

Fast Quantum Process Tomography via Riemannian Gradient Descent

Daniel Volya, Andrey Nikitin and Prabhat Mishra
University of Florida, Gainesville, Florida, USA

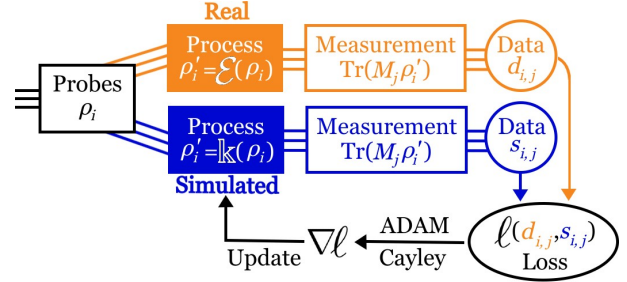
Abstract—Constrained optimization plays a crucial role in the fields of quantum physics and quantum information science and becomes especially challenging for high-dimensional complex structure problems. One specific issue is that of quantum process tomography, in which the goal is to retrieve the underlying quantum process based on a given set of measurement data. In this paper, we introduce a modified version of stochastic gradient descent on a Riemannian manifold that integrates recent advancements in numerical methods for Riemannian optimization. This approach inherently supports the physically driven constraints of a quantum process, takes advantage of state-of-the-art large-scale stochastic objective optimization, and has superior performance to traditional approaches such as maximum likelihood estimation and projected least squares. The data-driven approach enables accurate, order-of-magnitude faster results, and works with incomplete data. We demonstrate our approach on simulations of quantum processes and in hardware by characterizing an engineered process on quantum computers.

Index Terms—Quantum process tomography, quantum characterization, quantum computing, Riemannian optimization.


I. INTRODUCTION

Quantum computers present a theoretical exponential speedup over classical computers for various problems. However, the capability of these quantum computers is largely limited by environmental noise. Numerous quantum devices experience undesired interactions with internal and external degrees-of-freedom, resulting in decoherence [1], [2]. Although there are promising algorithmic approaches to preserve coherence using quantum error correcting codes, their effectiveness further depends on appropriate understanding and formulation of the underlying noise [3], [4]. Thus, accurately and effectively describing the noise in current quantum computing hardware continues to be a crucial component for achieving their long-term technological potential.

Certain noise can be characterized as a quantum channel, which is a completely positive and trace-preserving mapping between quantum states. Quantum channels are often represented as a set of matrices that transform a quantum state ρ to another state ρ' . As shown in Fig. 1a, given the initial and resulting states, it is possible to reconstruct the quantum channel through tomography. Tomography refers to the procedure of creating a model by merging various partial cross-sections or slices [5]. Although each slice provides a limited viewpoint, the data obtained from numerous independent experiments allows for the construction of a comprehensive model for the system. This paper specifically focuses on quantum process tomography (QPT), but it is worth mentioning that the same approach can be applied to measurement tomography [6], [7], Hamiltonian tomography [8], [9], and gate-set tomography [10], [11].



(a) Quantum process tomography estimates the blackbox process via measurement data $d_{i,j}$. The data is the result of conducting measurements M_j on states produced by a channel \mathcal{E} acting on known probe states ρ_i .



(b) We estimate a quantum process via gradient descent of an objective loss function $\ell(\cdot)$ on a Riemannian manifold, defined as a set of Kraus operators \mathbb{K} . We adapt an iterative Cayley transformation to perform a retraction of the gradient $\nabla \ell$, approximating the exponential map of a tangent vector onto the manifold.

Fig. 1: Overview of proposed quantum process tomography.

There are promising approaches for QPT, including maximum-likelihood estimation [12], [13], Bayesian estimation [14], [15], projected gradient descent [16], projected least squares [17], [18], compressed sensing through convex optimization [19]–[21], variational QPT [22], and trace regression models [23]. The key challenge, however, is the exponentially growing size of the process representation, e.g., the Choi representation is a complex-valued $4^n \times 4^n$ matrix for n qubits [24]. This makes it difficult to not only estimate the process from noisy data, but also to derive the source of noise and its interpretation. Furthermore, to ensure high-fidelity results, the informationally-complete set (ICS) of measurements and input states are required, which also grows exponentially in size. As a result, all experiments so far have predominantly applied QPT to a maximum of 3 qubits [25]. Recently, methods such as ShadowQPT have emerged, which relax these constraints by approximating the quantum channel with classical shadows.

Although these methods have successfully been performed on 4-qubit systems [26], their solution lacks the properties of a quantum channel, making them unphysical.

To address the scalability concerns, we utilize a strategy similar to deep learning, which has proven to be successful in various areas of physics [27]–[30]. Deep learning involves optimizing complex computational models, such as neural networks, tensor networks [31], [32], and differential equation solvers [33], using gradient-based methods. Automatic differentiation enables the computational process to be differentiable with respect to its constituent elements, such as a hidden layer in deep learning. As a result, the computational model can be fine-tuned with automatic differentiation using gradient-based techniques. In the context of quantum physics and QPT, the elements that are typically optimized are often limited to a specific class, such as the underlying group or symmetry structure. Therefore, any optimization that is considered valid must comply with the corresponding constraints. This pertains to Riemannian optimization, which alters gradient-based methods to guarantee that the intermediate solutions maintain the desired property, such as unitarity, completely positive trace preserving (CPTP), etc.

