, \Omega] = \mathcal{L}_{K^*K}[\Omega^*, \Delta]$, by the symmetry of second derivatives. Using this symmetry property, we obtain the full Hessian (A24), which turns out to be

$$\begin{aligned} \text{Hess}(\Delta, \Omega) &= 2 \text{Re}\{\mathcal{L}_{KK}[\Delta, \Omega] + \mathcal{L}_{K^*K}[\Delta^*, \Omega]\} \\ &\quad + \text{Re}\{\text{Tr}[\mathcal{L}_K^T(\Delta K^\dagger\Omega + \Omega K^\dagger\Delta)] - \text{Tr}[\mathcal{L}_K^T K(\Delta^\dagger \Pi\Omega + \Omega^\dagger \Pi\Delta)]\}, \end{aligned} \quad (\text{A26})$$

where it is helpful to note that Eq. (A26) is related to Eq. (A25) via a symmetrization of the \mathcal{L}_K term. \blacksquare

Theorem 1: Let $\text{vec}(\Delta)$ be the row-major vectorization of $\Delta \in T_K \text{St}(D, d)$. Furthermore, let T and $\tilde{\mathcal{L}}_{KK}$ be defined by $T \text{vec}(X) = \text{vec}(X^T)$ and $\mathcal{L}_{KK}(\Delta, \cdot) = \tilde{\mathcal{L}}_{KK}^T \text{vec}(\Delta)$. Then the solution Δ of the linear equation in $\text{vec}(\Delta)$ and $\text{vec}(\Delta^*)$ given by

$$\begin{aligned} &(\tilde{\mathcal{L}}_{K^*K}^\dagger - (K \otimes K^T) T \tilde{\mathcal{L}}_{KK}^T - \frac{1}{2} \mathbb{1} \otimes (K^T \mathcal{L}_K) - \frac{1}{2} (K \mathcal{L}_K^T) \otimes \mathbb{1} - \frac{1}{2} \Pi \otimes (\mathcal{L}_K^\dagger K^*)) \text{vec}(\Delta) \\ &+ (\tilde{\mathcal{L}}_{KK}^\dagger - (K \otimes K^T) T \tilde{\mathcal{L}}_{K^*K}^T + \frac{1}{2} (\mathcal{L}_K^* \otimes K^T) T + \frac{1}{2} (K \otimes \mathcal{L}_K^\dagger) T) \text{vec}(\Delta^*) \\ &= -\frac{1}{2} \text{vec}(G) \end{aligned} \quad (\text{A27})$$

is the update direction along the geodesic given in Lemma 2 for the complex Newton method of a real function \mathcal{L} at position $K \in \text{St}(D, d)$.

Proof. The update direction Δ for the standard Newton method [48] is determined through the equation

$$\text{Hess}(\Delta, \Omega) = -\langle G, \Omega \rangle_K \quad \text{for all } \Omega \in T_K \text{St}(D, d), \quad (\text{A28})$$

which can be solved by rewriting the left-hand side as $\text{Hess}(\Delta, \Omega) = \langle f(\Delta), \Omega \rangle_K$ (for some yet to be determined f) and setting $\Omega = P_T(X)$ with arbitrary matrix X . This leads to

$$\langle f(\Delta), P_T(X) \rangle_K = -\langle G, P_T(X) \rangle_K, \quad (\text{A29})$$

$$\langle P_T(f(\Delta)), X \rangle_K = -\langle G, X \rangle_K \quad \text{for all } X \in \mathbb{C}^{D \times d}; \quad (\text{A30})$$

in the second line the scalar product is extended from the canonical scalar product initially defined on $T_K \text{St}(D, d)$ to $\mathbb{C}^{D \times d}$ and we use the same notation for both. The second line follows from the fact that any matrix X can be decomposed as $X = P_T(X) + P_N(X)$ and from $\langle A, B \rangle_K = 0$ for $A \in T_K \text{St}(D, d)$ and $B \in N_K \text{St}(D, d)$. To determine $f(\Delta)$, we split it into three terms $f(\Delta) = f_{KK}(\Delta) + f_{K^*K}(\Delta) + f_K(\Delta)$, where $f_{KK}(\Delta)$, $f_{K^*K}(\Delta)$, and $f_K(\Delta)$ depend only on \mathcal{L}_{KK} , \mathcal{L}_{K^*K} , and \mathcal{L}_K , respectively [compare Eq. (A26)].

We first look at the term $2 \text{Re}\{\mathcal{L}_{KK}[\Delta, \Omega]\} = 2 \text{Re}\{\text{Tr}(\mathcal{L}_{KK}[\Delta, \cdot]^T \Omega)\}$, where $\mathcal{L}_{KK}[\Delta, \cdot]$ is in $\mathbb{C}^{D \times d}$.

To solve $2 \text{Re}\{\text{Tr}(\mathcal{L}_{KK}[\Delta, \cdot]^T \Omega)\} = \langle f_{KK}(\Delta), \Omega \rangle_K$ for all $\Omega \in T_K \text{St}(D, d)$ and to find f_{KK} , we use Lemma 3 and obtain

$$f_{KK}(\Delta) = 2(\mathcal{L}_{KK}[\Delta, \cdot]^* - K \mathcal{L}_{KK}[\Delta, \cdot]^T K). \quad (\text{A31})$$

The same argument can be made for the \mathcal{L}_{K^*K} term, leading to

$$f_{K^*K}(\Delta^*) = 2(\mathcal{L}_{K^*K}[\Delta^*, \cdot]^* - K \mathcal{L}_{K^*K}[\Delta^*, \cdot]^T K). \quad (\text{A32})$$

To identify $f_K(\Delta)$, we rewrite the second line in Eq. (A26) as

$$\begin{aligned} & \text{Re}\{\text{Tr}[\mathcal{L}_K^T(\Delta K^\dagger \Omega + \Omega K^\dagger \Delta)] - \text{Tr}[\mathcal{L}_K^T K(\Delta^\dagger \Pi \Omega + \Omega^\dagger \Pi \Delta)]\} \\ &= \text{Re}\{\text{Tr}[(\mathcal{L}_K^T \Delta K^\dagger + K^\dagger \Delta \mathcal{L}_K^T - \mathcal{L}_K^T K \Delta^\dagger \Pi - (\Pi \Delta \mathcal{L}_K^T K)^\dagger) \Omega]\} \\ &\stackrel{!}{=} \text{Re}\{\text{Tr}[f_K(\Delta)^\dagger \Gamma \Omega]\}, \end{aligned}$$

where we have used the fact that $\text{Re}\{\text{Tr}[AB^\dagger]\} = \text{Re}\{\text{Tr}[A^\dagger B]\}$. Thus, we find that

$$\begin{aligned} f_K(\Delta) &= [(\mathcal{L}_K^T \Delta K^\dagger + K^\dagger \Delta \mathcal{L}_K^T - \mathcal{L}_K^T K \Delta^\dagger \Pi - (\Pi \Delta \mathcal{L}_K^T K)^\dagger) \Gamma^{-1}]^\dagger \\ &= 2K \Delta^\dagger \mathcal{L}_K^* + \Gamma^{-1} \mathcal{L}_K^* \Delta^\dagger K - \Pi \Delta K^\dagger \mathcal{L}_K^* - \Pi \Delta \mathcal{L}_K^T K \end{aligned}$$

by using $\Pi = \Pi^\dagger$, $\Pi \Gamma^{-1} = \Gamma^{-1} \Pi = \Pi$, $K^\dagger \Gamma^{-1} = 2K^\dagger$.

