

Agnostic Process Tomography

Chirag Wadhwa^{1,*}, Laura Lewis¹, Elham Kashefi^{1,2} and Mina Doosti¹

¹*School of Informatics, University of Edinburgh, Edinburgh, United Kingdom*

²*Laboratoire d'Informatique de Paris 6, CNRS, Sorbonne Université, Paris, France*



(Received 27 May 2025; accepted 27 October 2025; published 24 December 2025)

Characterizing a quantum system by learning its state or evolution is a fundamental problem in quantum physics and learning theory with a myriad of applications. Recently, as a new approach to this problem, the task of agnostic state tomography was defined, in which one aims to approximate an arbitrary quantum state by a simpler one in a given class. Generalizing this notion to quantum processes, we initiate the study of agnostic process tomography: given query access to an unknown quantum channel Φ and a known concept class $\mathcal{C}$ of channels, output a quantum channel that approximates Φ as well as any channel in the concept class $\mathcal{C}$, up to some error. In this work, we propose several natural applications for this new task in quantum machine learning, quantum metrology, classical simulation, and error mitigation. In addition, we give efficient agnostic process tomography algorithms for a wide variety of concept classes, including Pauli strings, Pauli channels, quantum junta channels, low-degree channels, and a class of channels produced by QAC⁰ circuits. The main technical tool we use is Pauli spectrum analysis of operators and superoperators. We also prove that, using ancilla qubits, any agnostic state tomography algorithm can be extended to one solving agnostic process tomography for a compatible concept class of unitaries, immediately giving us efficient agnostic learning algorithms for Clifford circuits, Clifford circuits with few T gates, and circuits consisting of a tensor product of single-qubit gates. Together, our results provide insight into the conditions and new algorithms necessary to extend the learnability of a concept class from the standard tomographic setting to the agnostic one.

DOI: [10.1103/q2nb-zg9m](https://doi.org/10.1103/q2nb-zg9m)

I. INTRODUCTION

A fundamental problem in quantum physics and quantum learning theory is to completely characterize a quantum system by learning its evolution, a task called *quantum process tomography* [1–7]. Quantum process tomography has numerous applications in benchmarking [1–4,8–11], error mitigation [12–14], quantum metrology [15,16], quantum control [17], and quantum cryptography [18,19]. Unfortunately, recovering a full description of an arbitrary quantum process requires exponential resources [6].

On the other hand, when the unknown process possesses some underlying structure, e.g., for Clifford circuits [20,21], Clifford circuits with few non-Clifford T gates [21], shallow circuits [22,23], junta unitaries and channels [24,25], matrix product operators [26], circuits with

a bounded number of gates [27], quantum boolean functions [28,29], low-degree unitaries and channels [30], and Pauli channels [31–37], or when one only wishes to predict some limited properties of this process [8,38–41], efficient algorithms can be developed. However, the algorithms for the former case only apply in the *realizable* setting: the unknown quantum process must *exactly* satisfy this known structure. In practice, this assumption may not hold (e.g., due to noisy access), rendering these algorithms ineffective.

In this work, we study how one may still be able to learn a process with a desirable underlying structure when given access to a more complex one. In particular, we aim to answer the following central question.

Given an unknown quantum process, can we learn its best approximation with a simple structure?

Ideally, once we have such an approximation, the simple representation can be used as an ansatz to obtain new information about the complex process. We outline potential applications in quantum machine learning, metrology, classical simulation, and error mitigation in Sec. II A. Notably, the analogous problem for quantum states [42–46] does not apply to many of these settings, motivating our generalization to quantum processes.

*Contact author: chirag.wadhwa@ed.ac.uk

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.

More precisely, the problem we aim to solve is as follows: given access to an unknown quantum process Φ and a known concept class $\mathcal{C}$ of channels, output the nearest approximation of Φ from the class $\mathcal{C}$. We refer to this task as *agnostic process tomography*, inspired by classical agnostic learning [47,48] and recent developments for agnostic state tomography [42–46]. We are primarily interested in cases when the concept class has a simple underlying structure and consists of efficiently implementable processes. One may also consider *improper* agnostic process tomography, in which the output of the algorithm is not in the concept class $\mathcal{C}$, but we emphasize that most applications require *proper* learning. We define the task formally in Sec. II.

In this work, we initiate the study of agnostic learning of quantum processes and present new efficient agnostic learning algorithms for a wide variety of concept classes, including Pauli strings, Pauli channels, quantum junta channels, low-degree channels, and a class of channels produced by QAC^0 circuits. We also learn the class of bounded-gate circuits, and the associated analysis reveals that when studying agnostic process tomography, one should focus on reducing the time complexity rather than the sample complexity, analogously to the state case [42]. In addition, we prove that any agnostic state tomography algorithm can be extended to perform improper agnostic process tomography of a compatible concept class when equipped with ancilla qubits. Applying this to the agnostic state tomography algorithms of Ref. [43], we immediately obtain efficient improper agnostic learning algorithms for the classes of Clifford circuits, Clifford circuits with few T gates, and circuits consisting of a tensor product of single-qubit gates. Thus, our results answer our central question in the affirmative for these classes and also partially answer an open question of Ref. [42], regarding whether one can suitably define a notion of agnostic learning for processes and prove guarantees for the class of Clifford circuits.

For the concept classes we consider, realizable learning algorithms have already been developed [20–22,24,25,27,28,30]. As a whole, our results provide new insight into when these realizable algorithms can be extended, with different technical proofs, to the agnostic setting. In particular, for simpler concept classes such as Pauli strings, Pauli channels, bounded-gate circuits and QAC^0 channels, our algorithms and analysis do not differ significantly from the realizable setting [22,27,28]. On the other hand, for juntas and low-degree quantum channels, the realizable algorithms [24,25,30] fail to generalize straightforwardly to the agnostic case, requiring us to develop new, robust algorithms. Specifically, for juntas, the realizable algorithms and proofs rely crucially on a quantity (namely, the influence) that only helps characterize the unknown unitary or channel if it is indeed a junta. Similarly, for low-degree unitaries, we show that a simple agnostic extension of a realizable algorithm [30] results in




FIG. 1. (Proper) Agnostic process tomography. (Left) Conventional tomography. In the realizable case, the unknown quantum process Φ is promised to fall in the class $\mathcal{C}$, and the goal is to find some $\Phi' \in \mathcal{C}$ that is close to Φ . (Right) Agnostic tomography. Here, the unknown process Φ can be any completely positive trace-preserving (CPTP) map. The goal is to find some $\Phi' \in \mathcal{C}$ that is not far from the closest channel Φ^* to Φ in $\mathcal{C}$.

a drastically increased sample and time complexity compared to both our new algorithm and the original, realizable algorithm. We elaborate further in Sec. V.

We define agnostic process tomography formally in Sec. II. Then, we describe general applications of agnostic process tomography in Sec. II A and our results in Sec. III. The main technical tool we use is Pauli spectrum analysis of operators and superoperators [22,24,28–30,40], and we detail the ideas behind our proofs in Sec. V.

II. PROBLEM DEFINITION

Formally, we define agnostic process tomography as follows, where we adopt the notation of Refs. [44,49]. We also depict the definition in Fig. 1.

Definition 1 (Agnostic process tomography). Let $1 > \epsilon$, $\delta > 0$ and $\alpha \geq 1$. Let $\mathcal{C}$ be a class of completely positive trace-preserving (CPTP) maps. A learning algorithm $\mathcal{A}$ is an $(\alpha, \epsilon, \delta)$ -agnostic learner with respect to the class $\mathcal{C}$ of CPTP maps if, given access to an arbitrary CPTP map Φ , $\mathcal{A}$ outputs a hypothesis Φ' such that

$$\text{err}(\Phi, \Phi') \leq \alpha \cdot \text{opt}(\Phi, \mathcal{C}) + \epsilon \quad (\text{II.1})$$

with probability at least $1 - \delta$, where err denotes an error function and

$$\text{opt}(\Phi, \mathcal{C}) = \min_{\mathcal{E} \in \mathcal{C}} \text{err}(\Phi, \mathcal{E}) \quad (\text{II.2})$$

is the optimal error of the unknown channel with respect to $\mathcal{C}$. *Agnostic process tomography* with respect to $\mathcal{C}$ is the task of $(\alpha, \epsilon, \delta)$ -agnostic learning with respect to $\mathcal{C}$. When the hypothesis Φ' is from the concept class $\mathcal{C}$, we say that $\mathcal{A}$ is *proper*. Otherwise, we say that it is *improper*. Moreover, if Φ is a unitary channel and $\mathcal{C}$ is a class of unitaries, we refer to the task as *agnostic unitary estimation*.

Here, we purposefully leave some aspects of the definition flexible to capture the most general setting. In

particular, we do not specify the access model for the unknown process Φ (e.g., query access, statistical query access [22,50–52], samples from the output distribution [53–55], etc.) nor the error function err (e.g., diamond distance, average-case distance, Frobenius distance, etc.). Moreover, for brevity, we may refer to an $(\alpha, \epsilon, \delta)$ -agnostic learner as an α -agnostic learner and similarly refer to $\text{opt}(\Phi, \mathcal{C})$ as opt when the other arguments are clear from context.

For our earlier motivation and the applications in Sec. II A, we focus on *proper* agnostic learning. We also remark that even if one succeeds at agnostic process tomography of a given concept class, the hypothesis will not be of practical use if the concept class does not model the unknown channel well, i.e., if opt itself is large. Namely, despite finding a hypothesis Φ' that is close to the optimal process Φ^* , the optimal process itself may not approximate Φ well. One may consider, e.g., running property testing algorithms [20,24,25,28,56,57] beforehand to check if the concept class is suitable.

Note that agnostic process tomography includes conventional process tomography [1–6,20–23,26–37] as a special case, when the unknown process is guaranteed to be in $\mathcal{C}$. However, agnostic process tomography also allows one to learn quantum processes in the presence of arbitrary noise. This greatly contrasts realizable process tomography algorithms, which are usually extremely noise-sensitive and require the unknown process to exactly satisfy some known conditions. While learning under noise is one motivation, agnostic process tomography is more broadly applicable, as discussed in the next section.

A. Example applications of our framework

As discussed previously, selecting an appropriate concept class is crucial to applying agnostic learning algorithms to practical scenarios. By carefully choosing a concept class, we highlight a wide range of problems in near-term quantum computation that can be modeled in the framework of *proper* agnostic process tomography. We hope that this offers fresh perspectives and tools for further advancements in these fields. For some of these applications, our proposed algorithms are already applicable, while others require new algorithms, motivating further exploration of agnostic process tomography beyond this work. We also emphasize that these applications are unique to the agnostic process tomography framework and often cannot be addressed using agnostic state tomography [42,43] alone.

1. Resource-constrained implementation of complex quantum processes

A key application of agnostic process tomography is approximately implementing complex quantum processes on devices with limited resources. Suppose a user with

limited local resources has cloud access to a large fault-tolerant quantum computer and a desired complex process to run. There are many reasons why the user may want to approximate this complex process by one that can be implemented on a weaker device. For instance, repeated access to the remote system could be prohibitively expensive. Instead, the user can obtain the best local approximation of the process by agnostic learning with respect to the class of circuits that can be implemented with the user's local resources. For example, if the user's device has fewer qubits, one could consider agnostic unitary estimation of quantum juntas (Appendix C 4), which learns the best few-qubit implementation of an arbitrary unitary. Other interesting classes may include Clifford circuits, Cliffords with a few T gates and circuits consisting of tensor products of single-qubit gates (Appendix C 2). We again emphasize that this requires carefully choosing the concept class, as the complex process may not be well-approximated by an arbitrary class. This application is depicted in Fig. 2(a).

2. Approximate classical simulation

We highlight a special case of the above application for approximately classically simulating complex quantum processes. For example, given an arbitrary process, finding the closest Clifford circuit allows for efficient classical simulation [58,59]. If the process is not too far from all Clifford circuits, which can be verified via property testing [56,57], this will offer accurate results. This holds more generally for any classically simulable class admitting an efficient agnostic process tomography algorithm. This application is depicted in Fig. 2(b).

3. A unifying view of quantum machine learning

Classically, the theoretical framework of agnostic learning captures the most general settings in machine learning (ML). Namely, learning the best parameters for an ML model is equivalent to agnostic learning the concept class of all functions obtained by varying parameters of this ML model. We argue that a similar equivalence holds for quantum machine learning (QML). In QML, functions are evaluated using parameterized quantum circuits [60], or more generally, quantum channels [such as linear combinations of unitaries (LCUs) [61–63]]. By choosing an appropriate access model and error, agnostic process tomography of the class of channels expressible by an ansatz allows one to find the best parameters for the QML model. In this sense, efficient agnostic process tomography of a model class implies efficient trainability. Thus, studying agnostic process tomography under various access models may lead to new insights into QML for both quantum and classical problems. For example, many QML algorithms learn from expectation values of




FIG. 2. Applications of agnostic process tomography. (a) Resource-constrained implementation. A client with limited computational resources has access to a quantum server, which can implement the client's desired complex quantum process $\mathcal{E}$. The client chooses a suitable concept class based on its available resources and provides it as input to an agnostic process tomography (APT) algorithm. The APT algorithm has access to $\mathcal{E}$ via the quantum server and can (potentially adaptively) query $\mathcal{E}$ on any input state. The APT algorithm then outputs the best approximate implementation $\hat{\mathcal{E}}$ of $\mathcal{E}$, which the client can implement on their local device. (b) Classical simulation. Consider a user with access to a quantum device, which can implement some quantum process $\mathcal{E}$, and the user wishes to classically simulate $\mathcal{E}$. Given the class of Clifford circuits (or any classically simulable class) and after querying $\mathcal{E}$ from the quantum device, the APT algorithm outputs a Clifford circuit $\hat{\mathcal{C}}$ that approximates $\mathcal{E}$, which can then be classically simulated. (c) Quantum machine learning and metrology. Consider a class $\mathcal{C}_\theta = \{\mathcal{E}(\theta_i)\}_i$ of parameterized quantum processes, which naturally occur in quantum machine learning (QML), e.g., parameterized quantum circuits (PQCs) and linear combination of unitaries (LCU), and quantum metrology. Running APT with respect to $\mathcal{C}_\theta$ and an appropriate access model (e.g., real-world data), one can use APT to find the unknown parameter θ^* corresponding to the QML model or metrology problem. (d) Error mitigation. Consider the parameterized class of noise channels $\{\mathcal{E}_\theta\}_\theta$ that can be mitigated via classical post-processing when θ is known. In addition, suppose we obtain noisy data $\hat{o}_{\text{real}}$ by running classically-simulable circuits on our noisy device and ideal data $\hat{o}_{\text{ideal}}$ by classically simulating the same circuits. Passing the concept class and data $\{(\hat{o}_{\text{real}}, \hat{o}_{\text{ideal}})\}_i$ into an APT algorithm, we learn the parameter θ that best models the noise. Then, on a new quantum circuit run on our noisy device, we can mitigate the noise via classical post-processing from our knowledge of θ .

observables on output states of parameterized quantum circuits, which can be modelled as quantum statistical queries [22,50,51,55,64,65]. Thus, agnostic process tomography from quantum statistical queries could be a promising direction for new insights in QML. This application is depicted in Fig. 2(c).