In particular, we model this process geometrically, as visualized in Fig. 1b. All possible quantum channels exist on a smooth continuous surface known as a manifold \mathcal{M} . Gradient-based methods find the point that minimizes the loss function by taking steps along the manifold. In our case, the loss function compares simulated data generated by applying the current channel \mathbb{k} with the real data that was generated by the true target channel \mathcal{E} . Furthermore, the unique properties of quantum channels restrict which channels are physically possible, changing the shape of our manifold. Specifically, quantum channels exist on the Stiefel manifold.

This paper makes the following major contributions:

- Formulates Riemannian optimization via gradient descent on the Stiefel manifold for QPT.
- Adapts stochastic Adam optimizer for QPT.
- Provides an open-source implementation Qutee.jl [34], with native accelerated computation via GPUs.
- Demonstrates the effectiveness of QPT with respect to process dimension and characterization on a set of real and simulated experiments.

The remainder of this paper is organized as follows. Section II provides formulation of quantum channels. Section III describes our proposed QPT framework. Section IV presents two case studies using simulation of quantum processes. Section V presents a case study with hardware-based evaluation of our QPT framework. Finally, Section VI concludes the paper.

II. QUANTUM CHANNELS

A quantum channel is generally represented as a completely positive (CP) and trace-preserving (TP) linear map \mathcal{E} that maps a quantum state represented by the density matrix ρ to another state represented by ρ' , i.e., $\rho' = \mathcal{E}(\rho)$. The adjoint of a quantum channel \mathcal{E}^\dagger generalizes the Heisenberg picture (under the Frobenius norm) and maps operators E to other




Fig. 2: Pauli-Liouville representation of simulated “noise.” A channel with Pauli noise $K_{\text{noise}} = [\sqrt{p}I \otimes I, \sqrt{(1-p)}\sigma_z \otimes \sigma_z]$ is applied after a unitary operation U . Process tomography is performed to fit the noise+ U channel from measurement data. The “noise” part defined as $\hat{P} \circ U^{-1} - I$. (a) $p = 0.1$, (b) $p = 0.25$, and (c) $p = 0.25$ with measurement data perturbed by a normal distribution with $\epsilon = 0.01$.

operators E' , i.e., $E' = \mathcal{E}^\dagger(E)$. This section details how \mathcal{E} can be represented in various forms to highlight specific attributes.

A. Kraus Representation

For computations, it is convenient to use the Kraus-operator representation (also known as the canonical form) of a quantum channel \mathcal{E} and its adjoint \mathcal{E}^\dagger . Formally, it is

$$\mathcal{E}(\rho) = \sum_{l=1}^k K_l \rho K_l^\dagger \quad \text{and} \quad \mathcal{E}^\dagger(E) = \sum_{l=1}^k K_l^\dagger E K_l \quad (1)$$

where \mathcal{E} acts on states ρ and \mathcal{E}^\dagger acts on operators E . The k complex-valued matrices $\{K_l\}$ of dimension $N \times N$ are called the Kraus operators, with the requirement that the channel is trace-preserving (for I identity) $\mathcal{E}^\dagger(I) = \sum_l K_l^\dagger K_l = I$. Note that the Kraus representation is not unique, as $F_i = \sum_j U_{i,j} K_j$ represents the same CPTP map defined by a unitary transformation U .

Additionally, it is more efficient to work with the minimum number of Kraus operators needed to represent a CPTP map, which is known as the *Kraus rank*. Given that a Hilbert space of dimension N contains N^2 linearly independent operators, the Kraus rank is no larger than N^2 [24]. Thus, we may convert the Kraus representation to an $N^2 \times N^2$ (Hermitian) Choi matrix and obtain the minimal representation by considering only eigenvectors with non-zero eigenvalues [24]. An alternative method is to compute the overlap matrix $C_{ij} = \text{Tr}(K_i K_j^\dagger)$ and diagonalize it $C = V^\dagger D V$. The new Kraus operators $\tilde{K}_i = \sum_j V_{ij} K_j$ will be the minimal representation, where some matrices will be zeros if the original representation was redundant. Finally, we define a matrix of dimension $kN \times N$

$$\mathbb{k} = [K_1 \ K_2 \ \dots \ K_k]^T \quad (2)$$

by column-stacking the k Kraus operators. This matrix \mathbb{k} is an element of the Stiefel manifold, satisfying $\mathbb{k}^\dagger \mathbb{k} = I$ (the TP condition).

B. Liouville and Choi Representation

Although we perform all computations using the Kraus representation, it is illustrative to discuss the properties of CPTP maps via their Liouville representation. The Liouville (superoperator) representation of \mathcal{E} is an $N^2 \times N^2$ matrix of the form $\hat{\mathcal{E}} = \sum_l K_l \otimes \bar{K}_l$ (\bar{K}_l being the complex conjugate of K_l) acting on a vectorized N^2 -dimensional density matrix $|\rho\rangle\rangle$ from the left. This representation, equivalent to the Kraus representation, conveniently allows for the examination of a channel's spectral properties along with performing channel composition, i.e., for channels \mathcal{E} and \mathcal{L} : $(\mathcal{E} \circ \mathcal{L})(\rho) = \mathcal{E}(\mathcal{L}(\rho)) = \hat{\mathcal{E}} \cdot \hat{\mathcal{L}} \cdot |\rho\rangle\rangle$ where \cdot is the standard matrix (vector) product. As a result, a specific component of the quantum process, such as noise, can be isolated.