Equation (A30) implies that $P_T(f(\Delta)) = -G$, and it remains to compute $P_T(f_K(\Delta))$, as $P_T(f_{KK}(\Delta)) = f_{KK}(\Delta)$ and $P_T(f_{K^*K}(\Delta^*)) = f_{K^*K}(\Delta^*)$ (see Lemma 3). After a straightforward computation using $\Gamma^{-1} = \mathbb{1} + KK^\dagger$, $\Pi \Gamma^{-1} = \Pi$, $P_T(\Pi Z) = P_T(Z)$, as well as $K^\dagger \Pi = \Pi K = 0$, we get

$$P_T(f_K(\Delta)) = -\Pi \Delta K^\dagger \mathcal{L}_K^* - 2 \text{skew}(\Delta \mathcal{L}_K^T) K - 2K \text{skew}(\mathcal{L}_K^T \Delta) \quad (\text{A33})$$

with $\text{skew}(A) = (A - A^\dagger)/2$.

Finally, by substituting Eqs. (A31)–(A33) into Eq. (A30), we obtain the Newton equation

$$\begin{aligned} & \mathcal{L}_{KK}[\Delta, \cdot]^* - K \mathcal{L}_{KK}[\Delta, \cdot]^T K + \mathcal{L}_{K^*K}[\Delta^*, \cdot]^* - K \mathcal{L}_{K^*K}[\Delta^*, \cdot]^T K \\ & - \frac{1}{2} \Pi \Delta K^\dagger \mathcal{L}_K^* - \text{skew}(\Delta \mathcal{L}_K^T) K - K \text{skew}(\mathcal{L}_K^T \Delta) \\ &= -G/2. \end{aligned} \quad (\text{A34})$$

This is a linear equation in Δ and Δ^* that can be solved via rewriting it as an equation in $\text{vec}(\Delta)$. Using row-major vectorization with $\text{vec}(AXB) = (A \otimes B^T) \text{vec}(X)$, and matrices T and $\tilde{\mathcal{L}}_{KK}$ defined as $T \text{vec}(X) = \text{vec}(X^T)$ and

$\mathcal{L}_{KK}(\Delta, \cdot) = \tilde{\mathcal{L}}_{KK}^T \text{vec}(\Delta)$, we arrive at the final equation for the single gate case

$$\begin{aligned} & (\tilde{\mathcal{L}}_{K^*K}^\dagger - (K \otimes K^T)T\tilde{\mathcal{L}}_{KK}^T - \frac{1}{2}\mathbb{1} \otimes (K^T\mathcal{L}_K) - \frac{1}{2}(K\mathcal{L}_K^T) \otimes \mathbb{1} - \frac{1}{2}\Pi \otimes (\mathcal{L}_K^\dagger K^*)) \text{vec}(\Delta) \\ & + (\tilde{\mathcal{L}}_{KK}^\dagger - (K \otimes K^T)T\tilde{\mathcal{L}}_{K^*K}^T + \frac{1}{2}(\mathcal{L}_K^* \otimes K^T)T + \frac{1}{2}(K \otimes \mathcal{L}_K^\dagger)T) \text{vec}(\Delta^*) \\ & = -\frac{1}{2} \text{vec}(G). \end{aligned}$$

This completes the proof. ■

Now, for the simultaneous optimization over all gates on $\text{St}(D, d)^{\times n}$, the Hessian as defined in Eq. (A24) is determined by including all terms in Eq. (A20). The Newton equation reads

$$\text{Hess}(\Delta_1 \oplus \cdots \oplus \Delta_n; \Omega_1 \oplus \cdots \oplus \Omega_n) = - \sum_{i=1}^n \langle G_i, \Omega_i \rangle_{\mathcal{K}_i} \quad (\text{A35})$$

for all $\Omega_i \in T_{\mathcal{K}_i} \text{St}(D, d)$. The terms in Eq. (A20) where $i = j$ are obtained from the single variable case. The mixed variable terms $f_{\mathcal{K}_i\mathcal{K}_j}(\Delta_i)$ and $f_{\mathcal{K}_i^*\mathcal{K}_j}(\Delta_i^*)$ still need to be determined. Analogous to Eq. (A32) we need to solve

$$\begin{aligned} & 2 \text{Re}\{\text{Tr}(\mathcal{L}_{\mathcal{K}_i^*\mathcal{K}_j}[\Delta_i^*, \cdot]^T \Omega_j)\} = \langle f_{\mathcal{K}_i^*\mathcal{K}_j}(\Delta_i), \Omega_j \rangle_{\mathcal{K}_j} \quad \text{for all } \Omega_j \in T_{\mathcal{K}_j} \text{St} \\ \text{and } & 2 \text{Re}\{\text{Tr}(\mathcal{L}_{\mathcal{K}_i\mathcal{K}_j}[\Delta_i, \cdot]^T \Omega_j)\} = \langle f_{\mathcal{K}_i\mathcal{K}_j}(\Delta_i), \Omega_j \rangle_{\mathcal{K}_j} \quad \text{for all } \Omega_j \in T_{\mathcal{K}_j} \text{St}. \end{aligned}$$

We can use Lemma 3 again and obtain

$$\begin{aligned} & f_{\mathcal{K}_i^*\mathcal{K}_j}(\Delta_i) = 2(\mathcal{L}_{\mathcal{K}_i^*\mathcal{K}_j}^*[\Delta_i, \cdot] - \mathcal{K}_j \mathcal{L}_{\mathcal{K}_i^*\mathcal{K}_j}[\Delta_i^*, \cdot]^T \mathcal{K}_j) \\ \text{and } & f_{\mathcal{K}_i\mathcal{K}_j}(\Delta_i) = 2(\mathcal{L}_{\mathcal{K}_i\mathcal{K}_j}^*[\Delta_i^*, \cdot] - \mathcal{K}_j \mathcal{L}_{\mathcal{K}_i\mathcal{K}_j}[\Delta_i, \cdot]^T \mathcal{K}_j), \end{aligned}$$

which satisfy $f_{\mathcal{K}_i^*\mathcal{K}_j}(\Delta_i^*) \in T_{\mathcal{K}_j} \text{St}(D, d)$ and $f_{\mathcal{K}_i\mathcal{K}_j}(\Delta_i) \in T_{\mathcal{K}_j} \text{St}(D, d)$. The full Newton equation on $\text{St}(D, d)^{\times n}$ in vectorized form then reads

$$\begin{aligned} & \bigoplus_{i=1}^n \left(\tilde{\mathcal{L}}_{\mathcal{K}_i^*\mathcal{K}_i}^\dagger - (\mathcal{K}_i \otimes \mathcal{K}_i^T)T\tilde{\mathcal{L}}_{\mathcal{K}_i\mathcal{K}_i}^T - \frac{1}{2}\mathbb{1} \otimes (\mathcal{K}_i^T\mathcal{L}_{\mathcal{K}_i}) - \frac{1}{2}(\mathcal{K}_i\mathcal{L}_{\mathcal{K}_i}^T) \otimes \mathbb{1} - \frac{1}{2}\Pi \otimes (\mathcal{L}_{\mathcal{K}_i}^\dagger \mathcal{K}_i^*) \right) \text{vec}(\Delta_i) \\ & + \bigoplus_{i=1}^n \sum_{j:j \neq i} (\tilde{\mathcal{L}}_{\mathcal{K}_j^*\mathcal{K}_i}^\dagger - (\mathcal{K}_i \otimes \mathcal{K}_i^T)T\tilde{\mathcal{L}}_{\mathcal{K}_j\mathcal{K}_i}^T) \text{vec}(\Delta_j) \\ & + \bigoplus_{i=1}^n \left(\tilde{\mathcal{L}}_{\mathcal{K}_i\mathcal{K}_i}^\dagger - (\mathcal{K}_i \otimes \mathcal{K}_i^T)T\tilde{\mathcal{L}}_{\mathcal{K}_i^*\mathcal{K}_i}^T + \frac{1}{2}(\mathcal{L}_{\mathcal{K}_i}^* \otimes \mathcal{K}_i^T)T + \frac{1}{2}(\mathcal{K}_i \otimes \mathcal{L}_{\mathcal{K}_i}^\dagger)T \right) \text{vec}(\Delta_i^*) \\ & + \bigoplus_{i=1}^n \sum_{j:j \neq i} (\tilde{\mathcal{L}}_{\mathcal{K}_j\mathcal{K}_i}^\dagger - (\mathcal{K}_i \otimes \mathcal{K}_i^T)T\tilde{\mathcal{L}}_{\mathcal{K}_j^*\mathcal{K}_i}^T) \text{vec}(\Delta_j^*) \\ & = -\frac{1}{2} \bigoplus_{i=1}^n \text{vec}(G_i). \end{aligned} \quad (\text{A36})$$