4. Agnostic quantum metrology

Quantum metrology is considered one of the most significant real-world applications of quantum technology [16]. Here, one aims to estimate an unknown parameter θ with precision limited by the fundamental bounds of measurement [15], given access to a known class

of parameterized quantum channels $\{\mathcal{E}_\theta\}_\theta$ called the *encoding*. By probing the channel and performing coherent measurements, one can construct an estimator $\hat{\theta}$ of the unknown parameter with desirable statistical properties. However, in practice, the encoding may be corrupted by unknown environmental noise. Due to the challenges of characterizing noise and its effect on parameter estimation, noisy quantum metrology is less well-studied, and results are restricted to specific classes of noise [66–72]. Many standard protocols are also highly noise-sensitive [16,73–75]. Considering the known class $\{\mathcal{E}_\theta\}_\theta$ as our concept class, one can model the problem of extracting a desirable estimator as an agnostic process tomography problem, which may provide us with new insights into metrology in noisy settings. While noiseless metrology is close to conventional tomography, the parameter-encoding channel becomes unknown when transformed by realistic noise, which is more suitable for agnostic learning. This perspective is also applicable in adversarial scenarios, where an unknown attack channel affects the encoding [76,77]. We are unsure how this approach would compare to the standard quantum limit (the Heisenberg limit is unattainable for general noise without quantum error correction [78]) but nevertheless believe this is interesting to explore. This application is depicted in Fig. 2(c).

5. Error mitigation

Error mitigation protocols are often based on machine learning [79,80] or require a strong characterization of the noise channel beforehand [12,81]. We argue that agnostic process tomography can be used to generalize a machine-learning-based error mitigation protocol called Clifford data regression [79]. In this setting, we wish to mitigate a fixed, arbitrary noise channel $\mathcal{E}$ on a quantum device. We assume that $\mathcal{E}$ is well-approximated by a parameterized class of noise channels $\{\mathcal{E}_\theta\}_\theta$ such that, when θ is known, the errors can be mitigated via classical post-processing. This is relevant when methods that map arbitrary noise to a known class, e.g., Pauli twirling [81–84], are unavailable. We propose the following protocol: (1) Execute various classically-simulable circuits on the quantum device to obtain noisy data. (2) Simulate the same circuits classically to obtain ideal, noiseless data. (3) Using the noisy and ideal data, perform agnostic process tomography with respect to $\{\mathcal{E}_\theta\}_\theta$. This procedure outputs an approximation $\mathcal{E}_{\theta'}$ of the best representative of the arbitrary noise channel $\mathcal{E}$ from $\{\mathcal{E}_\theta\}_\theta$. Then, when implementing new circuits on this device, we can use $\mathcal{E}_{\theta'}$ to mitigate noise via classical post-processing. Our new protocol can be seen as providing the pre-characterization of error often needed for error mitigation protocols such as probabilistic error cancellation [81]. Note that in this setting, our access to the unknown process is through the pairs of ideal and noisy data rather than query access.

We consider the class of Pauli channels $\sum_x p_x \sigma_x(\cdot) \sigma_x$, parameterized by an underlying probability distribution $\{p_x\}_x$, as an example of our protocol for error mitigation. Our agnostic tomography algorithm for Pauli channels in Appendix D3 only requires data obtained by performing measurements in the basis of Choi states of the Pauli operators. The approximated noise can then be used to mitigate errors in new data using standard techniques for Pauli channel noise [81,85]. Note a limitation is that our agnostic learning algorithm for Pauli channels requires the use of ancilla qubits. This example also indicates that for this protocol to work, agnostic process tomography cannot be used in a black box manner. Instead, the agnostic algorithm should dictate what data to gather from the noisy quantum device. While we do not show an immediate application, we believe that this approach can result in new protocols for error mitigation for specific classes of noise. This application is depicted in Fig. 2(d).

III. RESULTS

We design new efficient agnostic learning algorithms for a wide variety of concept classes. We consider the setting where we are given query access to the unknown quantum channel, i.e., we can apply the unknown channel to any (potentially adaptively-chosen) input state. Moreover, we consider the error function to be given by either the root mean squared trace distance d_{avg} or the normalized Frobenius distance d_F , depending on the concept class. The root mean squared trace distance is defined as

$$d_{\text{avg}}(\mathcal{E}_1, \mathcal{E}_2) \triangleq \sqrt{\mathbb{E}_{|\psi\rangle \sim \mu_S} [d_{\text{tr}}(\mathcal{E}_1(|\psi\rangle\langle\psi|), \mathcal{E}_2(|\psi\rangle\langle\psi|))^2]}, \quad (\text{III.1})$$

where the expectation is taken over Haar-random states, and d_{tr} denotes the trace distance $d_{\text{tr}}(\rho, \sigma) \triangleq \|\rho - \sigma\|_1 / 2$. For operators on an n -qubit space, the normalized Frobenius distance is

$$d_F(A, B) \triangleq \frac{1}{\sqrt{2^n}} \|A - B\|_F, \quad (\text{III.2})$$

where $\|\cdot\|_F$ is the Frobenius norm (see Appendix A2). Meanwhile, for superoperators in an n -qubit space, the normalized Frobenius distance is

$$d_F(\Phi, \Psi) \triangleq \frac{1}{2^n \sqrt{2}} \|J(\Phi) - J(\Psi)\|_F, \quad (\text{III.3})$$

where $J(\Phi)$ denotes the Choi representation of Φ , i.e., the unnormalized Choi state. Note that all of the above distances are average-case. While one may consider stronger worst-case distances, e.g., diamond distance, for most of the classes we consider, even existing realizable algorithms

TABLE I. Summary of results. We omit the δ -dependence of the sample and time complexities, all of which scale with $\log(1/\delta)$. Here, n denotes the system size and ϵ denotes the approximation error from the optimal representative of the given concept class. The error is measured with respect to either d_{avg} (Definition A11) or d_F (Definitions A8 and A9). AUE and APT from AST denotes our results in converting an agnostic state tomography algorithm into one for agnostic unitary estimation and process tomography, respectively. Here, N/T denotes the sample and time complexity, respectively, of the underlying agnostic state tomography algorithm. The complexities are made explicit for specific cases in Appendices C2 and D2. For APT of G -gate circuits, $\text{opt}^{1/4}$ indicates that we get the bound $\text{err} \leq \text{opt}^{1/4} + \epsilon$ and similarly for APT from AST. For QAC^0 channels, d/s is the depth and size of the underlying QAC^0 circuit, respectively.

Concept Class	Result	Queries (N)	Time
G-gate Circuits (Corollary C2)	Proper 1-agnostic (d_{avg})	$\tilde{\mathcal{O}}(G \log(n)/\epsilon^4)$	Inefficient
AUE from AST (Appendix C2)	Improper 1-agnostic (d_F)	$N(2n, \epsilon^2)$	$T(2n, \epsilon^2)$
Pauli Strings (Theorem C1)	Proper 1-agnostic (d_{avg})	$\mathcal{O}(1/\epsilon^4)$	$\mathcal{O}(N)$
Unitary k -Juntas (Theorem C2)	Proper 2-agnostic (d_F)	$\mathcal{O}\left(\frac{k \log(n)}{\epsilon^4} + \frac{k16^k}{\epsilon^2}\right)$	$\mathcal{O}\left(\frac{n^k k \log(n)}{\epsilon^4}\right)$
Degree- d Unitaries (Theorem C3)	Improper 1-agnostic (d_F)	$\mathcal{O}(d \cdot 9^d n^{2d} \log(n)/\epsilon^2)$	$\mathcal{O}(N)$
(a) Summary of agnostic unitary estimation guarantees.			
Concept Class	Result	Queries (N)	Time
G-gate Circuits (Corollary D2)	Proper, $\text{opt}^{1/4}(d_{\text{avg}})$	$\tilde{\mathcal{O}}(G \log(n)/\epsilon^4)$	Inefficient
APT from AST (Appendix D2)	Improper, $2^{1/4} \text{opt}^{1/2}(d_F)$	$N(2n, \epsilon^2)$	$T(2n, \epsilon^2)$
Pauli Strings (Theorem D1)	Proper 1-agnostic (d_F)	$\mathcal{O}(1/\epsilon^4)$	$\mathcal{O}(N)$
Pauli Channels (Theorem D2)	Proper 1-agnostic (d_F)	$\mathcal{O}(1/\epsilon^2)$	$\mathcal{O}(N \log N)$
k -Junta Channels (Theorem D3)	Improper 1-agnostic (d_F)	$\mathcal{O}\left(\frac{k2^{12k} n^k \log n}{\epsilon^4}\right)$	$\mathcal{O}(N)$
Degree- d Channels (Theorem D4)	Improper 1-agnostic (d_F)	$\mathcal{O}(d \cdot 81^d n^{4d} \log(n)/\epsilon^2)$	$\mathcal{O}(N)$
QAC^0 Channels (Theorem D5)	Improper 1-agnostic (d_F)	$n^{\mathcal{O}(\log^d(s^4/\epsilon^4))}$	$n^{\mathcal{O}(\log^d(s^4/\epsilon^4))}$
(b) Summary of agnostic process tomography guarantees.			

only succeed with respect to such average-case distances [24,25,27,30].

For agnostic unitary estimation, we consider the classes of circuits consisting of G two-qubit gates [27], Pauli strings, unitary k -juntas, and degree- d unitaries. Pauli strings are unitaries of the form $\{I, X, Y, Z\}^{\otimes n}$. Unitary k -juntas are unitaries acting nontrivially only on k out of n qubits (Definition C2). Also, degree- d unitaries are unitaries whose Pauli coefficients of order greater than d are 0 (Definition C3). We note that, typically, k and d are considered at most logarithmic in n , making our algorithms computationally and statistically efficient (see Table I).

For agnostic process tomography, we consider similar classes, namely circuits consisting of G gates, Pauli strings, Pauli channels, k -junta channels, degree- d channels, and QAC^0 channels [22]. Here, Pauli channels are defined as those of the form $\sum_x p_x \sigma_x(\cdot) \sigma_x$ for a probability distribution $\{p_x\}_x$. k -junta channels and degree- d channels are defined similarly to the unitary case (see Definitions D2 and D3). Finally, we also consider n -to-1 QAC^0 channels [22], which are defined by applying a unitary implemented by a QAC^0 circuit [86] and tracing out all but one qubit (see Definition D5). Our algorithms for Pauli channels and juntas can be readily applied to the

resource-constrained implementation and error mitigation applications described in Sec. II A.

Moreover, under compatibility conditions on the concept classes, we show that one can convert any agnostic state tomography algorithm (with respect to fidelity) into agnostic unitary estimation and process tomography algorithms (with respect to the normalized Frobenius distance d_F) by learning the Choi state of the unknown unitary or channel. More precisely, if $\mathcal{C}$ is the concept class for agnostic state tomography, this conversion holds when, for each unitary U in the concept class for agnostic unitary estimation and process tomography, the Choi state of U is in the class $\mathcal{C}$. Note that for agnostic process tomography, the concept class must be unitary. This immediately gives us 1-agnostic unitary estimation algorithms for Clifford circuits, Clifford circuits with few T gates, and circuits consisting of tensor products of single-qubit gates from existing agnostic state tomography algorithms [43] with a quadratic increase with respect to the error ϵ . We obtain similar results for agnostic process tomography with a weaker guarantee on the error ($\text{err} \leq 2^{1/4} \text{opt}^{1/2} + \epsilon$). We note that we can avoid the quadratic increase in ϵ and obtain 1-agnostic unitary estimation and process tomography with respect to entanglement infidelity, defined as the infidelity between two Choi states [Eq. (C30)]. This measure has been considered

widely in (realizable) unitary estimation [6,87–93], but we focus on d_F for consistency with our other results.

In Table I, we summarize our results for agnostic unitary estimation and process tomography, respectively. Aside from the G -gate circuit case, all of our algorithms are sample and time-efficient.

For both agnostic unitary estimation and agnostic process tomography, the algorithms for circuits consisting of G gates are sample-efficient but computationally inefficient. In fact, we prove a more general result (see Propositions 1 and 3): for any finite-sized class $\mathcal{C}$, one can design an agnostic learning algorithm that utilizes $\mathcal{O}(\log |\mathcal{C}|)$ queries to the unknown quantum process but has time complexity scaling with $|\mathcal{C}|$. Qualitatively, this tells us that efficient agnostic tomography is easy from a sample complexity perspective, but it can be difficult to achieve with respect to time complexity. This is similar to an analogous argument for agnostic state tomography [42] based on shadow tomography [94]. Thus, for agnostic process tomography, one should also focus on computational efficiency rather than just sample efficiency. This is in contrast to the realizable setting, where results typically only focus on sample complexity.