Since the Liouville representation consists of complex elements, it can be convenient to represent $\hat{\mathcal{E}}$ (when $N = 4^n$) in the Pauli basis, where all the matrix elements are real. The Pauli-Liouville may be written as $\hat{\mathcal{P}}_{i,j} = \text{Tr} \left[E_j^\dagger \mathcal{E}(E_i) \right]$ where operators $\{E_i\}$ are the Pauli operators. For example, Fig. 2 displays a heatmap of the Pauli-Liouville matrix where a unitary channel is extracted to reveal only a “noise” part.

Furthermore, the Liouville matrix is related to the Choi matrix \hat{J} through an involution $\hat{J}_{ij,kl} = \hat{\mathcal{E}}_{ik,jl}$. By the Choi-Jamiołkowski isomorphism, a quantum channel may be viewed as a quantum state. Therefore, one way to define distances between channels is via state fidelity

$$F(\rho, \sigma) = \frac{1}{N^2} \left(\text{Tr} \sqrt{\sqrt{\rho} \sigma \sqrt{\rho}} \right)^2 \quad (3)$$

where ρ and σ are two quantum channels as Choi matrices.

III. QUANTUM PROCESS TOMOGRAPHY USING RIEMMANIAN OPTIMIZATION

This section describes our proposed QPT framework via Riemmanian optimization. Assuming the convention in (2) of column-stacked Kraus operators that reside in the Stiefel manifold \mathcal{M} , and a loss function $\ell(\cdot)$, we formulate the task of QPT as the following optimization problem

$$\min_{\mathbb{K} \in \mathcal{M}} \ell(\mathbb{K}). \quad (4)$$

A general loss function for the QPT task may be defined as

$$\ell(\mathbb{K}) = \sum_{i,j} \left[d_{i,j} - \text{Tr} \left(M_j \sum_c K_c \rho_i K_c^\dagger \right) \right]^2, \quad (5)$$

where M_j denotes the j -th measurement operator, ρ_i denotes the i -th initial state, and $d_{i,j}$ denotes readout of the experiment. In other words, the optimization problem is to find a set of Kraus operators that minimizes the difference between experimental and simulated data. A graphical representation of this process is presented in Fig. 1a.

This section is organized as follows. First, we outline the steps for using the output of automatic differentiation to obtain a gradient residing in the tangent bundle. Next, we discuss an efficient method for performing the gradient descent via retraction. Finally, we discuss optimization opportunities.

A. Stiefel Manifold, Tangents, and Automatic Differentiation

We recall that the Stiefel manifold consists of all $n \times p$, $n \geq p$ unitary matrices

$$\mathcal{M}(n, p) = \{\mathbb{K} \in \mathbb{C}^{n \times p} : \mathbb{K}^\dagger \mathbb{K} = I_p\}. \quad (6)$$

The tangent space at a point $\mathbb{K} \in \mathcal{M}(n, p)$ is given by

$$T_{\mathbb{K}}\mathcal{M}(n, p) = \{\mathbb{K}\Omega + \mathbb{K}_\perp \Theta : \Omega = -\Omega^\dagger, \Theta \in \mathbb{C}^{(n-p) \times p}\} \quad (7)$$

where $\mathbb{K}_\perp \in \mathbb{C}^{n \times (n-p)}$ is orthonormal and $\text{span}(\mathbb{K}_\perp) = (\text{span}(\mathbb{K}))^\perp$. The Stiefel manifold becomes a Riemannian manifold by introducing a Riemannian metric g , which is a smoothly varying inner product in the tangent space. There are two natural inner products for the tangent space: the Euclidean inner product inherited from the embedding space $\mathbb{C}^{n \times p}$ and the canonical inner product. That is, given $A, B \in T_{\mathbb{K}}\mathcal{M}(n, p)$, we may opt for either:

$$\langle A, B \rangle = \left\{ \text{Tr}(A^\dagger B) \text{ or } \text{Tr} \left(A^\dagger \left(I - \frac{1}{2} X X^\dagger \right) B \right) \right\}. \quad (8)$$

The distinction arises in the norm of a vector $A = \mathbb{K}\Omega + \mathbb{K}_\perp \Theta$,

$$\|A\| = \sqrt{\langle A, A \rangle} = \begin{cases} \sqrt{\|\Omega\|^2 + \|\Theta\|^2} & (\text{inherited}) \\ \sqrt{\frac{1}{2}\|\Omega\|^2 + \|\Theta\|^2} & (\text{canonical}) \end{cases} \quad (9)$$

The factor of $\frac{1}{2}$ gives an equal weighting of both Ω and Θ ($\Omega = -\Omega^\dagger$ produces an extra factor of 2.) We assume the canonical inner product throughout the remainder of this paper.