We are now faced with an equation of the type $A\mathbf{x} + B\mathbf{x}^* = \mathbf{c}$, a solution to which can be obtained by solving

$$\begin{pmatrix} A & B \\ B^* & A^* \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{x}^* \end{pmatrix} = \begin{pmatrix} \mathbf{c} \\ \mathbf{c}^* \end{pmatrix}. \quad (\text{A37})$$

Equation (A36) is an equation on the tangent space and can be solved by finding a basis therein. However, in order to avoid a basis change at every step, we choose to solve it on the ambient space by setting $\Delta_i = P_{T_i}(\Delta_i)$ and $\Delta_i^* = P_{T_i}^*(\Delta_i^*)$.

The matrix equation for the update directions Δ_i is now given by

$$\underbrace{\begin{pmatrix} H_{G \leftarrow \Delta} \oplus_i P_{T_i} & H_{G \leftarrow \Delta^*} \oplus_i P_{T_i}^* \\ H_{G \leftarrow \Delta^*}^* \oplus_i P_{T_i}^* & H_{G \leftarrow \Delta}^* \oplus_i P_{T_i} \end{pmatrix}}_{=:H} \begin{pmatrix} \text{vec}(\Delta_1) \\ \dots \\ \text{vec}(\Delta_n) \\ \text{vec}(\Delta_1^*) \\ \dots \\ \text{vec}(\Delta_n^*) \end{pmatrix} = -\frac{1}{2} \begin{pmatrix} \text{vec}(G_1) \\ \dots \\ \text{vec}(G_n) \\ \text{vec}(G_1^*) \\ \dots \\ \text{vec}(G_n^*) \end{pmatrix}, \quad (\text{A38})$$

where the submatrices $H_{G \leftarrow \Delta}$ and $H_{G \leftarrow \Delta^*}$ can be identified from Eq. (A36).

The update directions for the saddle-free Newton method (Algorithm 1) are calculated by applying

$$(|H| + \lambda \mathbb{1})^{-1} \quad (\text{A39})$$

to the right-hand side of Eq. (A38). In practice, we use $\frac{1}{2}(H + H^\dagger)$ as the Hessian, since there exist more efficient methods for diagonalizing a Hermitian matrix compared to an arbitrary matrix.

3. Complex Euclidean gradient and Hessian

To compute the Riemannian Hessians for optimization on the Stiefel manifolds, the complex Euclidean gradients and Euclidean Hessians are needed. Here we go into more detail on their derivations for the least-squares objective function. In this section we also detail how the terms $\mathcal{L}_{\mathcal{K}\mathcal{K}}$ and $\mathcal{L}_{\mathcal{K}^*\mathcal{K}}$ and their conjugates from Appendix A 2 are calculated.

As the tensor network whose contraction yields $p_{j|i}$ in the cost function (20) is parameterized in terms of matrix variables and their conjugates, we use Wirtinger calculus for the derivatives and treat the conjugate variables as independent. A method for finding all the relevant terms in the Hessian for a scalar function of complex matrix variables is outlined in, e.g., Ref. [102], and we summarize it in the following. See also Ref. [103] for a short derivation of the complex derivative and Hessian in the context of optimization in complex Euclidean space.

Our objective function is in general not analytic, as is the case with real-valued functions of complex variables. This can be seen for the simplest case with one unitary gate U and $\rho = E = |0\rangle\langle 0|$, where $\mathcal{L} = |\langle 0|U|0\rangle|^4 = |U_{00}|^4 = (U_{00} * U_{00}^*)^2$. However, the derivatives with respect to the real and imaginary parts of the matrix variables exist and one can define formal derivatives for $f : \mathbb{C}^{M \times N} \times \mathbb{C}^{M \times N} \rightarrow \mathbb{R}$ via

$$\begin{aligned} \frac{\partial f(Z, Z^*)}{\partial Z} &:= \frac{\partial f(Z, Z^*)}{\partial \text{Re}[Z]} - i \frac{\partial f(Z, Z^*)}{\partial \text{Im}[Z]}, \\ \frac{\partial f(Z, Z^*)}{\partial Z^*} &:= \frac{\partial f(Z, Z^*)}{\partial \text{Re}[Z]} + i \frac{\partial f(Z, Z^*)}{\partial \text{Im}[Z]}, \end{aligned}$$

where $\partial f(Z, Z^*)/\partial Z \in \mathbb{C}^{M \times N}$ with $[\partial f(Z, Z^*)/\partial Z]_{ij} = \partial f(Z, Z^*)/\partial Z_{ij}$. These formal derivatives have nice properties, for instance $\partial f(Z, Z^*)/\partial Z^*$ is the direction of maximum increase of f and $\partial f(Z, Z^*)/\partial Z^* = 0$ identifies a stationary point of f ; see, e.g., Theorems 3.2 and 3.4 of Ref. [102]. Furthermore, the product rule and the chain rule apply as they do for real-valued matrix variables.

As laid out in Ref. [102, Lemma 5.2], we can write the second-order Taylor series of f as

$$\begin{aligned} &f(Z + dZ, Z^* + dZ^*) \\ &= f(Z, Z^*) + \left(\frac{\partial}{\partial \text{vec}(Z)} f(Z, Z^*) \right) d \text{vec}(Z) + \left(\frac{\partial}{\partial \text{vec}(Z^*)} f(Z, Z^*) \right) d \text{vec}(Z^*) \\ &\quad + \frac{1}{2} [d \text{vec}^T(Z^*) d \text{vec}^T(Z)] \begin{bmatrix} f_{ZZ^*} & f_{Z^*Z^*} \\ f_{ZZ} & f_{Z^*Z} \end{bmatrix} \begin{bmatrix} d \text{vec}(Z) \\ d \text{vec}(Z^*) \end{bmatrix} + r(dZ, dZ^*), \end{aligned} \quad (\text{A40})$$

where the higher-order contribution $r(dZ, dZ^*)$ satisfies

$$\lim_{(dZ, dZ^*) \rightarrow 0} \frac{r(dZ, dZ^*)}{\|(dZ, dZ^*)\|_F^2} = 0. \quad (\text{A41})$$




FIG. 8. Tensor network representation of first derivative.

The second-order derivatives are defined via

$$f_{ZZ} = \frac{\partial}{\partial \text{vec}(Z)^T} \frac{\partial}{\partial \text{vec}(Z)} f(Z, Z^*, \dots; y),$$

and similarly for $f_{Z^*Z^*}$, f_{Z^*Z} , and f_{ZZ^*} . The vectorization is to be understood as joining together of indices in a fixed order. For instance, $\text{vec} : \mathbb{C}^{n \times d^2 \times d \times d} \rightarrow \mathbb{C}^{nd^4}$ vectorizes \mathcal{K} , where the individual d -dimensional legs are the matrix indices of the Kraus operators, and the d^2 index numbers the different Kraus operators. Note that $\text{vec}(\partial f(Z, Z^*, y)/\partial Z) = \partial f(Z, Z^*, y)/\partial \text{vec}(Z)$.