IV. Discussion

Our work initiates the study of agnostic process tomography, which we define as learning the best approximation of an unknown process from a simple class. We design efficient agnostic learning algorithms for various concept classes, leveraging the powerful tools of Pauli spectrum analysis. We also propose several near-term applications for agnostic process tomography. We believe there is still substantial work to be done in this area, and we outline some open questions below.

First, can one design an efficient algorithm to map an arbitrary superoperator to its closest quantum channel? This is desirable because many of our improper agnostic learning algorithms can be made proper using such a subroutine, and many applications of agnostic process tomography (Sec. II A) require proper learning algorithms. As a significant first step towards this question, we have reduced the problem to that of proving a particular mathematical condition holds (see Appendix D 7) but were unable to prove or disprove it. Solving this question would be a significant result, as it would also lead to proper learning algorithms for realizable process tomography problems, such as those in Refs. [22,30].

Second, can we improve the error bounds for agnostic process tomography of bounded-gate circuits and agnostic process tomography from agnostic state tomography? Notably, as a result of the distance inequalities we used, our learning algorithms achieve error bounds of $\text{opt}^{1/4}$ and $2^{1/4}\text{opt}^{1/2}$, respectively, rather than opt . We believe these

bounds can be improved to 1-agnostic process tomography and leave this as an open problem. Also, can we avoid the quadratic increase in error when converting agnostic state tomography algorithms to those for agnostic process tomography, as in quantum process tomography [6]?

Third, is it possible to develop efficient agnostic learning algorithms for juntas that match the realizable lower bounds [24,25]? Our proposed algorithms for juntas have worse complexities than the realizable setting due to needing to estimate the weight of subsets of qubits, rather than the influence.

Next, for learning low-degree quantum objects, can one obtain better complexities by using the algorithm of Ref. [30] in the agnostic setting? We show in Appendix C 5 that straightforwardly utilizing the algorithm from Ref. [30] results in an inefficient algorithm with worse performance than ours. However, we believe that by proving tighter upper bounds on the norms in the BH inequalities, this could be improved.

Finally, we list several broad open questions. Could different definitions of agnostic tomography lead to new applications, e.g., comparing the performances of agnostic learners for different concept classes? For instance, the definition of Ref. [49] may lend itself to this, where they define agnostic learning with respect to two classes rather than just one. Can one develop agnostic learning algorithms with respect to worst-case distance metrics, e.g., diamond distance? All results in our work only hold for average-case distances. Also, many of our proposed algorithms require preparing the Choi state of the unknown channel. Instead, can one develop agnostic learning algorithms with fewer resources? As agnostic process tomography is at least as hard as in the realizable setting, lower bounds in realizable process tomography can be transferred to the agnostic setting as well. Can one prove stronger worst-case agnostic lower bounds? Finally, are there other physically-motivated classes that admit efficient agnostic process tomography algorithms? Some interesting classes include the evolution of well-studied Hamiltonians [95–98] or classes of parameterized quantum circuits [60] with specific ansatz [99–101].

V. Methods

Fourier analysis of Boolean functions [102] is widely used in theoretical computer science and classical learning theory. In particular, it has been used to develop efficient learning algorithms for functions with a sufficiently sparse Fourier decomposition [103–105] and in classical agnostic learning [47,106–108].

Pauli spectrum analysis [28] generalizes Fourier analysis for quantum operators (see Appendix A 2 for a more thorough discussion). Analogous to the Fourier decomposition, the key idea is that one can uniquely decompose an operator A acting on an n -qubit system into a linear

combination of Pauli strings:

$$A = \sum_{x \in \{0,1,2,3\}^n} \hat{A}_x \sigma_x, \quad \hat{A}_x \triangleq \frac{1}{2^n} \text{tr}(A \sigma_x), \quad (\text{V.1})$$

where $\sigma_x \in \{I, X, Y, Z\}^{\otimes n}$. Recently, Ref. [25] further generalized this to the *Fourier spectrum of superoperators* (see Appendix A 3), where a superoperator Φ in an n -qubit space is uniquely decomposed as

$$\begin{aligned} \Phi &= \sum_{x,y \in \{0,1,2,3\}^n} \hat{\Phi}(x,y) \Phi_{x,y}, \\ \hat{\Phi}(x,y) &= \frac{1}{4^n} \text{tr}(J(\Phi_{x,y})^\dagger J(\Phi)), \quad \Phi_{x,y}(\rho) = \sigma_x \rho \sigma_y, \end{aligned} \quad (\text{V.2})$$

where $J(\Phi)$ denotes the Choi representation of Φ , i.e., the unnormalized Choi state. As in classical learning theory, many quantum learning algorithms utilize Pauli and Fourier spectrum analysis [22,24,25,28–30,40].

In our work, we leverage both formalisms. Most of our algorithms involve learning the significant Pauli and Fourier coefficients. However, we can only estimate these coefficients up to an additive error, which *a priori* may not satisfy the necessary properties of a valid unitary or channel, resulting in improper algorithms.

To resolve this, for agnostic unitary estimation, we “project” our improper hypothesis to its closest unitary using a technique from Ref. [23]. We show that this outputs the closest unitary with respect to the normalized Frobenius distance d_F , resulting in proper agnostic unitary estimation. However, this problem is more difficult for agnostic process tomography. We discuss a potential approach based on convex optimization [109] for projecting an improper hypothesis onto a valid quantum channel [110] in Appendix D 7. We are unable to fully solve the problem due to one difficult-to-verify condition for the convergence guarantees of Ref. [109] to hold (Conjecture 1). Nevertheless, we made substantial progress towards resolving this open question, which is also recurrent in the realizable process tomography literature [22,24,25]. Now, we summarize our results and the techniques used for each class.

A. Bounded-gate circuits

First, we consider any finite-sized class $\mathcal{C}$ of unitaries. Similarly to Ref. [27], we use Clifford classical shadows [111] to estimate the average infidelity between the unknown quantum channel Φ and all $U_i \in \mathcal{C}$. Then, we output the unitary in $\mathcal{C}$ with the smallest estimated distance. Thus, we solve proper agnostic unitary estimation for $\mathcal{C}$ with respect to the root mean squared trace distance d_{avg} and show this uses $\mathcal{O}(\log |\mathcal{C}|)$ queries to Φ and time $\mathcal{O}(|\mathcal{C}|)$. For many classes of interest, this results in a

sample-efficient but computationally inefficient algorithm, emphasizing the importance of computational efficiency for agnostic learning. We show a weaker version of this result for agnostic process tomography of $\mathcal{C}$, which stems from converting from average infidelity to average trace distance for mixed inputs. As a corollary, by utilizing known covering number bounds [27], we obtain agnostic learning algorithms for the class of unitaries implementable by G two-qubit gates. Note that in the realizable setting, [27] proves computational lower bounds, implying that the time complexity for these algorithms cannot be improved further. We prove these results in Appendices C 1 and D 1.

B. Agnostic process tomography from agnostic state tomography

Recently, agnostic state tomography for several concept classes with respect to fidelity [42,43] has been studied. We show that this translates into guarantees for agnostic unitary estimation and process tomography for a compatible concept class $\mathcal{C}$ of unitaries with respect to d_F with a quadratic increase in the error ϵ . In particular, we require that the Choi state of each $U \in \mathcal{C}$ is in the concept class for agnostic state tomography. Our algorithm is simple: apply the agnostic state tomography algorithm to the Choi state of the unknown quantum process. Specifically, we obtain improper 1-agnostic unitary estimation but a weaker error bound for agnostic process tomography (see Table I), again due to converting between infidelity and trace distance for mixed inputs. The key step in our proof shows a relationship between the infidelity and the normalized Frobenius distance between two Choi states. We prove this result in Appendices C 2 and D 2.

C. Pauli strings

Next, we consider the class of Pauli strings $\{I, X, Y, Z\}^{\otimes n}$. We use the fact that one can efficiently sample from the distribution formed by the Pauli and Fourier coefficients [28]. After collecting samples, we simply output the most frequently sampled Pauli string. Thus, we achieve proper 1-agnostic unitary estimation and process tomography. We prove this in Appendices C 3 and D 3.

D. Pauli channels

Pauli channels are a naturally occurring and widely studied [31–37,81,85,112–115] class of quantum channels. Pauli channels are of the form $\sum_x p_x \sigma_x(\cdot) \sigma_x$ for a probability distribution $\{p_x\}_x$. Similarly to Pauli strings, we can sample from $\{p_x\}_x$. Then, agnostic tomography of Pauli channels reduces to empirically estimating $\{p_x\}_x$. Thus, we obtain proper 1-agnostic process tomography of Pauli channels. We prove this in Appendix D 3.

E. Junta unitaries and channels

A k -junta unitary/channel is an n -qubit unitary/channel that acts only on k qubits (Definitions C2 and D2). [24,25] devised realizable algorithms for learning junta unitaries/channels. Here, they estimate the *influence* [Eq. (A16)] of each qubit, then select a highly influential subset of qubits and perform full state tomography on the Choi state restricted to these qubits.

In the agnostic setting, we show that the optimal error is not tightly characterized by the most influential subset, but by the highest *weight* subset (Definitions A5 and A7). In the realizable setting, the relevant subset of qubits has the same weight and influence. However, we show that these notions are distinct in the agnostic setting (Example C1). We also prove that the hypothesis's error is characterized by the chosen subset and depends on its weight, not its influence. Thus, our algorithms aim to identify the highest weight subset of qubits. However, identifying this subset requires more queries and time than identifying the most influential one, giving us worse complexities than the realizable algorithms.

In the unitary case, we first estimate the weights of all subsets of $[n]$ of cardinality k using samples from the distribution formed by the Pauli coefficients. Then, we consider the subset S with the greatest estimated weight. Next, we estimate all the coefficients of the unitary that are supported only on S . We then construct a hypothesis using these estimates and project it to a unitary as described previously. Despite being less efficient than tomography on qubits in S , our method makes it easier to obtain a proper hypothesis. Our method results in improper 1-agnostic unitary estimation of unitary juntas. After applying the unitary projection to the improper hypothesis, we obtain proper 2-agnostic unitary estimation. We prove this formally in Appendix C 4.

For agnostic tomography of junta channels, we find that identifying the subset with the highest Fourier weight is much harder than in the unitary case. Here, we can only sample from the distribution defined by the diagonal Fourier coefficients $\{\hat{\Phi}(x, x)\}_x$ while the weight also depends on the off-diagonal coefficients. As a result, our algorithm learns *all* low-degree coefficients and uses them to estimate the weights. Explicitly, we learn all coefficients $\hat{\Phi}(x, y)$ for $|x|, |y| \leq k$, where $|x|$ denotes the size of the support of the string $x \in \{0, 1, 2, 3\}^n$. From here, the algorithm is similar to the unitary junta case. We achieve improper 1-agnostic process tomography of junta channels. We discuss a potential approach for producing a proper hypothesis in Appendix D 7. We prove this in Appendix D 4.

F. Low-degree unitaries and channels

Low-degree unitaries and channels are defined as those with all their Pauli and Fourier weight on low-degree

terms. Explicitly, a unitary has degree d if when $|x| > d$, then $\hat{U}_x = 0$ and similarly for channels (Definitions C3 and D3). A straightforward idea is to learn all of their low-degree coefficients. It turns out that this simple algorithm results in improper 1-agnostic learning of low-degree objects. Our unitary projection method mentioned previously is computationally inefficient for low-degree unitaries and its output may not be low-degree, leaving us with an improper hypothesis.

Reference [30] develops realizable algorithms leveraging Bohnenblust-Hille (BH) inequalities for operators [116] and superoperators [30]. These inequalities provide upper bounds on norms of the spectra of low-degree objects. Using these bounds, Ref. [30] shows that it suffices to learn significantly fewer Pauli and Fourier coefficients, resulting in algorithms with complexities independent of n . Using similar arguments for agnostic unitary estimation, we show that the algorithm of Ref. [30] is inefficient and only performs better than the trivial algorithm when the optimal error of the class of low-degree unitaries is at most inverse exponential in n , i.e., when the problem is almost realizable. We remark that improved bounds for the BH inequality may result in better guarantees for this algorithm and leave this open for future work. As this algorithm does not perform well in the simpler unitary case, we do not analyze its performance for agnostic process tomography. We prove these results in Appendices C 5 and D 5.

G. QAC⁰ channels

The circuit class QAC⁰ consists of constant-depth circuits with controlled-Z gates and arbitrary single-qubit gates [86]. In Ref. [22], the authors study the class of channels formed by applying a QAC⁰ circuit and tracing out all but one qubit. Reference [22] shows that the Choi representations of such channels have Pauli weight concentrated on low-degree terms. Similarly to Ref. [22], we learn the low-degree coefficients of the Choi representation of an arbitrary channel using classical shadows [111], resulting in improper 1-agnostic process tomography for this class of QAC⁰ channels. We show this in Appendix D 6.

ACKNOWLEDGMENTS

The authors thank Richard R. Allen, Alexandru Cojocaru, Ryan Sweke, Bo Yang, and Sisi Zhou for insightful discussions. The authors are supported by the Quantum Advantage Pathfinder (QAP) with Grant No. EP/X026167/1, the Hub in Quantum Computing and Simulation with Grant No. EP/T001062/1, and the UK Engineering and Physical Sciences Research Council (EPSRC). L.L. is also supported by a Marshall Scholarship.

All authors contributed to the main ideas and the writing of the manuscript. C.W. and L.L. carried out the technical

proofs and detailed analysis. E.K. and M.D. supervised the project and provided conceptual guidance throughout.

DATA AVAILABILITY

No data were created or analyzed in this study.

APPENDIX A: PRELIMINARIES

For $k \in \mathbb{N}$, we write $[k] \triangleq \{1, \dots, k\}$. Throughout, let n denote the system size. Let $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$ denote the Hilbert space on these n qubits, and let $\mathcal{B}(\mathcal{H})$ denote the set of bounded linear operators on $\mathcal{H}$. Let $\mathcal{U}_d \subseteq \mathcal{B}(\mathcal{H})$ denote the set of d -dimensional unitaries. A map from $\mathcal{B}(\mathcal{H})$ to itself is known as a *superoperator* on $\mathcal{B}(\mathcal{H})$. A *quantum channel* on $\mathcal{B}(\mathcal{H})$ is a superoperator that is completely positive and trace-preserving (CPTP). For any unitary, we sometimes use a calligraphic font to denote the associated quantum channel.