We use reverse (adjoint) mode automatic differentiation to compute the loss gradient $\ell : \mathcal{M}(n, p) \rightarrow \mathbb{R}$, which we will denote by $\nabla \ell$. However, automatic differentiation assumes that the underlying space is Euclidean. Thus, the calculated gradient $\nabla \ell \in \mathbb{C}^{n \times p}$ does not adhere to the properties of tangent space of the Stiefel manifold in (7). To obtain a gradient that resides in the tangent space, we use the fact that the Stiefel manifold may be embedded in the Euclidean space. Therefore, the gradient vector in the Stiefel tangent space can be expressed as

$$\nabla_{T_{\mathbb{K}}\mathcal{M}} \ell = \text{proj}_{\mathbb{K}}(\nabla \ell) \quad (10)$$

where $\text{proj}_{\mathbb{K}}$ projects the Euclidean gradient onto the Stiefel tangent space and is defined as

$$\text{proj}_{\mathbb{K}}(X) = X - \mathbb{K}(\mathbb{K}X^\dagger + \mathbb{K}^\dagger X)/2. \quad (11)$$

B. Gradient Descent and Retraction

Gradient-based methods typically assume an initial guess x for a parameter and then iteratively evaluate the loss function, compute the gradient, and then update the parameter

$$x \leftarrow x - \beta \nabla \ell(x). \quad (12)$$

Similarly, the gradient-based approach to solve the optimization problem in (4) assumes an initial guess for the quantum channel $\mathbb{K} \in \mathcal{M}$ and iteratively updates the guess. However, unlike the simple update rule of (12), extra care must be taken to ensure that each individual update to \mathbb{K} remains on the manifold.

Let $\gamma(\tau)$ denote a curve on the manifold such that $\gamma(0) = \mathbb{k}$ and $\dot{\gamma}(0) = \nabla_{T_{\mathbb{k}}\mathcal{M}}\ell(\mathbb{k})$. The curve $\gamma(\tau)$ is considered a geodesic if it satisfies the exponential map

$$\gamma(\tau) = \text{Exp}_{\mathbb{k}}(\tau \nabla_{T_{\mathbb{k}}\mathcal{M}}\ell(\mathbb{k})), \text{ for } \tau \in [0, 1]. \quad (13)$$

In other words, the exponential map takes a vector in tangent space to a point on the manifold such that it remains on the shortest path. Recall that the objective is to minimize a function while remaining constrained to the manifold. One way to view this is to find the geodesic on the manifold that moves us from the starting point to the optimal point that minimizes the function. However, this is difficult to formulate generally, and instead we must rely on approximate algorithms such as line search. In this case, via the gradient-based method, we perform a step in the tangent space and are still constrained to the manifold via the exponential mapping.

The exponential map is not computationally efficient. We can approximate the exponential map to m -th order in a Taylor or Padé expansion. Other approximations include matrix decomposition, such as QR and singular value decomposition (SVD). Such approximation of the exponential is known as retraction. A common choice for a retraction is the Cayley transform (a first-order Padé expansion), which defines a curve

$$Y(\tau) = \left(I + \frac{\tau}{2}W\right)^{-1} \left(I - \frac{\tau}{2}W\right) \mathbb{k} \quad (14)$$

where W is a skew-hermitian matrix, i.e. $W^\dagger = -W$. By choosing $W = \bar{W} - \bar{W}^\dagger$, where $\bar{W} = \nabla \ell(\mathbb{k}) \mathbb{k}^\dagger - \frac{1}{2} \mathbb{k} (\mathbb{k}^\dagger \nabla \ell(\mathbb{k}) \mathbb{k}^\dagger)$, the transformation implicitly projects the gradient onto the tangent space in the descent direction. Here, W is the matrix operator for projection in (11).

For large models \mathbb{k} , the Cayley transform in (14) cannot be performed efficiently due to the matrix inversion. It is possible to simplify the computation of the inverse via the Sherman-Morrison-Woodbury formula, however, this is only efficient to do so for matrices where $n \ll p$. Instead, we adapt the fixed-point iteration method for the Cayley transform [35], which requires only matrix multiplications while also serving as an efficient approximation. The iterative update to the Cayley transform is defined as

$$Y(\tau)^i \leftarrow \mathbb{k} + \frac{\tau}{2} W(\mathbb{k} + Y(\tau)^{i-1}) \quad (15)$$

where i denotes the iterative step. Figure 3 highlights the computational and error improvement in iterative Cayley.

C. Momentum Updates and Vector Transport

Data collected in the real world are subject to randomness. This means that the measurements we obtain are assumed to come from an underlying probability distribution. Therefore, our objective is to minimize the expected value $\mathbb{E}[\ell(\mathbb{k})]$, where $\ell(\cdot)$ depends on stochastic measurement data $d_{i,j}$. We may consider the gradient $\nabla \ell$ as a random variable for which obtaining the true first (mean) and second (standard deviation) moments is costly. Adam, a successful technique for gradient-based optimization of stochastic objective functions, estimates and updates the moments via random subsamples




Fig. 3: Iterative-based Cayley transform has better performance (time and accuracy) with respect to traditional Cayley transform. The accuracy (dashed lines) are computed by solving the exact exponential map and taking the norm difference with respect to approximated retraction.

(minibatches) of data points. Adam further accumulates the gradients from previous iterations into a momentum term, which reduces the oscillations in $\ell(\cdot)$ arising from randomness and hence accelerates convergence. However, Adam is only suitable for optimizations in a Euclidean space. We adopt Adam for the task of QPT on the Stiefel manifold.