For the optimizations over A , \mathcal{K} , and B , we need the first and second derivatives of \mathcal{L} by the respective variables and their conjugates. Let $Z \in \{A, A^*, \mathcal{K}, \mathcal{K}^*, B, B^*\}$ and $Y \in \{Z, Z^*\}$. Then

$$\begin{aligned} \frac{\partial}{\partial Z} \mathcal{L}(Z, \dots; y) &= \frac{2}{m} \sum_{\mathbf{i}} (p_{\mathbf{i}}(Z, \dots) - y_{\mathbf{i}}) \frac{\partial p_{\mathbf{i}}}{\partial Z}, \\ \frac{\partial}{\partial Y} \frac{\partial}{\partial Z} \mathcal{L}(Z, Y, \dots; y) &= \frac{\partial}{\partial Y} \frac{2}{m} \sum_{\mathbf{i}} (p_{\mathbf{i}}(Z, Y, \dots) - y_{\mathbf{i}}) \frac{\partial p_{\mathbf{i}}}{\partial Z} \\ &= \frac{2}{m} \sum_{\mathbf{i}} \frac{\partial p_{\mathbf{i}}(Z, Y, \dots)}{\partial Y} \frac{\partial p_{\mathbf{i}}(Z, Y, \dots)}{\partial Z} + \frac{2}{m} \sum_{\mathbf{i}} (p_{\mathbf{i}}(Z, Y, \dots) - y_{\mathbf{i}}) \frac{\partial^2 p_{\mathbf{i}}(Z, Y, \dots)}{\partial Y \partial Z}, \end{aligned}$$

meaning that derivatives of the objective function reduce to the derivatives of the tensor p . Taking the derivative of a tensor network with respect to one of its constituent tensors can be easily done in the pictorial representation by removing the respective tensor. For instance, $\partial p_{\mathbf{i}}/\partial \mathcal{K}^*$ can be calculated as shown in Fig. 8, using the product rule. Care has to be taken for the order of open indices when removing a tensor. In practice, we do not calculate the full tensor $\partial p/\partial \mathcal{K}^*$ of size n^ℓ and only compute $\partial p_{\mathbf{i}}/\partial \mathcal{K}^*$ for $\mathbf{i} \in I$, since usually $|I| \ll n^\ell$.

4. Mean variation error dependence on the choice of objective function

A well-motivated alternative to the least-squares objective function defined in Eq. (20) is the likelihood function

$$L_I(A, \mathcal{K}, B|y) := \prod_{\mathbf{i} \in I} \prod_{j \in [n_E]} p_{j|\mathbf{i}}(A, \mathcal{K}, B)^{k_{j|\mathbf{i}}}, \quad (\text{A42})$$

where $y_{j|\mathbf{i}} = k_{j|\mathbf{i}}/m$ again denotes the relative number out of m times outcome j is measured for sequence \mathbf{i} ; see also Eq. (28). The likelihood function at a given model parametrization (A, \mathcal{K}, B) and for measurement results y is precisely the probability of observing y , given the model probabilities $p(A, \mathcal{K}, B)$. To simplify the optimization, often the logarithm of the likelihood function is chosen, since it shares the same maxima. This loglikelihood function, which is also used in the final optimization procedure of pyGSTi [18] is then given by

$$\log L_I(A, \mathcal{K}, B|y) := m \sum_{\mathbf{i} \in I} \sum_j y_{j|\mathbf{i}} \log[p_{j|\mathbf{i}}(A, \mathcal{K}, B)]. \quad (\text{A43})$$

In Fig. 9 we show the effects of augmenting mGST (which is by default run on the least-squares objective function for numerical reasons) with the loglikelihood function after a least-squares estimate is found. To do this, we use Algorithm 2 with the negative loglikelihood function (A43) as the objective function.

We observe that, for the XYI-gate set, which we use to compare mGST and pyGSTi in the main text, optimizing the loglikelihood function decreases the mean variation error on the data sequences as well as on all sequences of the same length. The improvement is stronger for fewer samples and becomes negligible at around 10^6 samples. Interestingly, loglikelihood optimization also improves the least-squares error on all sequences, at the cost of slightly increasing it on the data sequences. This indicates that in the sample count range of 10^2 – 10^6 , optimizing the loglikelihood function leads to less overfitting.

5. Noise mitigation of shadow estimation with GST characterization

In Sec. V we numerically demonstrated how the results of a low-rank GST experiment can be used to correct estimation protocols based on inverting an informationally complete POVM. Such protocols are recently referred to as shadow estimation [46]. We now give a short mathematical description of the method and explain how low-rank GST estimates can be included in a scalable way.

In the following, we use also bra-ket notation for the space of linear operators $L(\mathcal{H})$ and its dual space as defined by the canonical isomorphism induced by the Hilbert Schmidt inner product $(O|\rho) = \text{Tr}(O^\dagger \rho)$. The quantum channel of a unitary U is written by the corresponding calligraphic letter, e.g., $\mathcal{U}|\rho) \equiv |U\rho U^\dagger)$.

We consider the task of estimating the expectation value of multiple observables in an unknown quantum state that we can repeatedly prepare on a quantum device. Being able to measure an informationally complete POVM $\{\Pi_x\}$ on the state, one can construct an estimator for an observable O . *Informational completeness* is equivalent to, in mathematical terms, the POVM constituting a *frame* for $L(\mathcal{H})$, and the associated frame operator $\mathcal{M} = \sum_x |\Pi_x\rangle\langle\Pi_x|$ being invertible; see, e.g., Ref. [104]. We can calculate the canonical dual frame to the POVM as $|\tilde{\Pi}_x) = \mathcal{M}^{-1}|\Pi_x)$. By construction we have the frame duality relation

$$\sum_x |\tilde{\Pi}_x)\langle\Pi_x| = \text{Id}_{L(\mathcal{H})}. \quad (\text{A44})$$

Thus, for any state ρ ,

$$(O|\rho) = \sum_x (O|\tilde{\Pi}_x)\langle\Pi_x|\rho). \quad (\text{A45})$$

By Born's rule, repeated measurements of the POVM yield independent and identically distributed samples $\Omega = (x_1, \dots, x_m)$ from the distribution with density $p_\rho(x) = \langle\Pi_x|\rho)$. Given Ω , we can calculate the empirical mean estimator

$$\hat{o} = \frac{1}{|\Omega|} \sum_{x \in \Omega} (O|\tilde{\Pi}_x), \quad (\text{A46})$$

and, by Eq. (A45), $\mathbb{E}[\hat{o}] = (O|\rho)$.

The sequence of dual frame elements $(\tilde{\Pi}_{x_1}, \dots, \tilde{\Pi}_{x_m})$ given by the measured samples Ω has been called the *classical shadow* of ρ in Ref. [46].