1. Choi state

A useful tool we will need is the Choi-Jamiołkowski isomorphism [117,118], which maps between unitaries acting on the Hilbert space $\mathcal{H}$ and pure states in $\mathcal{H} \otimes \mathcal{H}$.

Definition A1 (Choi state of a unitary). Let U be a unitary acting on an n -qubit Hilbert space $\mathcal{H}$. Then, the *Choi state* is defined by applying U on one half of a maximally entangled state over $2n$ qubits. In particular, define the maximally entangled state

$$|\Phi^+\rangle = \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i, i\rangle. \quad (\text{A1})$$

Then, the Choi state is defined as

$$|v(U)\rangle \triangleq (U \otimes I)|\Phi^+\rangle = \frac{1}{\sqrt{2^n}} \sum_{i,j=0}^{2^n-1} U_{ij} |i, j\rangle. \quad (\text{A2})$$

The Choi state can be further generalized for quantum channels, mapping quantum channels to mixed states.

Definition A2 (Choi state of a quantum channel). Let $\mathcal{E}$ be a quantum channel acting on $\mathcal{B}(\mathcal{H})$. Then, the associated *Choi state* is defined by applying $\mathcal{E}$ on one half of a maximally entangled state over $2n$ qubits. In particular, the Choi state is defined as

$$v(\mathcal{E}) = (\mathcal{E} \otimes \mathcal{I})(|\Phi^+\rangle\langle\Phi^+|) = \frac{1}{2^n} \sum_{i,j=0}^{2^n-1} \mathcal{E}(|i\rangle\langle j|) \otimes |i\rangle\langle j|. \quad (\text{A3})$$

For an arbitrary superoperator Ψ on $\mathcal{B}(\mathcal{H})$, we will also consider the *Choi representation*, obtained by applying

Ψ on one half of the unnormalized, maximally entangled state, i.e.,

$$J(\Psi) = \sum_{i,j=0}^{2^n-1} \Psi(|i\rangle\langle j|) \otimes |i\rangle\langle j|. \quad (\text{A4})$$

The Choi representation of a quantum channel satisfies the following necessary and sufficient properties.

Lemma A1 (Properties of Choi representations of quantum channels [119]). Let Ψ be a superoperator from $\mathcal{B}(\mathcal{H}_1)$ to $\mathcal{B}(\mathcal{H}_2)$. Then Ψ is a quantum channel if and only if

$$J(\Psi) \succcurlyeq 0 \quad \text{and} \quad \text{tr}_1(J(\Psi)) = I_2, \quad (\text{A5})$$

where tr_1 denotes the partial trace over $\mathcal{H}_1$ and I_2 denotes the identity on $\mathcal{H}_2$.

2. Pauli spectrum of operators

A key technique we use throughout our proofs is the analysis of the Pauli spectrum of quantum operators, first introduced in Refs. [28] to generalize Boolean Fourier analysis [102] to the quantum setting. Pauli analysis was also studied in Refs. [22,24,29]. Throughout this work, we mostly follow the notation of Ref. [28,102].

The set of bounded linear operators $\mathcal{B}(\mathcal{H})$ can be viewed as a Hilbert space when endowed with the Hilbert-Schmidt inner product

$$\langle A, B \rangle \triangleq \text{tr}(A^\dagger B) = \sum_{i,j \in [2^n]} A_{ij}^* B_{ij}, \quad (\text{A6})$$

where we use A_{ij} to denote the (i, j) -th entry of A as a matrix. We use $\|\cdot\|_F$ to denote the norm induced by this inner product, namely the Hilbert-Schmidt or Frobenius norm. In particular, we can write

$$\|A\|_F^2 = \langle A, A \rangle = \sum_{i,j \in [2^n]} |A_{ij}|^2. \quad (\text{A7})$$

Let $\{\sigma_0, \sigma_1, \sigma_2, \sigma_3\}$ denote the Pauli matrices, i.e.,

$$\begin{aligned} \sigma_0 = I &= \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, & \sigma_1 = X &= \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \\ \sigma_2 = Y &= \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, & \sigma_3 = Z &= \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \end{aligned} \quad (\text{A8})$$

For a vector $x \in \{0, 1, 2, 3\}^n$, we denote $\sigma_x \triangleq \sigma_{x_1} \otimes \dots \otimes \sigma_{x_n}$ and write $\text{supp}(x) \triangleq \{i \in [n] : x_i \neq 0\}$. We will denote the *degree* of x by $|x| = |\text{supp}(x)|$. Because the Paulis $\{\sigma_x\}_{x \in \{0,1,2,3\}^n}$ form an orthogonal basis for $\mathcal{B}(\mathcal{H})$ with respect to the Hilbert-Schmidt inner product, then we can

uniquely decompose any operator with respect to this basis.

Definition A3 (Pauli spectrum [28]). For any $A \in \mathcal{B}(\mathcal{H})$, there exists a unique decomposition as

$$A = \sum_{x \in \{0,1,2,3\}^n} \hat{A}_x \sigma_x, \quad \hat{A}_x \triangleq \frac{1}{2^n} \langle \sigma_x, A \rangle. \quad (\text{A9})$$

Here, $\hat{A}_x$ is called the *Pauli coefficient* of A , and $\{\hat{A}_x\}_{x \in \{0,1,2,3\}^n}$ is called the *Pauli spectrum* of A .

As in classical Fourier analysis, there are quantum analogs of Parseval and Plancherel's identities.

Lemma A2 (Proposition 10 in Ref. [28]). Let $A, B \in \mathcal{B}(\mathcal{H})$. Then, Parseval's identity is

$$\frac{1}{2^n} \|A\|_F^2 = \sum_{x \in \{0,1,2,3\}^n} |\hat{A}_x|^2. \quad (\text{A10})$$

Moreover, Plancherel's identity is

$$\frac{1}{2^n} \langle A, B \rangle = \sum_{x \in \{0,1,2,3\}^n} \hat{A}_x^* \hat{B}_x. \quad (\text{A11})$$

Following the notation in Ref. [22,28], we define the *Pauli weight of an operator at level k* . This is similar to the classical definition of Fourier weight [102].

Definition A4 (Pauli weight at level k [22,28]). Let $A \in \mathcal{B}(\mathcal{H})$ and $0 \leq k \leq n$. Then, the *Pauli weight of A at level k* is defined by

$$w^{\leq k}(A) \triangleq \sum_{x \in \{0,1,2,3\}^n, |x|=k} |\hat{A}_x|^2. \quad (\text{A12})$$

We similarly define $w^{\leq k}(A)$, $w^{>k}(A)$, etc.

For a subset $S \subseteq [n]$, we also denote

$$\mathcal{P}_S \triangleq \{x \in \{0,1,2,3\}^n : \text{supp}(x) \subseteq S\}. \quad (\text{A13})$$

Notice that

$$|\mathcal{P}_S| = 4^{|S|}. \quad (\text{A14})$$

We also use the shorthand $\mathcal{P}_n \triangleq \mathcal{P}_{[n]} = \{0,1,2,3\}^n$. We also define the Pauli weight of $A \in \mathcal{B}(\mathcal{H})$ on a given subset $S \subseteq [n]$, similar to the Pauli weight at level k .

Definition A5 (Pauli weight on a subset). Let $A \in \mathcal{B}(\mathcal{H})$ and $S \subseteq [n]$. Then, the *Pauli weight of A on the subset S* is defined by

$$w_S(A) \triangleq \sum_{x \in \mathcal{P}_S} |\hat{A}_x|^2. \quad (\text{A15})$$

This definition is reminiscent of the influence on a set of qubits. Recall from Ref. [24] that the influence of a unitary U on a subset $S \subseteq [n]$ is defined as

$$\text{Inf}_S[U] = \sum_{\substack{x \in \{0,1,2,3\}^n \\ \text{supp}(x) \cap S \neq \emptyset}} |\hat{U}_x|^2. \quad (\text{A16})$$

Here, we see that the Pauli weight on S differs from the influence due to the set of $x \in \{0,1,2,3\}^n$ that the summation is over.

The Choi state of a unitary has a nice relationship with its Pauli decomposition. Let $U = \sum_{x \in \{0,1,2,3\}^n} \hat{U}_x \sigma_x$ be the Pauli decomposition of some unitary U . Then, the Choi state of U decomposes as

$$|v(U)\rangle = \sum_{x \in \{0,1,2,3\}^n} \hat{U}_x |v(\sigma_x)\rangle, \quad (\text{A17})$$

where $|v(\sigma_x)\rangle$ denotes the Choi state of the Pauli σ_x .

3. Fourier spectrum of superoperators

One can also extend the Pauli analysis from the previous section to superoperators, which was initiated in Ref. [25]. This is also explored in Ref. [22], but we mainly follow the presentation of Ref. [25]. The space of superoperators on $\mathcal{B}(\mathcal{H})$ can be viewed as a Hilbert space when endowed with the inner product

$$\langle \Phi, \Psi \rangle \triangleq \langle J(\Phi), J(\Psi) \rangle = \text{tr}(J(\Phi)^\dagger J(\Psi)). \quad (\text{A18})$$

This inner product induces the norm

$$\|\Phi\|_F \triangleq \sqrt{\langle J(\Phi), J(\Phi) \rangle} = \|J(\Phi)\|_F. \quad (\text{A19})$$

Consider the superoperators $\{\Phi_{x,y}\}_{x,y \in \mathcal{P}_n}$

$$\Phi_{x,y}(\rho) = \sigma_x \rho \sigma_y, \quad \forall x, y \in \mathcal{P}_n. \quad (\text{A20})$$

These superoperators form an orthogonal basis for the space of superoperators on $\mathcal{B}(\mathcal{H})$ with the norm given by

$$\|\Phi_{x,y}\|_F = 2^n, \quad \forall x, y \in \mathcal{P}_n. \quad (\text{A21})$$

Because $\{\Phi_{x,y}\}_{x,y}$ forms an orthogonal basis, we can uniquely decompose any superoperator with respect to this basis.

Definition A6 (Fourier expansion of superoperators; Definition 7 in Ref. [25]). For any superoperator Φ on $\mathcal{B}(\mathcal{H})$, there exists a unique decomposition as

$$\Phi = \sum_{x,y \in \mathcal{P}_n} \hat{\Phi}(x,y) \Phi_{x,y}, \quad \hat{\Phi}(x,y) = \frac{1}{4^n} \langle \Phi_{x,y}, \Phi \rangle. \quad (\text{A22})$$

Here, $\hat{\Phi}_{x,y}$ is a Fourier coefficient of Φ , and $\{\hat{\Phi}_{x,y}\}_{x,y \in \mathcal{P}_n}$ is called the *Fourier spectrum* of Φ . Denote by $\hat{\Phi}$ the $4^n \times 4^n$ -matrix with entries $\hat{\Phi}(x,y)$.

The matrix of Fourier coefficients $\hat{\Phi}$ satisfies the following properties.

Lemma A3 (Properties of Fourier coefficients; Lemma 8 in Ref. [25]). For any superoperator Φ on $\mathcal{B}(\mathcal{H})$, the matrix of Fourier coefficients $\hat{\Phi}$ satisfies the following properties.

(1) There exists a unitary U such that

$$\hat{\Phi} = \frac{1}{2^n} U^\dagger J(\Phi) U. \quad (\text{A23})$$

(2) Then, $\hat{\Phi} \succcurlyeq 0$ if and only if $J(\Phi) \succcurlyeq 0$.

(3) If Φ is a quantum channel, then

$$0 \leq \hat{\Phi}(x,x) \leq 1, \quad \forall x \in \mathcal{P}_n \quad \text{and} \quad \sum_{x \in \mathcal{P}_n} \hat{\Phi}(x,x) = \text{tr}(\hat{\Phi}) = 1. \quad (\text{A24})$$

We remark that for a unitary with Pauli spectrum $\{\hat{U}_x\}_x$, the Fourier coefficients of the associated quantum channel Φ are given by

$$\hat{\Phi}(x,y) = \hat{U}_x \hat{U}_y^*. \quad (\text{A25})$$

Similar to the case of the Pauli spectrum, we once again have analogs of Parseval's and Plancherel's identities.

Lemma A4 (Parseval's and Plancherel's for superoperators). Let Φ, Ψ be quantum channels on $\mathcal{B}(\mathcal{H})$. Then, Parseval's identity is

$$\frac{1}{4^n} \|\Phi\|_F^2 = \sum_{x,y \in \mathcal{P}_n} |\hat{\Phi}(x,y)|^2. \quad (\text{A26})$$

Moreover, Plancherel's identity is

$$\frac{1}{4^n} \langle \Phi, \Psi \rangle = \sum_{x,y \in \mathcal{P}_n} \hat{\Phi}(x,y)^* \hat{\Psi}(x,y). \quad (\text{A27})$$

Also similarly to the Pauli spectrum, we will define the weight at level k of a superoperator.

Definition A7 (Fourier weight at level k). Let Φ be a superoperator on $\mathcal{B}(\mathcal{H})$ and $0 \leq k \leq n$. Then, the *Fourier weight of Φ at level k* is defined by

$$w^k(A) \triangleq \sum_{x,y \in \mathcal{P}_n, |x|=|y|=k} |\hat{\Phi}(x,y)|^2. \quad (\text{A28})$$

The weight of a subset, the weights of order at least k , and so on can be defined similarly to Definitions A5 and A4.

The Choi state of a quantum channel Φ also has a nice relationship with the Fourier coefficients. In particular, the Choi representation is given by

$$J(\Phi) = \sum_{x,y \in \mathcal{P}_n} \hat{\Phi}(x,y) J(\Phi_{x,y}), \quad (\text{A29})$$

and the Choi state is given by

$$\begin{aligned} v(\Phi) &= \frac{1}{2^n} J(\Phi) = \frac{1}{2^n} \sum_{x,y \in \mathcal{P}_n} \hat{\Phi}(x,y) J(\Phi_{x,y}) \\ &= \sum_{x,y \in \mathcal{P}_n} \hat{\Phi}(x,y) |v(\sigma_x)\rangle \langle v(\sigma_y)|, \end{aligned} \quad (\text{A30})$$

which can be seen by normalizing $J(\Phi)$ by 2^n , as Φ is trace-preserving.