The key idea is to move a vector in the tangent space along the search paths on the manifold. In general, this is done using a parallel transport that preserves the geometrical properties of the vector from one tangent space to another. However, this requires the solution of a costly differential equation. Instead, some geometric conditions may be relaxed to perform an efficient *vector transport*. Since the Stiefel manifold is a submanifold of the Euclidean space, the tangent space is a subspace of the Euclidean space. Hence, a vector transport on the Stiefel manifold is a projection onto the tangent space. Following [35], using the property that the projection is a linear map, the key Adam steps may first be performed in Euclidean space, followed by a projection:

$$\alpha \text{proj}_{\mathbb{k}_k}(M_k) + \beta \text{proj}_{\mathbb{k}_k}(\nabla \ell(\mathbb{k}_k)) = \text{proj}_{\mathbb{k}_k}(\alpha M_k + \beta \nabla \ell(\mathbb{k}_k))$$

Namely, a linear combination of the gradient $\nabla \ell(\mathbb{k})$ and the momentum M_k for step $k-1$ is taken with real “learning” coefficients α and β . Instead of explicitly performing a projection, the iterative Cayley transform in (15) implicitly projects the vector onto the tangent space while also performing a parameter update.

IV. CASE STUDY: SIMULATION OF QUANTUM PROCESSES

In this section, we present the results of our QPT optimization on a set of problems with varying difficulty: randomized channels and controlling coherent states of a harmonic oscillator. In both examples, we use a random initial guess \mathbb{k} for the channel. Furthermore, all optimizations are performed on an Nvidia RTX6000 GPU with 48GB of GDDR6 of memory.

A. Random Channels with Insufficient Measurements

In our first case study, we explore the capabilities of Riemannian optimization of quantum processes with respect to dimensions, Kraus rank, and measurement error. We follow a direct QPT approach, where 6^n probes ρ_i and 6^n measurements M_j are selected as eigenstates of the generalized




Fig. 4: (a) Mean loss of the optimization as noise ϵ is increased for different channel ranks. Reconstruction error saturates as a power curve with respect to noise ϵ (note the log-log scale.) (b) Channel fidelity at $\epsilon = 0.01$ for a reconstructed quantum channel with varying Kraus rank as well as access to only a fraction of data. Kraus rank is the dominant factor for fidelity, while the fraction of data is responsible for fidelity margin.

TABLE I: Comparing execution time of QPT methods.

# Qubits	Hilbert space	Basis size	Qiskit time	Our time
2	4	16	4.153s	3.424s
3	8	64	76.134s	4.791s
4	16	256	1355.841s	61.766s

Gellmann matrices. For a random channel $\mathcal{E}_{\text{rand}}$, the simulated data is computed as

$$d_{i,j} = \text{Tr}(M_j \mathcal{E}_{\text{rand}}(\rho_i)) + \mathcal{N}(0, \epsilon) \quad (16)$$

where $\mathcal{N}(0, \epsilon)$ introduces measurement error following a normal distribution with variance ϵ . We generate a random channel \mathbb{k} by computing the orthonormal part of the QR decomposition of a random $rn \times n$ matrix where each element is sampled from the normal distribution.

Figure 4a shows simulation results for a 3-qubit system undergoing a random process $\mathcal{E}_{\text{rand}}$. The Hilbert space is $2^3 = 8$, with a maximal Kraus rank of $4^3 = 64$. Various Kraus ranks are fitted, with different measurement noise. The computed loss with lower Kraus ranks saturate sooner, as the fitted channel is simply not expressive enough to represent the full-rank process. Figure 4b further exemplifies the importance of Kraus rank. The fidelity of the reconstructed channel (3) with respect to original channel is computed, where the reconstructed channel was computed with a fraction of data, namely $\sqrt{\nu}6^n$ probes and $\sqrt{\nu}6^n$ measurements are used given a ratio $\nu \in (0, 1]$. High fidelity is observed even with the absence of a fraction of data. We are able to reconstruct the quantum channel in a few seconds, even at maximal rank. In comparison, other tomography techniques become limiting even at 3-qubits. Table I compares our method to state-of-the-art libraries, such as Qiskit [36] (via constrained maximum-likelihood estimate), for several qubit sizes.

B. Universal Control of a Harmonic Oscillator

Our second case study considers a continuous-variable quantum system, where information is encoded in an infinite-dimensional Hilbert space, following a similar example in [30]. Due to the nature of an infinite-dimensional system, selecting




Fig. 5: Wigner functions of a probe state undergoing ideal simulation and tomographic reconstruction with a Fock space $N = 64$. The data are measured values of displaced parity operator $\Pi_d(\beta)$, but with fewer sample points β to mimic an experimental setup.

an appropriate Hilbert space cut-off is crucial for correctly describing a state numerically. In our case, we consider a Fock space cut-off of $N = 64$, which exceeds previously considered dimension cut-offs.