A practical implementation of an informationally complete POVM on a digital quantum computer can be realized with measurements in randomly selected bases from a sufficiently large group. To be explicit, we consider the simplest and perhaps most well-known example: the measurement in a randomly chosen multiqubit Pauli basis. The POVM can be implemented by applying a random (different) local Clifford rotation on every qubit and measuring in the computational basis. For informational completeness, it is sufficient to choose the rotations uniformly from the set $C = \{\text{Id}, H, HS\}$, where H is the Hadamard gate and S the phase gate. In our notation, we consider POVM effects $\Pi_x = \Pi_{\mathbf{g}, \mathbf{b}} = \bigotimes_l \Pi_{g_l, b_l}$ indexed by $C^n \times \{0, 1\}^n$ that are the tensor products of the local POVM effects $\Pi_{g_l, b_l} = \frac{1}{3} g_l^\dagger |b_l\rangle\langle b_l| g_l$ with $g_l \in C$. Let $\{\hat{\sigma}_k \mid k \in \{0, 1, 2, 3\}\}$ denote the Pauli matrices normalized in the Frobenius norm. The frame operator is given by

$$\begin{aligned} 3^n \mathcal{M} &= \frac{1}{3^n} \left(\sum_{U \in \{\mathbb{1}, H, HS\}} \mathcal{U}^\dagger (|\hat{\sigma}_0\rangle\langle\hat{\sigma}_0| + |\hat{\sigma}_3\rangle\langle\hat{\sigma}_3|) \mathcal{U} \right)^{\otimes n} \\ &= \frac{1}{3^n} \left(3|\hat{\sigma}_0\rangle\langle\hat{\sigma}_0| + \sum_i |\hat{\sigma}_i\rangle\langle\hat{\sigma}_i| \right)^{\otimes n} \\ &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1/3 & 0 & 0 \\ 0 & 0 & 1/3 & 0 \\ 0 & 0 & 0 & 1/3 \end{pmatrix}^{\otimes n}, \end{aligned} \quad (\text{A47})$$

where the matrix in the last line is represented in the Pauli basis. Since \mathcal{M}^{-1} acts on qubit l as $\mathcal{M}_l^{-1}(X) = 3X - \text{Tr}(X)\mathbb{1}$ for any X , we find that $\tilde{\Pi}_{g_l, b_l} = \bigotimes_{i=1}^n (3g_l^\dagger |b_l\rangle\langle b_l| g_l - \mathbb{1})$ [46].

Huang *et al.* [46] showed that, when using random Pauli basis measurements, the variance of the mean estimator for estimating local observables does not scale with the system size. Using a median-of-means estimator to boost the confidence, Huang *et al.* [46] further established that the expectation value of M different k -local observables can be estimated to ϵ -additive precision from $\mathcal{O}(\log(M)4^k/\epsilon^2)$ state copies.

Experimental implementations of the POVM are prone to errors, effectively implementing a noisy POVM with effects Π_x^\dagger . In the envisioned implementation here, noise sources affect the implementation of the gates C and the noise induces a bias in the estimator for the observables. However, if the noise is characterized to some extent, we can correct the estimators for this bias. To this end, let \mathcal{M}^\dagger be the (half-sided) noisy frame operator $\mathcal{M}^\dagger =$

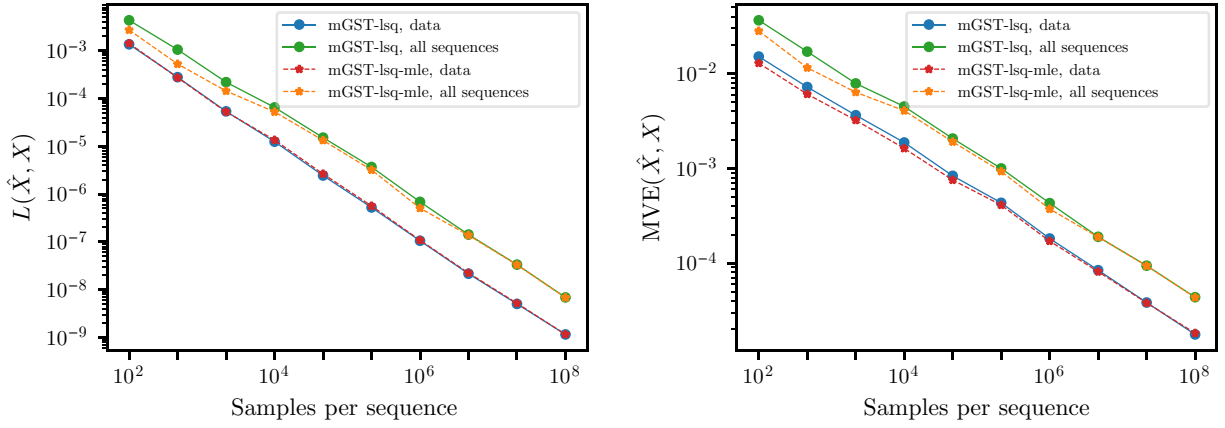


FIG. 9. Effect of optimizing the loglikelihood function after the least-squares objective function. The results of only the least-squares optimization are denoted by mGST-lsq and those of additional loglikelihood optimization by mGST-lsq-mle. The left plot shows the least-squares objective function $\mathcal{L}(\hat{\mathcal{X}}, \mathcal{X})$ on the data sequences (used for estimation), as well as on all sequences of a given length l (here $l = 7$). The right plot uses the MVE, again on data sequences or all sequences. The underlying gate set is given by the XYI model with depolarizing noise of strength $p = 0.01$ on each gate and $p = 0.01$ on the initial state, as well as random unitary rotations $e^{i\gamma H}$ with $H \sim \text{GUE}$ and $\gamma = 0.01$ on each gate. To ensure that mGST-lsq is fully converged, we set the desired relative precision to $\epsilon = 10^{-5}$ in the convergence criterion, see Eq. (30).

$\sum_x |\Pi_x\rangle\langle\Pi_x|$. If we know \mathcal{M}^\dagger in the classical postprocessing, we can calculate a dual frame to (Π_x^\dagger) by $|\tilde{\Pi}_x^\dagger\rangle = \mathcal{M}^{\dagger-1}|\Pi_x\rangle$. Note that using the “half-sided noisy” frame operator instead of the frame operator of the noisy POVM yields an expression of a dual frame in terms of the ideal POVM and not the noisy POVM. Using $\{|\tilde{\Pi}_x^\dagger\rangle\}$ instead of the ideal dual-frame in Eq. (A46) yields unbiased estimators of observables even in the presence of noise, thus effectively mitigating the noise.

This motivates our approach to noise-mitigated shadow estimation. Having extracted a noise model via gate set tomography, we can numerically estimate \mathcal{M}^\dagger and, thus, construct (approximately) unbiased estimators. Our method is summarized in Protocol 1 below.

The results in Sec. V demonstrate our scheme numerically in simple but already practically relevant settings that we describe in the following. When the multiqubit unitaries and computational basis measurement implementing the POVM factorize into local tensor products, so does the ideal frame operator \mathcal{M} and the dual frame (shadow). But due to correlated gate-dependent noise, \mathcal{M}^\dagger might not exhibit this computationally tractable structure. Furthermore, characterizing the implementation of exponentially many multiqubit unitaries and the basis measurements without additional assumptions is infeasible. In practice, however, noise-induced correlations and crosstalk might still predominantly affect a limited number of qubits simultaneously. For example, when noise predominantly affects neighboring qubits, we can use the implementation \mathcal{X} of a gate set including $C \times C$ on neighboring qubits extracted via mGST to calculate \mathcal{M}^\dagger . To this end, let $\mathcal{G}_{g_1, g_2}^{(i, i+1)}$ denote the two-qubit process implementing the gate $g_1 \times g_2$ on

qubits i and $i + 1$. For simplicity, we ignore errors in the computational basis measurement. We set $\Pi_{(g_1, g_2), (b_1, b_2)}^{i, (i+1)} = \frac{1}{9}(\mathcal{G}_{g_1, g_2}^{(i, i+1)})^\dagger |b_1, b_2\rangle\langle b_1, b_2| (\mathcal{G}_{g_1, g_2}^{(i, i+1)})$ and numerically calculate

$$\mathcal{M}_{i, i+1}^\dagger = \sum_{g_1, g_2 \in C, b_1, b_2 \in \{0, 1\}} |\Pi_{g_1, b_1}\rangle\langle\Pi_{g_2, b_2}| (\Pi_{(g_1, g_2), (b_1, b_2)}^{i, (i+1)}) |. \quad (\text{A48})$$

This amounts to calculating a 16×16 matrix in the Pauli basis that can be easily inverted. The noise-mitigated single-shot estimators thus read

$$\langle O | \tilde{\Pi}_{\mathbf{g}, \mathbf{b}}^\dagger \rangle = \langle O | \bigotimes_{i=1}^{n/2} (M_{2i, 2i+1}^\dagger)^{-1} |\Pi_{g_{2i}, b_{2i}}\rangle\langle\Pi_{g_{2i+1}, b_{2i+1}}|. \quad (\text{A49})$$

With this expression at hand, the rest of the protocol consists of computing the mean or median of means from a collection of single-shot estimators, following the standard method of shadow estimation [46].