4. Distance metrics and covering nets

In this section, we review definitions of various distance metrics and important properties needed for our proofs. In Appendix A2, we introduced the Frobenius norm $\|A\|_F^2 = \text{tr}(A^\dagger A)$. We now define normalized variants of the Frobenius distance.

Definition A8 (Normalized Frobenius distance between operators). Let $A, B \in \mathcal{B}(\mathcal{H})$. We define the *normalized Frobenius distance* as

$$d_F(A, B) \triangleq \frac{1}{\sqrt{2^{n+1}}} \|A - B\|_F. \quad (\text{A31})$$

We have the following relation between the Pauli spectra of two operators and their normalized Frobenius distance.

Lemma A5 (Frobenius distance and Pauli coefficients). For any $A, B \in \mathcal{B}(\mathcal{H})$,

$$d_F(A, B)^2 = \frac{1}{2} \sum_{x \in \mathcal{P}_n} |\hat{A}_x - \hat{B}_x|^2. \quad (\text{A32})$$

Proof. This follows by direct computation.

$$d_F(A, B)^2 = \frac{1}{2^{n+1}} \|A - B\|_F^2 \quad (\text{A33})$$

$$= \frac{1}{2^{n+1}} \left(\sum_{i,j \in [2^n]} |A_{ij} - B_{ij}|^2 \right) \quad (\text{A34})$$

$$= \frac{1}{2^{n+1}} \left(\sum_{i,j \in [2^n]} |A_{ij}|^2 + |B_{ij}|^2 - A_{ij}^* B_{ij} - A_{ij} B_{ij}^* \right) \quad (\text{A35})$$

$$= \frac{1}{2^{n+1}} (\|A\|_F^2 + \|B\|_F^2 - \langle A, B \rangle - \langle B, A \rangle) \quad (\text{A36})$$

$$= \frac{1}{2} \left(\sum_{x \in \mathcal{P}_n} |\hat{A}_x|^2 + |\hat{B}_x|^2 - \hat{A}_x^* \hat{B}_x - \hat{A}_x \hat{B}_x^* \right) \quad (\text{A37})$$

$$= \frac{1}{2} \sum_{x \in \mathcal{P}_n} |\hat{A}_x - \hat{B}_x|^2. \quad (\text{A38})$$

Here, the second and fourth lines follow by definition of the Frobenius norm and the Hilbert-Schmidt inner product. The second to last line follows from Lemma A2. ■

Based on the notion of a Frobenius norm for superoperators introduced in Appendix A3, we now define a similar distance between superoperators.

Definition A9 (Normalized Frobenius distance between superoperators [25]). Let Φ, Ψ be two superoperators on $\mathcal{B}(\mathcal{H})$. Then, we define the normalized Frobenius distance as

$$d_F(\Phi, \Psi) \triangleq \frac{1}{2^n \sqrt{2}} \|\Phi - \Psi\|_F \triangleq \frac{1}{2^n \sqrt{2}} \|J(\Phi) - J(\Psi)\|_F. \quad (\text{A39})$$

Note that the normalization ensures that the distance is in $[0, 1]$ for two quantum channels. We now show a relation between this distance and the Fourier spectrum of superoperators.

Lemma A6 (Normalized Frobenius distance of superoperators and Fourier coefficients; Corollary 28 in Ref. [25]). Let Φ, Ψ be two superoperators on $\mathcal{B}(\mathcal{H})$. Then,

$$d_F(\Phi, \Psi)^2 = \frac{1}{2} \sum_{x,y \in \mathcal{P}_n} |\hat{\Phi}(x,y) - \hat{\Psi}(x,y)|^2. \quad (\text{A40})$$

Proof. We give an explicit proof via computation.

$$d_F(\Phi, \Psi)^2 = \frac{1}{2^{2n+1}} \|\Phi - \Psi\|_F^2 \quad (\text{A41})$$

$$= \frac{1}{2^{2n+1}} \langle \Phi - \Psi, \Phi - \Psi \rangle \quad (\text{A42})$$

$$= \frac{1}{2^{2n+1}} (\|\Phi\|_F^2 + \|\Psi\|_F^2 - \langle \Phi, \Psi \rangle - \langle \Psi, \Phi \rangle) \quad (\text{A43})$$

$$= \frac{1}{2} \sum_{x,y \in \mathcal{P}_n} \left(|\hat{\Phi}(x,y)|^2 + |\hat{\Psi}(x,y)|^2 - \hat{\Phi}(x,y)^* \hat{\Psi}(x,y) - \hat{\Phi}(x,y) \hat{\Psi}(x,y)^* \right) \quad (\text{A44})$$

$$= \frac{1}{2} \sum_{x,y \in \mathcal{P}_n} |\hat{\Phi}(x,y) - \hat{\Psi}(x,y)|^2, \quad (\text{A45})$$

where the first three lines follow from the definitions of the normalized Frobenius distance and Hilbert-Schmidt norm for superoperators, and the fourth line follows from Lemma A4. ■

We now define distances between two quantum states, which we will shortly use to define average-case distances between quantum channels.

Definition A10 (Trace distance and fidelity). The trace distance between two quantum states is given by

$$d_{\text{tr}}(\rho, \sigma) \triangleq \frac{1}{2} \|\rho - \sigma\|_1. \quad (\text{A46})$$

The fidelity between two quantum states is given by

$$F(\rho, \sigma) \triangleq \text{tr} \left(\sqrt{\sqrt{\rho} \sigma \sqrt{\rho}} \right)^2. \quad (\text{A47})$$

We call $1 - F$ the *infidelity* of two states.

The trace distance and fidelity are related in the following ways [120]. For any two quantum states ρ, σ ,

$$1 - \sqrt{F(\rho, \sigma)} \leq d_{\text{tr}}(\rho, \sigma) \leq \sqrt{1 - F(\rho, \sigma)}. \quad (\text{A48})$$

When one of the states is pure, e.g., $\sigma = |\psi\rangle\langle\psi|$, then

$$1 - F(|\psi\rangle\langle\psi|, \rho) \leq d_{\text{tr}}(|\psi\rangle\langle\psi|, \rho). \quad (\text{A49})$$

When both states are pure, then

$$1 - F(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) = d_{\text{tr}}(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|)^2. \quad (\text{A50})$$

In addition to the normalized Frobenius distance, we can also consider average-case distances. We follow Ref. [27]

to define an average distance between two quantum channels with respect to a measure.

Definition A11 (Average distance [27]). For any measure μ over d -dimensional states, we define the *average distance* between two channels d_μ as root mean squared trace distance of output states over inputs from μ

$$d_\mu(\mathcal{E}_1, \mathcal{E}_2) \triangleq \sqrt{\mathbf{E}_{\rho \sim \mu} [d_{\text{tr}}(\mathcal{E}_1(\rho), \mathcal{E}_2(\rho))^2]}. \quad (\text{A51})$$

In particular, when the average is taken over the Haar measure μ_S , we denote the distance by d_{avg}

$$d_{\text{avg}}(\mathcal{E}_1, \mathcal{E}_2) \triangleq \sqrt{\mathbf{E}_{|\psi\rangle \sim \mu_S} [d_{\text{tr}}(\mathcal{E}_1(|\psi\rangle\langle\psi|), \mathcal{E}_2(|\psi\rangle\langle\psi|))^2]}. \quad (\text{A52})$$

Proof.

$$d_\mu(\mathcal{E}_1, \mathcal{E}_3)^2 = \mathbf{E}_{\rho \sim \mu} [d_{\text{tr}}(\mathcal{E}_1(\rho), \mathcal{E}_3(\rho))^2] \quad (\text{A54})$$

$$\leq \mathbf{E}_{\rho \sim \mu} [(d_{\text{tr}}(\mathcal{E}_1(\rho), \mathcal{E}_2(\rho)) + (d_{\text{tr}}(\mathcal{E}_2(\rho), \mathcal{E}_3(\rho)))^2] \quad (\text{A55})$$

$$= d_\mu(\mathcal{E}_1, \mathcal{E}_2)^2 + d_\mu(\mathcal{E}_2, \mathcal{E}_3)^2 + 2\mathbf{E}_{\rho \sim \mu} [d_{\text{tr}}(\mathcal{E}_1(\rho), \mathcal{E}_2(\rho))d_{\text{tr}}(\mathcal{E}_2(\rho), \mathcal{E}_3(\rho))] \quad (\text{A56})$$

$$\leq d_\mu(\mathcal{E}_1, \mathcal{E}_2)^2 + d_\mu(\mathcal{E}_2, \mathcal{E}_3)^2 + 2\sqrt{\mathbf{E}_{\rho \sim \mu} [d_{\text{tr}}(\mathcal{E}_1(\rho), \mathcal{E}_2(\rho))^2]} \sqrt{\mathbf{E}_{\rho \sim \mu} [d_{\text{tr}}(\mathcal{E}_2(\rho), \mathcal{E}_3(\rho))^2]} \quad (\text{A57})$$

$$= (d_\mu(\mathcal{E}_1, \mathcal{E}_2) + d_\mu(\mathcal{E}_2, \mathcal{E}_3))^2, \quad (\text{A58})$$

where in the first inequality we use the triangle inequality of trace distance, and in the last inequality we use the Cauchy-Schwarz inequality. $\blacksquare$

We can also relate the average-case distance d_{avg} to the Pauli coefficients via the following lemma.

Lemma A8 (Average distance and Pauli coefficients). For any n -qubit unitaries U, V , we have

$$d_{\text{avg}}(U, V)^2 = \frac{2^n}{2^n + 1} \left(1 - \left| \sum_{x \in \mathcal{P}_n} \hat{U}_x^* \hat{V}_x \right|^2 \right). \quad (\text{A59})$$

Proof. By properties of the Haar integral (e.g., Example 48 in Ref. [123]), we have

$$d_{\text{avg}}(U, V)^2 = \frac{4^n - |\text{tr}(U^\dagger V)|^2}{2^n(2^n + 1)}. \quad (\text{A60})$$

Then, by Plancherel's identity (Lemma A2), we have

$$\text{tr}(U^\dagger V) = \langle U, V \rangle = 2^n \sum_{x \in \mathcal{P}_n} \hat{U}_x^* \hat{V}_x. \quad (\text{A61})$$

However, because Haar-random states are difficult to generate, one may turn to other notions of random states with respect to which to measure an average-case distance. One such physically motivated class of states is locally scrambled ensembles [27, 121, 122].

We will make use of the fact that for any measure μ over states, d_μ satisfies the triangle inequality. We will follow the proof of Ref. [27], where this was shown for unitaries.

Lemma A7 (Triangle inequality of d_{avg} ; Lemma 6 in Ref. [27]). For any quantum channels $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$ acting on $\mathcal{B}(\mathcal{H})$, and any measure μ over states in $\mathcal{B}(\mathcal{H})$, we have

$$d_\mu(\mathcal{E}_1, \mathcal{E}_3) \leq d_\mu(\mathcal{E}_1, \mathcal{E}_2) + d_\mu(\mathcal{E}_2, \mathcal{E}_3). \quad (\text{A53})$$

The result clearly follows from this. $\blacksquare$

A related topic is that of covering nets from high-dimensional probability theory. We can define a covering net with respect to any of the distance metrics introduced above.

Definition A12 (Covering net and number and metric entropy). Let (X, d) be a metric space. Let $K \subseteq X$ be a subset and $\epsilon > 0$. Then, define the following.

- (a) $N \subseteq K$ is an ϵ -covering net of K if for any $x \in K$, there exists a $y \in N$ such that $d(x, y) \leq \epsilon$.
- (b) The *covering number* $\mathcal{N}(K, d, \epsilon)$ of K is the smallest possible cardinality of an ϵ -covering net of K .
- (c) The *metric entropy* is $\log \mathcal{N}(K, d, \epsilon)$.

A useful result about covering nets that we will need in our proofs is bounds on the metric entropy of covering nets for unitaries of bounded gate complexity [27].

Lemma A9 (Covering nets for unitaries of bounded gates with respect to d_F ; Theorem 8 and Corollary 1 in Ref. [27]). Let $U^G \subseteq U(2^n)$ be the set of n -qubit unitaries that can be implemented by G two-qubit gates. Then for any $\epsilon \in (0, 1]$, there exist universal constants $c_1, c_2, C > 0$ such that for $1 \leq G/C \leq 4^{n+1}$, the metric entropy of U^G with respect to the normalized Frobenius distance d_F can be bounded as

$$\begin{aligned} \frac{G}{4C} \log\left(\frac{c_1}{\epsilon}\right) &\leq \log \mathcal{N}(U^G, d_F, \epsilon) \\ &\leq 16G \log\left(\frac{c_2 G}{\epsilon}\right) + 2G \log n. \end{aligned} \quad (\text{A62})$$

Moreover, the metric entropy of U^G with respect to the average-case distance d_{avg} can be bounded as

$$\begin{aligned} \frac{G}{4C} \log\left(\frac{c_1}{8\epsilon}\right) - \log\left(\frac{c_2}{2\epsilon}\right) &\leq \log \mathcal{N}(U^G, d_{\text{avg}}, \epsilon) \\ &\leq 16G \log\left(\frac{c_2 G}{\epsilon}\right) + 2G \log n. \end{aligned} \quad (\text{A63})$$

APPENDIX B: USEFUL LEMMAS

In this section, we compile several lemmas that we use repeatedly throughout our proofs. Our algorithms for both agnostic unitary estimation and agnostic process tomography rely on sampling from a probability distribution to estimate it empirically. In our agnostic unitary estimation algorithms, we will often sample from the distribution $\{|\hat{U}_x|^2\}_x$ formed by the Pauli coefficients to estimate it empirically. On the other hand, for APT, we will sample from the distribution formed by the diagonal Fourier coefficients $\{\hat{\Phi}(x, x)\}_x$. For guarantees on the estimation, we make use of the following well-known result from distribution learning theory. We refer to Theorem 9 in Ref. [124] for a proof and follow the presentation of Lemma 16 in Ref. [30].