Such systems (single bosonic mode) may be modeled as quantum harmonic oscillators with associated creation \hat{a}^\dagger and annihilation operators \hat{a} . Coherent states $|\alpha\rangle$ are the unique eigenstates with corresponding eigenvalue α of the annihilation operator $\hat{a}|\alpha\rangle = \alpha|\alpha\rangle$. Universal control of a harmonic oscillator can be achieved by combining displacement operation

$$D(\alpha) = \exp(\alpha\hat{a}^\dagger - \alpha^*\hat{a}), \quad |\alpha\rangle = D(\alpha)|0\rangle \quad (17)$$

and a selective number-dependent arbitrary phase (SNAP) gate

$$S(\vec{\phi}) = \sum_n e^{i\phi_n} |n\rangle \langle n| \quad (18)$$

which adds a real phase ϕ_n to the Fock state $|n\rangle$. This can be shown by observing that the generators and commutators of $D(x)$ and $S(\phi)$ generate the full Lie algebra $u(N)$ for any truncated oscillator space [37].

We simulate a setup where the coherent probe states $\rho_i = |\alpha\rangle \langle \alpha|$ undergo a displacement + SNAP process, and are then measured via a displaced-parity operator [38]

$$\Pi_d(\beta) = \sum_{n=0}^N (-1)^n D(\beta) |n\rangle \langle n| D(\beta)^\dagger. \quad (19)$$

Similarly, the Wigner functions are approximated using expectation values of the parity operator [39]. Figure 5 shows the Wigner functions for one instance of the complete process. In other words, an initial probe state ρ_i undergoes the simulated process to a final state ρ'_i . The computation of the expectation values for the displaced-parity operator is performed on a discretized grid. This was done to mimic an experimental scenario in which conducting a full Wigner tomography is infeasible. We find good results with a rank-fitted channel $r = 2$, despite the fact that the process is purely unitary ($r=1$). This indicates, non-unitary behavior, which arises from the discretization of measurement expectation values.




Fig. 6: Circuit of quantum nondemolishing measurement.

V. STEADY STATES OF NONDEMOLISHMENT MEASUREMENT IN QUANTUM HARDWARE

In the following experiment, we assume a setup as shown in Fig. 6. Namely, a system that is coupled to a detector

$$\rho = |0\rangle_D \langle 0|_D \otimes \rho_S \quad (20)$$

where the system is in an unknown mixed state ρ_S and the detector is in a known pure state $|0\rangle_D$. The system evolves for a duration dt subject to a two-body Hamiltonian H ,

$$\rho' = U \rho U^\dagger, \quad U = \exp(-iHdt). \quad (21)$$

The detector is then measured and promptly reset to its known state $|0\rangle$. As a result of this procedure, the system is disrupted, resulting in a new state given by a partial trace

$$\rho_S = \text{Tr}_D [U(|0\rangle_D \langle 0|_D \otimes \rho_S) U^\dagger]. \quad (22)$$


In this ideal scenario, we may express the evolution of ρ_S as a quantum channel \mathcal{E} in terms of two Kraus operators

$$K_1 = \langle 0|_D U |0\rangle_S \quad \text{and} \quad K_2 = \langle 0|_D U |1\rangle_S. \quad (23)$$

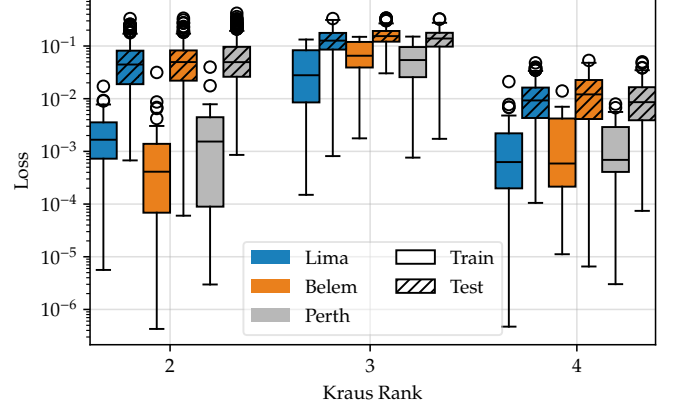
The evolution then follows (1). The repeated application of \mathcal{E} , will cause the system to converge to a steady state.

We run the circuit in Fig. 6 on three different IBM Quantum computers: *ibmq_lima*, *ibmq_belem*, *ibmq_perth*. The system and detector are selected as neighboring qubits coupled via a cross-resonance (CR) gate. We idle the qubits for a duration dt , corresponding to an idealized evolution $U_{[dt]}$. After the detector measurement, the system is measured on a different basis by performing a single qubit rotation gate followed by measurement in the computational basis. This process is repeated in which the measurement of the system is only performed in the end, providing an estimate of the quantum state ρ_S after each non-demolishing measurement. In addition, we performed 30 trials.

The process is considered to be strictly Markovian, and that $U_{[dt]}$ is stable from one iteration to the next. Assuming the loss function of (5), we fit a channel \mathcal{E} with ranks $r = \{2, 3, 4\}$ in 20 trials, leaving 10 trials to test the fitted channel. As shown in Fig. 7b, the maximum rank ($r = 4$) provides the best results from optimization against the test data. While (23) states that the ideal scenario is described with a Kraus rank $r = 2$, in reality the detector state is not perfectly initialized to $|0\rangle$ and may instead be in a state $|1\rangle$ or a general superposition. Hence, the maximum rank of $r = 4$ is required to capture faulty initialization. In addition, stochastic and non-Markovian effects are present on quantum devices which can no longer be described by a quantum channel.