Chen *et al.* [101] proposed a complimentary approach for robust shadow estimation, also inferring an approximation of the noisy frame operator from a separate calibration experiment. Under the assumption of gate-independent noise, the authors derived a 2^n parameter expression for \mathcal{M}^\dagger as a Pauli-noise channel and devised (SPAM-robust) RB-style experiments to learn arbitrarily many of its parameters, where each parameter corresponds to an irreducible representation of the local Clifford group. We find, in the gate-dependent noise model used here, that the frame operators significantly deviate from being a Pauli-noise

input: Target observable O , native local gate set \mathcal{G} with $C \subseteq \mathcal{G}$	
1 Perform 2-qubit-mGST on implementation of \mathcal{G} for qubit pairs $(1, 2), \dots, (n-1, n)$	} GST
2 Gauge optimize mGST estimators to unitary target gates	
3 for $i \in [N]$ do	} Classical Shadows data acquisition
4 Select setting $\mathbf{g} \in C^{\otimes n}$ uniformly at random	
5 Measure $\mathcal{U}_{\mathbf{g}} \rho\rangle$ in the standard basis	
6 Save setting \mathbf{g} and outcome \mathbf{b}	
7 end	
8 Construct $\mathcal{M}^{\mathbf{b}}$ from mGST gate estimates	} Combined post processing
9 Compute single shot estimators $\{\hat{o}_i = (O \mathcal{M}^{\mathbf{b}-1} \Pi_{\mathbf{g}, \mathbf{b}})\}_{i=1}^N$	
10 return $\hat{O} = \text{mean or median-of-means}(\{o_i\})$	

Protocol 1. GST-mitigated shadow estimation.

channel. For this reason, this particular setting is more amenable to GST-mitigated shadows than to the protocol of Ref. [101]. The plots on the left in Fig. 7 show that already a typical $n = 2$ frame operator with gate-dependent noise does not adhere to being diagonal in the Pauli basis.

Ultimately, we envision that different robust and self-consistent noise and error characterization protocols, such as mGST for local coherent errors and RB for incoherent noise strength in different irreducible representations, can be combined to arrive at accurate and scalable estimates of the effective frame operator in the presence of noise.

-
- [1] J. Eisert, D. Hangleiter, N. Walk, I. Roth, D. Markham, R. Parekh, U. Chabaud, and E. Kashefi, Quantum certification and benchmarking, *Nat. Rev. Phys.* **2**, 382 (2020).
 - [2] M. Kliesch and I. Roth, Theory of Quantum System Certification, *PRX Quantum* **2**, 010201 (2021).
 - [3] J. Emerson, R. Alicki, and K. Życzkowski, Scalable noise estimation with random unitary operators, *J. Opt. B* **7**, S347 (2005).
 - [4] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Randomized benchmarking of quantum gates, *Phys. Rev. A* **77**, eid 012307 (2008).
 - [5] E. Magesan, J. M. Gambetta, and J. Emerson, Characterizing quantum gates via randomized benchmarking, *Phys. Rev. A* **85**, 042311 (2012).
 - [6] J. Helsen, I. Roth, E. Onorati, A. H. Werner, and J. Eisert, A general framework for randomized benchmarking, *ArXiv:2010.07974* [quant-ph].
 - [7] S. Kimmel, M. P. da Silva, C. A. Ryan, B. R. Johnson, and T. Ohki, Robust Extraction of Tomographic Information via Randomized Benchmarking, *Phys. Rev. X* **4**, 011050 (2014).
 - [8] S. Kimmel and Y. K. Liu, in *2017 International Conference on Sampling Theory and Applications (SampTA)* (IEEE, Tallinn, Estonia, 2017), p. 345.
 - [9] I. Roth, R. Kueng, S. Kimmel, Y. K. Liu, D. Gross, J. Eisert, and M. Kliesch, Recovering Quantum Gates from Few Average Gate Fidelities, *Phys. Rev. Lett.* **121**, eid 170502 (2018).
 - [10] S. T. Flammia and J. J. Wallman, Efficient estimation of Pauli channels, *ACM Trans. Quantum Comput.* **1**, 1 (2020).
 - [11] S. T. Merkel, J. M. Gambetta, J. A. Smolin, S. Poletto, A. D. Córcoles, B. R. Johnson, C. A. Ryan, and M. Steffen, Self-consistent quantum process tomography, *Phys. Rev. A* **87**, eid 062119 (2013).
 - [12] C. Stark, Self-consistent tomography of the state-measurement Gram matrix, *Phys. Rev. A* **89**, 052109 (2014).
 - [13] C. Stark, Simultaneous estimation of dimension, states and measurements: Computation of representative density matrices and POVMs, *ArXiv:1210.1105*.
 - [14] R. Blume-Kohout, J. King Gamble, E. Nielsen, J. Mizrahi, J. D. Sterk, and P. Maunz, Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit, *ArXiv:1310.4492*.
 - [15] R. Blume-Kohout, J. K. Gamble, E. Nielsen, K. Rudinger, J. Mizrahi, K. Fortier, and P. Maunz, Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography, *Nat. Commun.* **8**, 14485 (2017).
 - [16] D. Greenbaum, Introduction to quantum gate set tomography, *ArXiv:1509.02921* (2015).
 - [17] E. Nielsen, K. Rudinger, T. Proctor, A. Russo, K. Young, and R. Blume-Kohout, Probing quantum processor performance with pyGSTi, *Quantum Sci. Technol.* **5**, eid 044002 (2020).
 - [18] E. Nielsen, J. K. Gamble, K. Rudinger, T. Scholten, K. Young, and R. Blume-Kohout, Gate set tomography, *Quantum* **5**, 557 (2021).
 - [19] J. P. Dehollain, J. T. Muhonen, R. Blume-Kohout, K. M. Rudinger, J. K. Gamble, E. Nielsen, A. Laucht, S. Simmons, R. Kalra, and A. S. Dzurak, *et al.*, Optimization of a solid-state electron spin qubit using gate set tomography, *New J. Phys.* **18**, 103018 (2016).
 - [20] K. Rudinger, C. W. Hogle, R. K. Naik, A. Hashim, D. Lobser, D. I. Santiago, M. D. Grace, E. Nielsen, T. Proctor, S. Seritan, S. M. Clark, R. Blume-Kohout, I. Siddiqi, and K. C. Young, Experimental Characterization of Crosstalk Errors with Simultaneous Gate Set Tomography, *PRX Quantum* **2**, eid 040338 (2021).
 - [21] C. Song, J. Cui, H. Wang, J. Hao, H. Feng, and Y. Li, Quantum computation with universal error mitigation