Lemma B1 (Empirical estimation of a distribution; Theorem 9 in Ref. [124]). Let $p = \{p(x)\}_x$ be a probability distribution over some set $\mathcal{X}$. Let p' be the empirical probability distribution obtained after sampling N times from p . Then, $N = \mathcal{O}(\log(1/\delta)/\epsilon^2)$ samples suffice to have, with probability at least $1 - \delta$,

- (1) $|p(x) - p'(x)| \leq \epsilon, \forall x \in \mathcal{X}$,
- (2) $\sqrt{\sum_{x \in \mathcal{X}} |p(x) - p'(x)|^2} \leq \epsilon$.

Many of our agnostic unitary estimation algorithms require learning a subset of the Pauli coefficients of a unitary. In the following lemma, we show how any Pauli coefficient can be estimated. This algorithm extends the

technique introduced in Lemma 24 of Ref. [28] for learning Hermitian unitary operators, where the Pauli coefficients are all guaranteed to be real. However, in general, the Pauli coefficients may be complex. We resolve this by using the approach from Ref. [28] to separately estimate the real and imaginary parts of each Pauli coefficient. This extension is also proposed in Lemma 17 of Ref. [30] without proof. For completeness, we provide a proof here.

Lemma B2 (Estimating any Pauli coefficient of a unitary; Lemma 17 in Ref. [30]). Given query access to an n -qubit unitary U and a string $x \in \mathcal{P}_n$ as input, there exists an algorithm that produces an estimate α_x of the Pauli coefficient $\hat{U}_x$ using $N = \mathcal{O}(\log(1/\delta)/\epsilon^2)$ queries to U such that

$$|\alpha_x - \hat{U}_x| \leq \epsilon, \quad (\text{B1})$$

with probability at least $1 - \delta$. Moreover, the algorithm runs in computational time $\mathcal{O}(\log(1/\delta)/\epsilon^2)$.

Proof. As discussed above, the algorithm estimates the real and imaginary parts of each Pauli coefficient separately by the same approach as in Lemma 24 of Ref. [28]. To do so, we need access to controlled- U . This can be implemented using query access to U alone by preparing n ancilla qubits, applying a controlled-SWAP between the ancilla and main register, applying U to the ancillas, another controlled-SWAP, and finally discarding the ancillas.

First, we estimate the real part. We proceed in the same way as Lemma 24 of Ref. [28]. We first prepare $|0\rangle |\Phi^+\rangle$, where $|\Phi^+\rangle$ denotes a Bell pair. Then, we apply a Hadamard gate to get

$$\frac{1}{\sqrt{2}} |0\rangle |\Phi^+\rangle + \frac{1}{\sqrt{2}} |1\rangle |\Phi^+\rangle. \quad (\text{B2})$$

Applying controlled- σ_x controlled on $|0\rangle$, we have

$$\frac{1}{\sqrt{2}} |0\rangle |v(\sigma_x)\rangle + \frac{1}{\sqrt{2}} |1\rangle |\Phi^+\rangle. \quad (\text{B3})$$

Next, we apply controlled- U as described above to obtain the state

$$\frac{1}{\sqrt{2}} |0\rangle |v(\sigma_x)\rangle + \frac{1}{\sqrt{2}} |1\rangle |v(U)\rangle. \quad (\text{B4})$$

Then, we apply another Hadamard gate to the first qubit to obtain

$$\frac{1}{2} |0\rangle (|v(\sigma_x)\rangle + |v(U)\rangle) + \frac{1}{2} |1\rangle (|v(\sigma_x)\rangle - |v(U)\rangle). \quad (\text{B5})$$

Finally, we measure the first qubit in the computational basis, giving us outcome probabilities

$$\begin{aligned}\Pr[0] &= \frac{1}{4}(\langle v(\sigma_x) | + \langle v(U) |)(|v(\sigma_x)\rangle + |v(U)\rangle) \\ &= \frac{1}{2} + \frac{(\hat{U}_x + \hat{U}_x^*)}{4} = \frac{1}{2} + \frac{\text{Re}(\hat{U}_x)}{2},\end{aligned}\quad (\text{B6})$$

and similarly

$$\Pr[1] = \frac{1}{2} - \frac{\text{Re}(\hat{U}_x)}{2}. \quad (\text{B7})$$

Using Hoeffding's inequality, we see that by repeating this procedure $\mathcal{O}(\log(1/\delta)/\epsilon^2)$ many times, we can estimate $\text{Re}(\hat{U}_x)$ up to accuracy $\epsilon/\sqrt{2}$ with probability at least $1 - \delta/2$.

To estimate the imaginary part, we replace Eq. (B4) with the following:

$$\frac{1}{\sqrt{2}}|0\rangle|v(\sigma_x)\rangle + \frac{i}{\sqrt{2}}|1\rangle|v(U)\rangle.$$

We again perform a Hadamard operation followed by a measurement of the control qubit. Using a similar analysis to the real case, we obtain measurement probabilities

$$\Pr[0] = \frac{1}{2} - \frac{1}{2}\text{Im}(\hat{U}_x) \text{ and } \Pr[1] = \frac{1}{2} + \frac{1}{2}\text{Im}(\hat{U}_x). \quad (\text{B8})$$

Then, by repeating this procedure $\mathcal{O}(\log(1/\delta)/\epsilon^2)$ many times, we can now estimate $\text{Im}(\hat{U}_x)$ up to accuracy $\epsilon/\sqrt{2}$ with probability at least $1 - \delta/2$. Using a union bound, we see that with probability at least $1 - \delta$, we have an ϵ -accurate estimate of $\hat{U}_x$.

As we spend $\mathcal{O}(1)$ time processing each sample in both procedures, the computational time is of the same order as the sample complexity. ■

Similarly, we can also learn any Fourier coefficient of a quantum channel. The same result is in Lemma 15 in Ref. [30], but we give a different proof that allows for an easier analysis of the computational time. This is because our proof makes the state preparation and measurement steps more explicit than that of Ref. [30].

Lemma B3 (Estimating any Fourier coefficient of a quantum channel). Given query access to a quantum channel Φ on $\mathcal{B}(\mathcal{H})$, and strings $x, y \in \mathcal{P}_n$ as input, there exists a method to produce an estimate $\alpha_{x,y}$ of the Fourier coefficient $\hat{\Phi}(x, y)$ using $N = \mathcal{O}(\log(1/\delta)/\epsilon^2)$ queries to Φ such that

$$|\alpha_{x,y} - \hat{\Phi}(x, y)| \leq \epsilon \quad (\text{B9})$$

with probability at least $1 - \delta$. Moreover, the algorithm runs in computational time $\mathcal{O}(\log(1/\delta)/\epsilon^2)$.

Proof. Recall that the Choi state of a quantum channel can be expanded in terms of its Fourier coefficients and the Choi states of the Pauli operators as

$$v(\Phi) = \sum_{x,y \in \mathcal{P}_n} \hat{\Phi}(x, y) |v(\sigma_x)\rangle \langle v(\sigma_y)|. \quad (\text{B10})$$

We first estimate the diagonal coefficients $\hat{\Phi}(x, x)$ and $\hat{\Phi}(y, y)$. From Eq. (B10), we have

$$\langle v(\sigma_x) | v(\Phi) | v(\sigma_x) \rangle = \hat{\Phi}(x, x). \quad (\text{B11})$$

Then, similarly to Lemma B2, by conducting $\mathcal{O}(\log(1/\delta)/\epsilon^2)$ SWAP tests [125–127] between $v(\Phi)$ and $|v(\sigma_x)\rangle \langle v(\sigma_x)|$, we can estimate $\hat{\Phi}(x, x)$ up to accuracy $\epsilon/(2\sqrt{2})$ with probability at least $1 - \delta/4$, and similarly for $\hat{\Phi}(y, y)$. Then, consider the state

$$|\psi(x, y)^+\rangle \triangleq \frac{|v(\sigma_x)\rangle + |v(\sigma_y)\rangle}{\sqrt{2}}. \quad (\text{B12})$$

Using Eq. (B10), we have

$$\begin{aligned}\langle \psi(x, y)^+ | v(\Phi) | \psi(x, y)^+ \rangle &= \frac{(\langle v(\sigma_x) | + \langle v(\sigma_y) |) v(\Phi) (|v(\sigma_x)\rangle + |v(\sigma_y)\rangle)}{2} \quad (\text{B13}) \\ &= \frac{\hat{\Phi}(x, x) + \hat{\Phi}(y, y) + \hat{\Phi}(x, y) + \hat{\Phi}(y, x)}{2} \quad (\text{B14}) \\ &= \frac{\hat{\Phi}(x, x) + \hat{\Phi}(y, y) + 2\text{Re}(\hat{\Phi}(x, y))}{2}, \quad (\text{B15})\end{aligned}$$

where the last line uses the fact that $\hat{\Phi}(y, x) = \hat{\Phi}(x, y)^*$, as $\hat{\Phi}$ is Hermitian. Now, with probability at least $1 - \delta/4$, we estimate the overlap between $v(\Phi)$ and $|\psi(x, y)^+\rangle \langle \psi(x, y)^+|$ up to accuracy $\epsilon/(2\sqrt{2})$, using $\mathcal{O}(\log(1/\delta)/\epsilon^2)$ SWAP tests. This allows us to estimate $\text{Re}(\hat{\Phi}(x, y))$ up to accuracy $\epsilon/\sqrt{2}$ since we already estimated $\hat{\Phi}(x, x)$ and $\hat{\Phi}(y, y)$ from the previous step.

Similarly, to estimate $\text{Im}(\hat{\Phi}(x, y))$, instead of $|\psi(x, y)^+\rangle$, we use the state

$$|\psi(x, y)^-\rangle \triangleq \frac{|v(\sigma_x)\rangle - i|v(\sigma_y)\rangle}{\sqrt{2}}, \quad (\text{B16})$$

which satisfies

$$\begin{aligned}\langle \psi(x, y)^- | v(\Phi) | \psi(x, y)^- \rangle &= \frac{\hat{\Phi}(x, x) + \hat{\Phi}(y, y) + 2\text{Im}(\hat{\Phi}(x, y))}{2}. \quad (\text{B17})\end{aligned}$$

As in the case of the real component, we use $\mathcal{O}(\log(1/\delta)/\epsilon^2)$ SWAP tests to estimate this overlap,

allowing us to estimate the imaginary component up to error $\epsilon/\sqrt{2}$. Thus, we obtain an estimate of $\hat{\Phi}(x, y)$ up to error ϵ . From a union bound, we have a total failure probability of at most δ . ■

Another tool we utilize often is the classical shadow formalism [111], which allows for the efficient prediction of expectation values of many observables from copies of a particular state. We use classical shadows to estimate Pauli coefficients as a subroutine in some of our learning algorithms.

Lemma B4 (Classical shadows; Theorem 1 in Ref. [111]). Given copies of a quantum state $\rho \in \mathcal{B}(\mathcal{H})$ and a list of Hermitian observables $\{O_1, \dots, O_M\} \subset \mathcal{B}(\mathcal{H})$, satisfying $\|O_i\|_\infty \leq 1$ for all $i \in [M]$, there exists an algorithm that produce estimates $\{\hat{o}_i\}$ such that

$$|\hat{o}_i - \text{tr}(O_i \rho)| \leq \epsilon, \forall i \in [M] \quad (\text{B18})$$

with probability at least $1 - \delta$. The method employs randomized measurements on copies of ρ , followed by classical post-processing. Depending on the measurement primitive, this aforementioned algorithm has the following sample and time complexity.

- (1) To achieve Eq. (B18), it suffices to perform *random Clifford measurements* on

$$N = \mathcal{O}\left(\frac{\log(M/\delta)}{\epsilon^2} \max_{i \in [M]} \text{tr}(O_i^2)\right) \quad (\text{B19})$$

copies of the unknown quantum state, but the classical post-processing time may be inefficient in general. When the observables have efficient representations, the algorithm is computationally efficient. For example, when the observable is a stabilizer state, the computational time is

$$\mathcal{O}(n^2 MN). \quad (\text{B20})$$

- (2) When all observables O_i are k -local, to achieve Eq. (B18), it suffices to perform *random Pauli measurements* on

$$N = \mathcal{O}\left(\frac{4^k \log(M/\delta)}{\epsilon^2}\right) \quad (\text{B21})$$

copies of the unknown quantum state, and the classical post-processing can be performed in time

$$\mathcal{O}(kNM). \quad (\text{B22})$$

- (3) When all observables O_i are k -local Paulis, to achieve Eq. (B18), it suffices to perform *random*

Pauli measurements on

$$N = \mathcal{O}\left(\frac{3^k \log(M/\delta)}{\epsilon^2}\right) \quad (\text{B23})$$

copies of the state, and the classical post-processing can be performed in time

$$\mathcal{O}(kNM). \quad (\text{B24})$$

Finally, the last tool we need is used to ensure that the output of our agnostic learning algorithms is unitary or CPTP. This is because many of our learning algorithms rely on estimating Pauli coefficients; hence these estimates may not result in an actual unitary or quantum channel. Thus, in order to achieve *proper* learning, in which the output of our learning algorithm lies in the concept class $\mathcal{C}$ that we want to learn with respect to, we need to approximate our learned operator by the nearest unitary or quantum channel. We refer to this as “projection.” The following lemma shows how to achieve this for unitary projection via the approach of Lemma 8 in Ref. [23]. Note that Lemma 8 in Ref. [23] uses the same algorithm but shows that it is a unitary projection with respect to the operator norm. We show that this also gives a unitary projection with respect to the (normalized) Frobenius distance.