(a) History of the loss $\ell(\cdot)$ throughout optimization with varying Kraus rank.



(b) Box-and-whisker plots of the final loss with varying Kraus ranks for training and testing sets of data.

Fig. 7: Results of optimization on IBM Quantum computers.

VI. CONCLUSION AND FUTURE WORK

We have introduced a data-driven approach for the constrained optimization task of quantum process tomography. Utilizing advanced stochastic objective optimizers, we can achieve outstanding results by considering the quantum process as residing in a Stiefel manifold and performing local updates that respect the geometry. We implement an iterative Cayley transform for an efficient retraction on the manifold, as well as adopt the implicit projection to include moment estimates of stochastic optimization. These considerations alleviate some numerical and memory issues, which distinguishes this approach from other tomographic methods. We demonstrated that our approach can quickly reconstruct quantum processes, such as scenarios with a Hilbert space of dimension 64. Furthermore, using random processes, we established the relationship between Kraus rank and measurement error for precision in estimation.

There are numerous promising avenues for future research. One interesting avenue is applying methods of Riemmanian optimization to other approximations and descriptions of quantum processes, such as tensor networks. Namely, such avenues may be utilized to characterize and deduce non-Markovian noise processes. Another important avenue is the consideration of error bars and uncertainty. In part, such considerations are important for distinguishing errors and confidence regions of the tomographic process itself.

ACKNOWLEDGMENTS

We acknowledge the use of IBM Quantum services for this work. The views expressed are those of the authors, and do not reflect the official policy or position of IBM or the IBM Quantum team. The computing systems used in this work is supported by DOE (DE-SC-0009883). This research is partially supported by the grants from DARPA (HR0011-24-3-0004) and NSF (CCF-1908131).