- on a superconducting quantum processor, *Sci. Adv.* **5**, eaaw5686 (2019).
- [22] S. Zhang, Y. Lu, K. Zhang, W. Chen, Y. Li, J.-N. Zhang, and K. Kim, Error-mitigated quantum gates exceeding physical fidelities in a trapped-ion system, *Nat. Commun.* **11**, eid 587 (2020).
- [23] L. Cincio, K. Rudinger, M. Sarovar, and P. J. Coles, Machine learning of noise-resilient quantum circuits, *ArXiv:2007.01210* (2020).
- [24] E. Knill, Quantum computing with realistically noisy devices, *Nature* **434**, 39 (2005).
- [25] P. Aliferis, D. Gottesman, and J. Preskill, Quantum accuracy threshold for concatenated distance-3 codes, *Quantum Inf. Comput.* **6**, 97 (2006).
- [26] A. W. Cross, D. P. DiVincenzo, and B. M. Terhal, A comparative code study for quantum fault-tolerance, *Quantum Inf. Comput.* **9**, 0541 (2009).
- [27] P. Aliferis and J. Preskill, Fibonacci scheme for fault-tolerant quantum computation, *Phys. Rev. A* **79**, eid 012332 (2009).
- [28] R. Kueng, D. M. Long, A. C. Doherty, and S. T. Flammia, Comparing Experiments to the Fault-Tolerance Threshold, *Phys. Rev. Lett.* **117**, 170502 (2016).
- [29] Y. R. Sanders, J. J. Wallman, and B. C. Sanders, Bounding quantum gate error rate based on reported average fidelity, *New J. Phys.* **18**, eid 012002 (2016).
- [30] J. J. Wallman, Bounding experimental quantum error rates relative to fault-tolerant thresholds, *ArXiv:1511.00727* (2015).
- [31] A. Shabani, R. L. Kosut, M. Mohseni, H. Rabitz, M. A. Broome, M. P. Almeida, A. Fedrizzi, and A. G. White, Efficient Measurement of Quantum Dynamics via Compressive Sensing, *Phys. Rev. Lett.* **106**, 100401 (2011).
- [32] S. T. Flammia, D. Gross, Y.-K. Liu, and J. Eisert, Quantum tomography via compressed sensing: Error bounds, sample complexity and efficient estimators, *New J. Phys.* **14**, 095022 (2012).
- [33] M. Kliesch, R. Kueng, J. Eisert, and D. Gross, Improving compressed sensing with the diamond norm, *IEEE Trans. Inf. Theory* **62**, 7445 (2016).
- [34] M. Kliesch, R. Kueng, J. Eisert, and D. Gross, Guaranteed recovery of quantum processes from few measurements, *Quantum* **3**, 171 (2019).
- [35] C. H. Baldwin, A. Kalev, and I. H. Deutsch, Quantum process tomography of unitary and near-unitary maps, *Phys. Rev. A* **90**, 012110 (2014).
- [36] R. Kueng, H. Rauhut, and U. Terstiege, Low rank matrix recovery from rank one measurements, *Appl. Comput. Harm. Anal.* **42**, 88 (2015).
- [37] A. V. Rodionov, A. Veitia, R. Barends, J. Kelly, D. Sank, J. Wenner, J. M. Martinis, R. L. Kosut, and A. N. Korotkov, Compressed sensing quantum process tomography for superconducting quantum gates, *Phys. Rev. B* **90**, eid 144504 (2014).
- [38] C. A. Riofrio, D. Gross, S. T. Flammia, T. Monz, D. Nigg, R. Blatt, and J. Eisert, Experimental quantum compressed sensing for a seven-qubit system, *Nat. Commun.* **8**, 15305 (2017).
- [39] A. Steffens, C. A. Riofrio, W. McCutcheon, I. Roth, B. A. Bell, A. McMillan, M. S. Tame, J. G. Rarity, and J. Eisert, Experimentally exploring compressed sensing quantum tomography, *Quantum Sci. Technol.* **2**, 025005 (2017).
- [40] A. Kyrillidis, A. Kalev, D. Park, S. Bhojanapalli, C. Caramanis, and S. Sanghavi, Provable compressed sensing quantum state tomography via non-convex methods, *npj Quantum Inf.* **4**, eid 36 (2018).
- [41] I. Roth, J. Wilkens, D. Hangleiter, and J. Eisert, Semi-device-dependent blind quantum tomography, *ArXiv:2006.03069* (2020).
- [42] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, Quantum State Tomography via Compressed Sensing, *Phys. Rev. Lett.* **105**, eid 150401 (2010).
- [43] D. Gross, Recovering low-rank matrices from few coefficients in any basis, *IEEE Trans. Inf. Theory* **57**, 1548 (2011).
- [44] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing* (Springer, Birkhäuser New York, 2013).
- [45] E. Nielsen, R. Blume-Kohout, L. Saldyt, J. Gross, T. L. Scholten, K. Rudinger, T. Proctor, J. K. Gamble, and A. Russo, Pygstio/pygsti: Version 0.9.9.3 (2020).
- [46] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [47] T. Abrudan, J. Eriksson, and V. Koivunen, Conjugate gradient algorithm for optimization under unitary matrix constraint, *Signal Process.* **89**, 1704 (2009).
- [48] A. Edelman, T. A. Arias, and S. T. Smith, The geometry of algorithms with orthogonality constraints, *SIAM J. Matrix Anal. Appl.* **20**, 303 (1998).
- [49] J. H. Manton, Optimization algorithms exploiting unitary constraints, *IEEE Trans. Signal Process.* **50**, 635 (2002).
- [50] Y. Sun, N. Flammarion, and M. Fazel, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., Vancouver, Canada, 2019), Vol. 32.
- [51] M. A. d. A. Bortoloti, T. A. Fernandes, O. P. Ferreira, and J. Yuan, Damped Newton's method on Riemannian manifolds, *J. Glob. Optim.* **77**, 643 (2020).
- [52] S. Wisdom, T. Powers, J. R. Hershey, J. Le Roux, and L. Atlas, in *Adv. Neural Inf. Process. Syst.* (Curran Associates, Inc., Barcelona, Spain, 2016), Vol. 29.
- [53] N. Boumal, P.-A. Absil, and C. Cartis, Global rates of convergence for nonconvex optimization on manifolds, *IMA J. Numer. Anal.* **39**, 1 (2019).
- [54] U. Helmke and J. B. Moore, *Optimization and Dynamical Systems* (Springer Science & Business Media, Springer London, 2012).
- [55] J. Manton, R. Mahony, and Y. Hua, The geometry of weighted low-rank approximations, *IEEE Trans. Signal Process.* **51**, 500 (2003).
- [56] L. Eldén and H. Park, A procrustes problem on the Stiefel manifold, *Numer. Math.* **82**, 599 (1999).
- [57] T. J. Bridges and S. Reich, Computing Lyapunov exponents on a Stiefel manifold, *Physica D* **156**, 219 (2001).
- [58] F. Lotte, L. Bougrain, A. Cichocki, M. Clerc, M. Congedo, A. Rakotomamonjy, and F. Yger, A review of classification algorithms for EEG-based brain-computer interfaces: A 10 year update, *J. Neural Eng.* **15**, 031005 (2018).