Lemma B5 (Unitary projection). There exists a map Proj_U that takes as input any operator $A \in \mathcal{B}(\mathcal{H})$ and outputs a unitary such that

$$d_F(A, \text{Proj}_U(A)) \leq d_F(A, B) \quad (\text{B25})$$

for any 2^n -dimensional unitary B and runs in time $\mathcal{O}(2^{3n})$. In other words, $\text{Proj}_U(A)$ is the closest unitary projection of A within normalized Frobenius distance:

$$\text{Proj}_U(A) = \arg \min_{B \in \mathcal{U}_{2^n}} d_F(A, B). \quad (\text{B26})$$

Proof. The projection Proj_U is defined as follows. Let $A = U \Sigma V^\dagger$ be the singular value decomposition of A , where Σ is a diagonal matrix consisting of non-negative real entries and $U, V \in \mathcal{U}_{2^n}$. Then, Proj_U is defined as

$$\text{Proj}_U(A) \triangleq UV^\dagger. \quad (\text{B27})$$

Intuitively, the projection computes the singular value decomposition of A and replaces all the singular values with 1. This technique was originally proposed in Ref. [23], where the authors showed that this is the closest unitary projection for any operator with respect to the operator norm. We instead show that this also gives a unitary projection for the (normalized) Frobenius distance. It

suffices to show that for all $B \in \mathcal{U}_d$

$$\|A - B\|_F^2 \geq \|A - UV^\dagger\|_F^2, \quad (\text{B28})$$

or equivalently, by the unitary invariance of the Frobenius norm,

$$\|\Sigma - U^\dagger BV\|_F^2 \geq \|\Sigma - I\|_F^2. \quad (\text{B29})$$

Denote $W \triangleq U^\dagger BV$. Now,

$$\|\Sigma - W\|_F^2 = \text{tr}((\Sigma - W^\dagger)(\Sigma - W)) \quad (\text{B30})$$

$$= \text{tr}(\Sigma^2 - W^\dagger \Sigma - \Sigma W + I) \quad (\text{B31})$$

where the first equality uses the fact that Σ is diagonal and non-negative, while the second line uses the fact that W is unitary.

Let $\{\Sigma_{ii}\}_{i \in [2^n]}$ denote the diagonal elements of Σ . Then, we have

$$\text{tr}(\Sigma W + W^\dagger \Sigma) = \sum_{i \in [2^n]} \langle i | (\Sigma W + W^\dagger \Sigma) | i \rangle \quad (\text{B32})$$

$$= \sum_{i \in [2^n]} \Sigma_{ii} \langle i | (W + W^\dagger) | i \rangle \quad (\text{B33})$$

$$= \sum_{i \in [2^n]} 2\Sigma_{ii} \cdot \text{Re}(\langle i | W | i \rangle) \quad (\text{B34})$$

$$\leq \sum_{i \in [2^n]} 2\Sigma_{ii} \quad (\text{B35})$$

$$= 2 \cdot \text{tr}(\Sigma). \quad (\text{B36})$$

Here, in the second line, we use that Σ is diagonal and so commutes with $W, W^\dagger$. In the inequality, we use that W is unitary. Finally, we get

$$\|\Sigma - W\|_F^2 \geq \text{tr}(\Sigma^2 - 2\Sigma + I) = \|\Sigma - I\|_F^2, \quad (\text{B37})$$

concluding the proof. The runtime can be seen as the runtime of any algorithm computing the SVD and multiplying $U, V^\dagger$. We refer to Lecture 31 of Ref. [128] for the analysis of algorithms based on the Golub-Kahan bidiagonalization procedure [129]. Namely, computing the SVD takes time $\mathcal{O}(m^3)$ for an $m \times m$ matrix. While such algorithms are only approximate, they demonstrate superlinear convergence, allowing one to efficiently obtain an output within machine precision and thus ignore this error in practice. ■

Showing such a result for CPTP projection is significantly more difficult, and we discuss it in detail in Appendix D 7.

APPENDIX C: AGNOSTIC UNITARY ESTIMATION (AUE)

In this section, we consider the task of agnostic unitary estimation (AUE). Recall that this is a special case of agnostic process tomography (APT) in which the unknown channel is assumed to be unitary and the concept class $\mathcal{C}$ consists entirely of unitaries. This section serves as a warm-up to Appendix D, as we will generalize many of the techniques presented here to APT later.

1. Bounded-gate circuits

In this section, we consider AUE when the concept class $\mathcal{C}$ consists of unitaries with a bounded gate complexity. In particular, we consider n -qubit unitaries generated by G two-qubit gates. This class of unitaries was studied in the realizable case in Ref. [27].

First, we prove a more general result about AUE for any class $\mathcal{C}$ of unitaries with finite size $|\mathcal{C}| < +\infty$. In the following proposition, we show that AUE with respect to the average-case distance metric d_{avg} and the class $\mathcal{C}$ can be accomplished using a sample complexity scaling as $\log |\mathcal{C}|$. Thus, if $\mathcal{C}$ is any class of unitaries generated by $\text{poly}(n)$ gates, $\log |\mathcal{C}| \leq \text{poly}(n)$ so that the sample complexity is efficient. Qualitatively, this says that efficient AUE is easy from a sample complexity perspective, but it can be difficult to achieve with respect to time complexity. This is similar in spirit to an analogous argument by Ref. [42] for agnostic state tomography. Their argument relies on shadow tomography [94], while ours is based on classical shadows and takes inspiration from a proof of Ref. [27].

Proposition 1 (Agnostic unitary estimation of finite-sized classes). There exists a learning algorithm for proper $(1, \epsilon, \delta)$ -agnostic unitary estimation with respect to error in average-case distance d_{avg} for any finite-sized class $\mathcal{C}$ of unitaries using

$$N = \mathcal{O}\left(\frac{\log(|\mathcal{C}|/\delta)}{\epsilon^4}\right) \quad (\text{C1})$$

queries to the unknown unitary U and running in time

$$\mathcal{O}\left(T \cdot \frac{|\mathcal{C}| \log(|\mathcal{C}|/\delta)}{\epsilon^4}\right), \quad (\text{C2})$$

where T is the time complexity of computing the inner product of two states given their exact description.

We note that the time complexity in the proposition depends on the complexity of computing the inner products of pairs of states. This is due to the use of Clifford classical shadows [111], which is not computationally efficient for arbitrary observables (see Lemma B4). The resulting complexity depends on the class of unitaries one aims

to learn and how efficiently their output states can be represented. However, since the time complexity already scales linearly with the size of the concept class, this algorithm would be computationally inefficient for most classes of interest, so we do not analyze the time complexity in detail.

Proof. The algorithm is the same as that of Sec. C.2(a) of Ref. [27]. Namely, we use Clifford classical shadows (Item 1 of Lemma B4) to estimate the distance $d_{\text{avg}}(U_i, U)$ between the unknown unitary U and all unitaries $U_i \in \mathcal{C}$. Then, our algorithm outputs the one with the minimal distance.

The analysis is the same as Sec. C.2(a) of Ref. [27], so we do not reproduce the full argument here. We start by preparing random input states from the ensemble

$$P = \text{Uniform}(|\psi\rangle = U|0\rangle, U \in \text{Cl}(n)) \quad (\text{C3})$$

of random stabilizer states, where $\text{Cl}(n)$ denotes the n -qubit Clifford group. Then, we use classical shadows (Lemma B4) to estimate distances with respect to this ensemble.

While Ref. [27] used random stabilizer product states as input, we use random stabilizer states instead. We note that changing the input distribution only changes the ensemble over which the distance is estimated, not the sample complexity for successful estimation. Now, by Eq. (C.8) of Ref. [27], using

$$N = 204 \frac{\log(2|\mathcal{C}|/\delta)}{\epsilon'^2} = \mathcal{O}\left(\frac{\log(|\mathcal{C}|/\delta)}{\epsilon'^2}\right) \quad (\text{C4})$$

queries to the unknown unitary, we see that we can construct unbiased estimators $\hat{o}_i$, such that $1 - \hat{o}_i$ estimates the square of the average distance with respect to P (Definition A11).

$$|d_P(U_i, U)^2 - (1 - \hat{o}_i)| \leq \epsilon', \quad \forall i \in [|\mathcal{C}|] \quad (\text{C5})$$

with probability at least $1 - \delta$ for some $\delta, \epsilon' > 0$. Then, we can select $i^* = \arg \min_i (1 - \hat{o}_i)$ and output U_{i^*} , i.e., the $U_i \in \mathcal{C}$ with the smallest estimated distance from U with respect to d_P . To relate this back to d_{avg} , recall that P forms a 3-design [130–132], implying that the average infidelity over P is the same as that over the Haar measure. Thus for any 2^n -dimensional unitaries U_1, U_2 ,

$$d_{\text{avg}}(U_1, U_2)^2 = \mathbf{E}_{|\psi\rangle \sim \mu_S} [d_{\text{tr}}(U_1|\psi\rangle, U_2|\psi\rangle)^2] \quad (\text{C6})$$

$$= \mathbf{E}_{|\psi\rangle \sim \mu_S} [1 - F(U_1|\psi\rangle, U_2|\psi\rangle)] \quad (\text{C7})$$

$$= \mathbf{E}_{|\psi\rangle \sim P} [1 - F(U_1|\psi\rangle, U_2|\psi\rangle)] \quad (\text{C8})$$

$$= \mathbf{E}_{|\psi\rangle \sim P} [d_{\text{tr}}(U_1|\psi\rangle, U_2|\psi\rangle)^2] \quad (\text{C9})$$

$$= d_P(U_1, U_2)^2, \quad (\text{C10})$$

where in the second and fourth lines we use Definition A10, and in the third line we use that P is a 3-design. We

can now bound the error of our hypothesis with respect to d_{avg} .

$$d_{\text{avg}}(U_{i^*}, U)^2 = d_P(U_{i^*}, U)^2 \quad (\text{C11})$$

$$\leq (1 - \hat{o}_{i^*}) + \epsilon' \quad (\text{C12})$$

$$= \min_i (1 - \hat{o}_i) + \epsilon' \quad (\text{C13})$$

$$\leq \min_i (d_P(U_i, U)^2 + \epsilon') + \epsilon' \quad (\text{C14})$$

$$\leq \min_i (d_P(U_i, U)^2) + 2\epsilon' \quad (\text{C15})$$

$$= \min_i (d_{\text{avg}}(U_i, U)^2) + 2\epsilon' \quad (\text{C16})$$

$$= \text{opt}(U, \mathcal{C})^2 + 2\epsilon'. \quad (\text{C17})$$

In the first line, we use Eq. (C10). In the second line, we use Eq. (C5). In the third line, we use the definition of i^* . In the fourth line, we use Eq. (C5) again. In the second to last line, we use Eq. (C10) again. Finally, in the last line, we use the definition of opt . Thus, by setting $\epsilon' = \epsilon^2/2$, we obtain

$$d_{\text{avg}}(U_{i^*}, U) \leq \text{opt}(U, \mathcal{C}) + \epsilon. \quad (\text{C18})$$

This gives us a 1-agnostic learner with the claimed sample complexity. The time complexity is due to the computation of the $\hat{o}_i$ for all elements of $\mathcal{C}$, each of which requires computing inner products between $\mathcal{O}(N)$ pairs of states. ■

While Proposition 1 only holds for concept classes of finite size, we now show that this can be easily extended to any class as long as it has bounded metric entropy (see Definition A12).

Corollary C1. Let $\epsilon, \epsilon', \delta > 0$. There exists a learning algorithm for proper $(1, \epsilon + \epsilon', \delta)$ -agnostic unitary estimation with respect to error in average-case distance d_{avg} for any class $\mathcal{C}$ with covering number $\mathcal{N}(\mathcal{C}, d_{\text{avg}}, \epsilon')$, using

$$N = \mathcal{O}\left(\frac{\log(\mathcal{N}(\mathcal{C}, d_{\text{avg}}, \epsilon')/\delta)}{\epsilon^4}\right) \quad (\text{C19})$$

queries to the unknown unitary U and running in time

$$\mathcal{O}(T \cdot \mathcal{N}(\mathcal{C}, d_{\text{avg}}, \epsilon') \cdot N), \quad (\text{C20})$$

where T is the time complexity of computing the inner product of two states given their exact description.

Proof. Because $\mathcal{C}$ has covering number $\mathcal{N}(\mathcal{C}, d_{\text{avg}}, \epsilon')$, there exists some ϵ' -covering net M with cardinality

$\mathcal{N}(\mathcal{C}, d_{\text{avg}}, \epsilon')$. Consider applying Proposition 1 to the ϵ' -cover M . This gives us a hypothesis $V \in M$ satisfying

$$d_{\text{avg}}(U, V) \leq \min_{W \in M} d_{\text{avg}}(U, W) + \epsilon. \quad (\text{C21})$$

Moreover, to find this hypothesis, we used

$$N = \mathcal{O}\left(\frac{\log(|M|/\delta)}{\epsilon^4}\right) = \mathcal{O}\left(\frac{\log(\mathcal{N}(\mathcal{C}, d_{\text{avg}}, \epsilon')/\delta)}{\epsilon^4}\right) \quad (\text{C22})$$

queries to the unknown unitary. The time complexity is also $\mathcal{O}(T \cdot \mathcal{N}(\mathcal{C}, d_{\text{avg}}, \epsilon') \cdot N)$ by Proposition 3.

Let W^* be the optimal hypothesis from $\mathcal{C}$, i.e.,

$$W^* \triangleq \arg \min_{W \in \mathcal{C}} d_{\text{avg}}(U, W) \quad \text{and} \quad \text{opt}(U, \mathcal{C}) = d_{\text{avg}}(U, W^*). \quad (\text{C23})$$

Then,

$$d_{\text{avg}}(U, V) \leq \min_{W \in M} d_{\text{avg}}(U, W) + \epsilon \quad (\text{C24})$$

$$\leq \min_{W \in M} (d_{\text{avg}}(U, W^*) + d_{\text{avg}}(W^*, W)) + \epsilon \quad (\text{C25})$$

$$\leq \text{opt}(U, \mathcal{C}) + \min_{W \in M} d_{\text{avg}}(W^*, W) + \epsilon \quad (\text{C26})$$

$$\leq \text{opt}(U, \mathcal{C}) + \epsilon' + \epsilon, \quad (\text{C27})$$

where the first inequality follows from the guarantees of Proposition 1, the second inequality follows from the triangle inequality for d_{avg} , the third inequality follows by the definition of W^* , and the fourth inequality uses that M is an ϵ' -covering net for $\mathcal{C}$. ■

Finally, we can obtain a sample complexity upper bound for AUE of unitaries with bounded gate complexity. This follows immediately from Lemma A9, which gives a bound on the metric entropy of unitaries with bounded gate complexity. However, by Corollary C1, the time complexity scales with the cardinality of the covering net, which is exponentially large (Lemma A9). This is not surprising, as even in the realizable case, there are lower bounds on the computational complexity of learning this class of unitaries [27] based on the hardness of RingLWE [133].