REFERENCES

- [1] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, A. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, no. 7779, pp. 505–510, Oct. 2019.
- [2] L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. F. Bulmer, F. M. Miatto, L. Neuhaus, L. G. Helt, M. J. Collins, A. E. Lita, T. Gerrits, S. W. Nam, V. D. Vaidya, M. Menotti, I. Dhand, Z. Vernon, N. Quesada, and J. Lavoie, “Quantum computational advantage with a programmable photonic processor,” *Nature*, vol. 606, no. 7912, pp. 75–81, Jun. 2022.
- [3] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, “Surface codes: Towards practical large-scale quantum computation,” *Phys. Rev. A*, vol. 86, no. 3, p. 032324, Sep. 2012.
- [4] J. Guillaud and M. Mirrahimi, “Repetition Cat Qubits for Fault-Tolerant Quantum Computation,” *Phys. Rev. X*, vol. 9, no. 4, p. 041053, Dec. 2019.
- [5] I. L. Chuang and M. A. Nielsen, “Prescription for experimental determination of the dynamics of a quantum black box,” *Journal of Modern Optics*, vol. 44, no. 11-12, pp. 2455–2467, Nov. 1997.
- [6] J. Z. Blumoff, K. Chou, C. Shen, M. Reagor, C. Axline, R. T. Brierley, M. P. Silveri, C. Wang, B. Vlastakis, S. E. Nigg, L. Frunzio, M. H. Devoret, L. Jiang, S. M. Girvin, and R. J. Schoelkopf, “Implementing and Characterizing Precise Multiqubit Measurements,” *Phys. Rev. X*, vol. 6, no. 3, p. 031041, Sep. 2016.
- [7] Y. Chen, M. Farahzad, S. Yoo, and T.-C. Wei, “Detector tomography on IBM quantum computers and mitigation of an imperfect measurement,” *Phys. Rev. A*, vol. 100, no. 5, p. 052315, Nov. 2019.
- [8] S. G. Schirmer and D. K. L. Oi, “Two-qubit Hamiltonian tomography by Bayesian analysis of noisy data,” *Phys. Rev. A*, vol. 80, no. 2, p. 022333, Aug. 2009.
- [9] J. H. Cole, “Hamiltonian tomography: The quantum (system) measurement problem,” *New J. Phys.*, vol. 17, no. 10, p. 101001, Sep. 2015.
- [10] L. Cincio, K. Rudinger, M. Sarovar, and P. J. Coles, “Machine Learning of Noise-Resilient Quantum Circuits,” *PRX Quantum*, vol. 2, no. 1, p. 010324, Feb. 2021.
- [11] E. Nielsen, J. K. Gamble, K. Rudinger, T. Scholten, K. Young, and R. Blume-Kohout, “Gate Set Tomography,” *Quantum*, vol. 5, p. 557, Oct. 2021.
- [12] J. Fiurášek and Z. Hradil, “Maximum-likelihood estimation of quantum processes,” *Phys. Rev. A*, vol. 63, no. 2, p. 020101, Jan. 2001.
- [13] M. F. Sacchi, “Maximum-likelihood reconstruction of completely positive maps,” *Phys. Rev. A*, vol. 63, no. 5, p. 054104, Apr. 2001.
- [14] K. Schultz, “Exponential families for Bayesian quantum process tomography,” *Phys. Rev. A*, vol. 100, no. 6, p. 062316, Dec. 2019.
- [15] J. M. Lukens, K. J. H. Law, A. Jasra, and P. Lougovski, “A practical and efficient approach for Bayesian quantum state estimation,” *New J. Phys.*, vol. 22, no. 6, p. 063038, Jun. 2020.
- [16] G. C. Knee, E. Bolduc, J. Leach, and E. M. Gauger, “Quantum process tomography via completely positive and trace-preserving projection,” *Phys. Rev. A*, vol. 98, no. 6, p. 062336, Dec. 2018.
- [17] M. Guță, J. Kahn, R. Kueng, and J. A. Tropp, “Fast state tomography with optimal error bounds,” *J. Phys. A: Math. Theor.*, vol. 53, no. 20, p. 204001, Apr. 2020.
- [18] T. Surawy-Stepney, J. Kahn, R. Kueng, and M. Guta, “Projected Least-Squares Quantum Process Tomography,” *Quantum*, vol. 6, p. 844, Oct. 2022.
- [19] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, “Quantum state tomography via compressed sensing,” *Phys. Rev. Lett.*, vol. 105, no. 15, p. 150401, Oct. 2010.
- [20] C. H. Baldwin, A. Kalev, and I. H. Deutsch, “Quantum process tomography of unitary and near-unitary maps,” *Phys. Rev. A*, vol. 90, no. 1, p. 012110, Jul. 2014.
- [21] 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*, vol. 90, no. 14, p. 144504, Oct. 2014.
- [22] G. Torlai, C. J. Wood, A. Acharya, G. Carleo, J. Carrasquilla, and L. Aolita, “Quantum process tomography with unsupervised learning and tensor networks,” *Nat Commun*, vol. 14, no. 1, p. 2858, May 2023.
- [23] H. Kadri, S. Ayache, R. Huusari, A. Rakotomamonjy, and L. Ralaivola, “Partial Trace Regression and Low-Rank Kraus Decomposition,” Aug. 2020.
- [24] M.-D. Choi, “Completely positive linear maps on complex matrices,” *Linear Algebra and its Applications*, vol. 10, no. 3, pp. 285–290, Jun. 1975.
- [25] M. AbuGhanem, “Full quantum process tomography of a universal entangling gate on an IBM’s Quantum Computer,” *arXiv.org*, Feb. 2024.
- [26] R. Levy, D. Luo, and B. K. Clark, “Classical shadows for Quantum Process Tomography on near-term quantum computers,” *Physical Review Research*, Jan. 2024.
- [27] G. Carleo and M. Troyer, “Solving the quantum many-body problem with artificial neural networks,” *Science*, vol. 355, no. 6325, pp. 602–606, Feb. 2017.
- [28] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, “Reconstructing quantum states with generative models,” *Nat Mach Intell*, vol. 1, no. 3, pp. 155–161, Mar. 2019.
- [29] O. Sharir, Y. Levine, N. Wies, G. Carleo, and A. Shashua, “Deep Autoregressive Models for the Efficient Variational Simulation of Many-Body Quantum Systems,” *Phys. Rev. Lett.*, vol. 124, no. 2, p. 020503, Jan. 2020.
- [30] S. Ahmed, F. Quijandria, and A. F. Kockum, “Gradient-Descent Quantum Process Tomography by Learning Kraus Operators,” *Phys. Rev. Lett.*, vol. 130, no. 15, p. 150402, Apr. 2023.
- [31] H.-J. Liao, J.-G. Liu, L. Wang, and T. Xiang, “Differentiable Programming Tensor Networks,” *Phys. Rev. X*, vol. 9, no. 3, p. 031041, Sep. 2019.
- [32] F. Pan, P. Zhou, S. Li, and P. Zhang, “Contracting Arbitrary Tensor Networks: General Approximate Algorithm and Applications in Graphical Models and Quantum Circuit Simulations,” *Phys. Rev. Lett.*, vol. 125, no. 6, p. 060503, Aug. 2020.
- [33] G. Torlai, J. Carrasquilla, M. T. Fishman, R. G. Melko, and M. P. A. Fisher, “Wave-function positization via automatic differentiation,” *Phys. Rev. Res.*, vol. 2, no. 3, p. 032060, Sep. 2020.
- [34] D. Volya, “RustyBamboo/QuTe.jl,” Aug. 2023.
- [35] J. Li, L. Fuxin, and S. Todorovic, “Efficient Riemannian Optimization on the Stiefel Manifold via the Cayley Transform,” Feb. 2020.
- [36] Qiskit contributors, “Qiskit: An open-source framework for quantum computing,” 2023.
- [37] S. Krastanov, V. V. Albert, C. Shen, C.-L. Zou, R. W. Heeres, B. Vlastakis, R. J. Schoelkopf, and L. Jiang, “Universal control of an oscillator with dispersive coupling to a qubit,” *Physical Review A*, vol. 92, no. 4, Oct. 2015.
- [38] R. F. Bishop and A. Vourdas, “Displaced and squeezed parity operator: Its role in classical mappings of quantum theories,” *Physical Review A*, vol. 50, no. 6, pp. 4488–4501, Dec. 1994.
- [39] A. Royer, “Wigner function as the expectation value of a parity operator,” *Physical Review A*, vol. 15, no. 2, pp. 449–450, Feb. 1977.