- [59] E. Chiumento and M. Melgaard, Stiefel and Grassmann manifolds in quantum chemistry, *J. Geom. Phys.* **62**, 1866 (2012).
- [60] X. Yu, J.-C. Shen, J. Zhang, and K. B. Letaief, Alternating minimization algorithms for hybrid precoding in millimeter wave MIMO systems, *IEEE J. Sel. Top. Signal Process.* **10**, 485 (2016).
- [61] F. Liu, C. Masouros, A. Li, H. Sun, and L. Hanzo, MU-MIMO communications with MIMO radar: from co-existence to joint transmission, *IEEE Trans. Wireless Commun.* **17**, 2755 (2018).
- [62] I. A. Luchnikov, M. E. Krechetov, and S. N. Filippov, Riemannian geometry and automatic differentiation for optimization problems of quantum physics and quantum technologies, *New J. Phys.* **23**, eid 073006 (2021).
- [63] D. Hangleiter, I. Roth, J. Eisert, and P. Roushan, Precise Hamiltonian identification of a superconducting quantum processor, *ArXiv:2108.08319*.
- [64] D. Hangleiter, I. Roth, D. Nagaj, and J. Eisert, Easing the Monte Carlo sign problem, *Sci. Adv.* **6**, eabb8341 (2020).
- [65] Y. N. Dauphin, R. Pascanu, C. Gulcehre, K. Cho, S. Ganguli, and Y. Bengio, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., Montréal, Canada, 2014), Vol. 27.
- [66] D. Süß, Ph.D. thesis, University of Cologne (2018).
- [67] M. Imaizumi, T. Maehara, and K. Hayashi, On tensor train rank minimization: Statistical efficiency and scalable algorithm, *Adv. Neural Inf. Process. Syst.* **30**, 3933 (2017).
- [68] H. Rauhut, R. Schneider, and Z. Stojanac, Low rank tensor recovery via iterative hard thresholding, *Linear Algebra Appl.* **523**, 220 (2017).
- [69] N. Ghadermarzy, Y. Plan, and Ö. Yilmaz, Near-optimal sample complexity for convex tensor completion, *Inf. Inference: J. IMA* **8**, 577 (2018).
- [70] M. Ashraphijuo and X. Wang, Characterization of deterministic and probabilistic sampling patterns for finite completability of low tensor-train rank tensor, *ArXiv:1703.07698*.
- [71] H. Rauhut and Ž. Stojanac, Tensor theta norms and low rank recovery, *Numer. Algorithms* **88**, 25 (2021).
- [72] B. Huang, C. Mu, D. Goldfarb, and J. Wright, Provable low-rank tensor recovery, *Optim.-Online* **4252**, 455 (2014).
- [73] C. Liu, H. Shan, and C. Chen, Tensor p-shrinkage nuclear norm for low-rank tensor completion, *Neurocomputing* **387**, 255 (2020).
- [74] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, *Nat. Phys.* **14**, 595 (2018).
- [75] J. Helsen, M. Ioannou, I. Roth, J. Kitzinger, E. Onorati, A. H. Werner, and J. Eisert, Estimating gate-set properties from random sequences, *ArXiv:2110.13178* (2021).
- [76] Y. Gu, R. Mishra, B.-G. Englert, and H. K. Ng, Randomized Linear Gate-Set Tomography, *PRX Quantum* **2**, 030328 (2021).
- [77] T. J. Evans, W. Huang, J. Yoneda, R. Harper, T. Tanttu, K. W. Chan, F. E. Hudson, K. M. Itoh, A. Saraiva, C. H. Yang, A. S. Dzurak, and S. D. Bartlett, Fast Bayesian Tomography of a Two-Qubit Gate Set in Silicon, *Phys. Rev. Appl.* **17**, eid 024068 (2022).
- [78] S. Montangero, *Introduction to Tensor Network Methods* (Springer, Cham, 2018).
- [79] M. Fannes, B. Nachtergaele, and R. Werner, Finitely correlated states on quantum spin chains, *Commun. Math. Phys.* **144**, 443 (1992).
- [80] S. Rommer and S. Östlund, Class of ansatz wave functions for one-dimensional spin systems and their relation to the density matrix renormalization group, *Phys. Rev. B* **55**, 2164 (1997).
- [81] F. Verstraete, D. Porras, and J. I. Cirac, Density Matrix Renormalization Group and Periodic Boundary Conditions: A Quantum Information Perspective, *Phys. Rev. Lett.* **93**, eid 227205 (2004).
- [82] I. V. Oseledets and E. E. Tyrtshnikov, Breaking the curse of dimensionality, or how to use SVD in many dimensions, *SIAM J. Sci. Comput.* **31**, 3744 (2009).
- [83] K. Kraus, *States, Effects and Operations: Fundamental Notions of Quantum Theory*, Lecture Notes in Physics (190. Springer, Berlin, 1983), Vol.
- [84] J. Lin, B. Buonacorsi, R. Laflamme, and J. J. Wallman, On the freedom in representing quantum operations, *New J. Phys.* **21**, eid 023006 (2019).
- [85] L. Rudnicki, Z. Puchala, and K. Zyczkowski, Gauge invariant information concerning quantum channels, *Quantum* **2**, 60 (2018).
- [86] T. Proctor, K. Rudinger, K. Young, M. Sarovar, and R. Blume-Kohout, What Randomized Benchmarking Actually Measures, *Phys. Rev. Lett.* **119**, eid 130502 (2017).
- [87] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011).
- [88] H. Schneider, Positive operators and an inertia theorem, *Numer. Math.* **7**, 11 (1965).
- [89] P. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds* (Princeton University Press, Princeton, 2009).
- [90] L. Grasedyck, M. Kluge, and S. Krämer, Variants of alternating least squares tensor completion in the tensor train format, *SIAM J. Sci. Comput.* **37**, A2424 (2015).
- [91] W. Wang, V. Aggarwal, and S. Aeron, Tensor completion by alternating minimization under the tensor train (TT) model, *ArXiv:1609.05587*.
- [92] C. Jin, P. Netrapalli, R. Ge, S. M. Kakade, and M. I. Jordan, On nonconvex optimization for machine learning: Gradients, stochasticity, and saddle points, *J. ACM* **68**, 1 (2021).
- [93] R. Brieger, I. Roth, and M. Kliesch, Python implementation of mGST, a compressive gate set tomography algorithm, <https://github.com/rabrie/mGST> (2021).
- [94] G. García-Pérez, M. A. C. Rossi, and S. Maniscalco, IBM Q Experience as a versatile experimental testbed for simulating open quantum systems, *npj Quantum Inf.* **6**, eid 1 (2020).
- [95] R. Vershynin, in *Compressed Sensing: Theory and Applications* (Cambridge University Press, Cambridge, 2012), p. 210.

- [96] L. Grasedyck and S. Krämer, Stable ALS approximation in the TT-format for rank-adaptive tensor completion, *Numer. Math.* **143**, 855 (2019).
- [97] A. Goeßmann, M. Götze, I. Roth, R. Sweke, G. Kutyniok, and J. Eisert, in *NeurIPS 2020—First Workshop on Quantum Tensor Networks in Machine Learning*, [ArXiv:2002.1238](https://arxiv.org/abs/2002.1238) (2020).
- [98] Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, J. R. McClean, and T. E. O’Brien, Quantum error mitigation, arXiv e-prints (2022).
- [99] S. Endo, Z. Cai, S. C. Benjamin, and X. Yuan, Hybrid quantum-classical algorithms and quantum error mitigation, *J. Phys. Soc. Jpn.* **90**, eid 032001 (2021).
- [100] D. E. Koh and S. Grewal, Classical shadows with noise, *Quantum* **6**, 776 (2022).
- [101] S. Chen, W. Yu, P. Zeng, and S. T. Flammia, Robust Shadow Estimation, *PRX Quantum* **2**, eid 030348 (2021).
- [102] A. Hjørungnes, *Complex-Valued Matrix Derivatives: With Applications in Signal Processing and Communications* (Cambridge University Press, USA, 2011), 1st ed.
- [103] A. Van Den Bos, Complex gradient and Hessian, *IEE Proc.-Vision Image Signal Process.* **141**, 380 (1994).
- [104] S. F. Waldron, *An Introduction to Finite Tight Frames* (Springer, Birkhäuser New York, 2018).