Corollary C2 (Agnostic unitary estimation of bounded-gate unitaries). There exists a (computationally inefficient) learning algorithm for proper $(1, \epsilon, \delta)$ -agnostic unitary estimation of the class of n -qubit unitaries of G two-qubit gates with respect to error in d_{avg} using

$$N = \tilde{\mathcal{O}}\left(\frac{G \log(n/\delta)}{\epsilon^4}\right) \quad (\text{C28})$$

queries.

2. AUE from agnostic state tomography

In this section, we prove a general result which shows that we can obtain an algorithm for improper AUE from existing agnostic state tomography algorithms with only a quadratic increase in sample and time complexity with respect to the error ϵ . This is a fairly general result that holds whenever the concept classes for our desired AUE algorithm and the agnostic state tomography are compatible, in the sense of Proposition 2. To properly state this proposition, we first give the definition of agnostic state tomography.

Definition C1 (Agnostic state tomography [42]). Let $0 < \epsilon, \delta < 1$. Let $\mathcal{D}$ be a class of quantum states. A learning algorithm $\mathcal{A}$ is an $(1, \epsilon, \delta)$ -agnostic learner with respect to fidelity and the concept class $\mathcal{D}$ if, given access to copies of an arbitrary quantum state ρ , $\mathcal{A}$ outputs a hypothesis state ρ' such that

$$F(\rho', \rho) \geq \max_{\sigma \in \mathcal{D}} F(\sigma, \rho) - \epsilon \quad (\text{C29})$$

with probability at least $1 - \delta$, where F denotes the fidelity. When the hypothesis ρ' is from the concept class $\mathcal{D}$, we say that $\mathcal{A}$ is *proper*. Otherwise, we say that it is *improper*.

Algorithms for agnostic state tomography with respect to fidelity are known for a variety of classes, e.g., stabilizer product states [42], stabilizer states, states with high stabilizer dimension, and discrete product states [43].

In the following proposition, we show how to convert any agnostic state tomography algorithm to an improper agnostic unitary estimation algorithm with effectively the same runtime, as long as the concept classes are compatible. This is satisfied for the corresponding unitary classes of all previously mentioned agnostic state tomography algorithms. Moreover, we note that since our algorithm is improper, the underlying agnostic state tomography algorithm does not have to be proper. However, we expect this to be helpful for obtaining proper learners, which we leave as an open question.

Proposition 2 (Improper AUE from agnostic state tomography). Suppose there exists a learning algorithm for (proper) $(1, \epsilon^2, \delta)$ -agnostic state tomography with respect to fidelity F for some class $\mathcal{D}$ of pure states. Let $\mathcal{C}$ be a concept class of unitaries. If $|v(U)\rangle \in \mathcal{D}$ for all $U \in \mathcal{C}$, then there exists a learning algorithm for improper $(1, \epsilon, \delta)$ -agnostic unitary estimation with respect to error in normalized Frobenius distance d_F for the class $\mathcal{C}$. Moreover, this requires the same sample and time complexity as the $(1, \epsilon^2, \delta)$ -agnostic state tomography algorithm.

We note that our result holds for agnostically learning the unitary channel associated to the unknown unitary.

Namely, if U is our unknown unitary, then our algorithm agnostically learns the channel $U(\cdot)U^\dagger$ with respect to d_F (the normalized Frobenius distance for superoperators from Definition A9). We note that this does not take into account any phase differences and so is slightly weaker than learning the unitary U itself with respect to d_F (the normalized Frobenius distance for operators from Definition A8). This is because our algorithm simply applies agnostic state tomography to the Choi state of the unknown process, inspired by Ref. [87], which does not contain this phase information.

Moreover, our analysis can be seen as extending that of Ref. [87] to the agnostic setting with respect to the normalized Frobenius distance d_F , rather than the (trivial) metric of entanglement infidelity defined as the infidelity between two Choi states:

$$\bar{F}(\Phi, \Psi) \triangleq 1 - F(v(\Phi), v(\Psi)), \quad (\text{C30})$$

where we adopt the (non-standard) notation of Ref. [6] and $v(\Phi)$ denotes the Choi state of Φ . We remark that we can also trivially obtain $(1, \epsilon, \delta)$ -agnostic unitary estimation with respect to the entanglement infidelity from $(1, \epsilon, \delta)$ -agnostic state tomography with respect to fidelity (notably, without the quadratic blowup with respect to ϵ) under the same compatibility condition. Entanglement infidelity has been considered widely in (realizable) unitary estimation [6, 87–93]. However, we focus on the normalized Frobenius distance d_F for consistency with our other results.

Proof. Let U denote the unknown unitary. The algorithm is simple: first, prepare the Choi state $|v(U)\rangle$ of U and then apply the $(1, \epsilon^2, \delta)$ -agnostic state tomography algorithm to $|v(U)\rangle$. The output of our algorithm is the density matrix of the pure state output by the state tomography algorithm.

The agnostic state tomography algorithm outputs some state $|\psi\rangle$ such that

$$F(v(U), |\psi\rangle\langle\psi|) \geq \max_{|\phi\rangle \in \mathcal{D}} F(v(U), |\phi\rangle\langle\phi|) - \epsilon^2, \quad (\text{C31})$$

where we use $v(U) \triangleq |v(U)\rangle\langle v(U)|$ to denote the density matrix corresponding to the pure Choi state. Note that this is consistent with the Choi state defined by the channel $\mathcal{E}(\rho) = U\rho U^\dagger$. We can equivalently write the above guarantee in terms of the infidelity

$$1 - |\langle v(U)|\psi\rangle|^2 \leq \min_{|\phi\rangle \in \mathcal{D}} (1 - |\langle v(U)|\phi\rangle|^2) + \epsilon^2. \quad (\text{C32})$$

We can convert this into a bound with respect to the normalized Frobenius distance d_F . Note that

$$\| |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_F^2 = 2(1 - |\langle\psi|\phi\rangle|^2) \quad (\text{C33})$$

for any pure states $|\psi\rangle, |\phi\rangle$. Moreover,

$$\frac{1}{2} \| |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_F^2 = 4^n d_F(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|)^2 \quad (\text{C34})$$

by definition of d_F (Definition A8), where $|\phi\rangle, |\psi\rangle$ are each on $2n$ qubits. Then, the above guarantee can be written as

$$4^n d_F(v(U), |\psi\rangle\langle\psi|)^2 \leq 4^n \min_{|\phi\rangle \in \mathcal{D}} d_F(v(U), |\phi\rangle\langle\phi|)^2 + \epsilon^2. \quad (\text{C35})$$

Taking the square root of both sides, we obtain

$$2^n d_F(v(U), |\psi\rangle\langle\psi|) \leq 2^n \min_{|\phi\rangle \in \mathcal{D}} d_F(v(U), |\phi\rangle\langle\phi|) + \epsilon. \quad (\text{C36})$$

Moreover, because $|v(V)\rangle \in \mathcal{D}$ for all $V \in \mathcal{C}$, we can upper bound this minimum by the minimum over all V in $\mathcal{C}$ instead. Thus, we have

$$2^n d_F(v(U), |\psi\rangle\langle\psi|) \leq 2^n \min_{V \in \mathcal{C}} d_F(v(U), v(V)) + \epsilon. \quad (\text{C37})$$

Now, we want to argue that agnostically learning the Choi state $v(U)$ of a unitary U with respect to $2^n d_F$ is equivalent to agnostically learning the unitary channel $U(\cdot)U^\dagger$ with respect to d_F . Note that the first d_F is the normalized Frobenius norm for operators (Definition A8) while the second d_F is the normalized Frobenius norm for superoperators (Definition A9). Then, notice

$$d_F(\mathcal{U}_1, \mathcal{U}_2) = \frac{1}{2^n \sqrt{2}} \|\mathcal{U}_1 - \mathcal{U}_2\|_F \quad (\text{C38})$$

$$= \frac{1}{2^n \sqrt{2}} \|J(U_1) - J(U_2)\|_F \quad (\text{C39})$$

$$= \frac{1}{\sqrt{2}} \|v(U_1) - v(U_2)\|_F \quad (\text{C40})$$

$$= 2^n d_F(v(U_1), v(U_2)). \quad (\text{C41})$$

Here, we use $\mathcal{U}$ to denote the unitary channel $U(\cdot)U^\dagger$. In the first and second lines, we use the definition of d_F for superoperators (Definition A9). In the third line, we use that the Choi representation is simply the unnormalized Choi state. Here, recall that $v(U) = |v(U)\rangle\langle v(U)|$, so the normalization factor is 2^n for this density matrix. Finally, in the last line, we use the definition of the normalized Frobenius distance for operators (Definition A8). Thus, this tells us that, indeed, agnostically learning the Choi state of a unitary with respect to $2^n d_F$ is the same as learning the corresponding unitary channel with respect to d_F . Hence, Eq. (C37) is exactly the guarantee needed for improper $(1, \epsilon, \delta)$ -AUE with respect to the class $\mathcal{C}$. ■

This gives a fairly general result for converting an agnostic state tomography algorithm with respect to $\mathcal{D}$ into an improper agnostic unitary estimation algorithm with respect to $\mathcal{C}$ when $|v(U)\rangle \in \mathcal{D}$ for all $U \in \mathcal{C}$. As remarked earlier, this holds for many classes of unitaries for which there already exist agnostic state tomography algorithms with respect to a suitable class. In particular, there exist agnostic state tomography algorithms for stabilizer states, states with high stabilizer dimension, and discrete product states [42,43]. Note that the corresponding classes of unitaries, i.e., Clifford circuits [134], Clifford + T circuits, and unitaries consisting of a tensor product of single-qubit gates, respectively, all satisfy the required property of Proposition 2. Thus, we can directly apply Proposition 2 with the state-of-the-art sample and time complexity bounds for agnostic state tomography to obtain improper AUE algorithms for the respective unitary classes. These are each direct corollaries of Corollary 6.3, Theorem 7.1, Corollary 8.3, and Corollary 9.3 in Ref. [43]. The results in Ref. [43] assume that $\text{opt} \geq \tau$ for some $\tau \geq \epsilon$, and their sample complexities depend on this parameter τ . When applying their results, we use that $\tau \geq \epsilon$ to eliminate the dependence on τ .

Corollary C3 (Improper AUE of Clifford circuits; via Corollary 6.3 in Ref. [43]). There exists a learning algorithm for improper $(1, \epsilon, \delta)$ -agnostic unitary estimation of the class of n -qubit Clifford unitaries with respect to error in d_F using

$$N = 2n \log(1/\delta) \left(\frac{1}{\epsilon} \right)^{\mathcal{O}(\log 1/\epsilon)} + \mathcal{O} \left(\frac{\log(1/\delta) + \log^2(1/\epsilon)}{\epsilon^4} \right) \quad (\text{C42})$$

queries to the unknown unitary and running in time

$$\mathcal{O} \left(n^2 \log(1/\delta) \left(n + \frac{\log(1/\delta)}{\epsilon^4} \right) \right) \cdot \left(\frac{1}{\epsilon} \right)^{\mathcal{O}(\log 1/\epsilon)}. \quad (\text{C43})$$

Corollary C4 (Improper AUE of Clifford + T circuits; via Theorem 7.1 in Ref. [43]). There exists a learning algorithm for improper $(1, \epsilon, \delta)$ -agnostic unitary estimation of the class of n -qubit unitaries consisting of Clifford gates and t T gates with respect to error in d_F using

$$N = 2n \log(1/\delta) \left(\frac{2^t}{\epsilon^2} \right)^{\mathcal{O}(\log 1/\epsilon)} \quad (\text{C44})$$

queries to the unknown unitary and running in time

$$\mathcal{O}(n^2 \log(1/\delta) N). \quad (\text{C45})$$

Corollary C5 (Improper AUE of product circuits; via Corollary 8.3 in Ref. [43]). Let $\mathcal{K}$ be a set of single-qubit gates for which $|\langle v(U_1) | v(U_2) \rangle|^2 \leq 1 - \mu$ for any distinct $U_1, U_2 \in \mathcal{K}$. There exists a learning algorithm for improper $(1, \epsilon, \delta)$ -agnostic unitary estimation of the class $\mathcal{C} = \mathcal{K}^{\otimes n}$ with respect to error in d_F using

$$N = \frac{\log^2(1/\delta) (2n|\mathcal{K}|)^{\mathcal{O}(\log(1/\epsilon)/\mu)}}{\epsilon^4} \quad (\text{C46})$$

queries to the unknown unitary and running in time $\mathcal{O}(N)$.

Corollary C6 (Improper AUE of Clifford product circuits; via Corollary 9.3 in Ref. [43]). There exists a learning algorithm for improper $(1, \epsilon, \delta)$ -agnostic unitary estimation of the class of tensor products of single-qubit Clifford gates with respect to error in d_F using

$$N = \log(2n) \log(1/\delta) \left(\frac{1}{\epsilon} \right)^{\mathcal{O}(\log 1/\epsilon)} + \mathcal{O} \left(\frac{\log^2(1/\epsilon) + \log(1/\delta)}{\epsilon^4} \right) \quad (\text{C47})$$

and running in time

$$\frac{4n^2 \log^2(1/\delta) (1/\epsilon)^{\mathcal{O}(\log 1/\epsilon)}}{\epsilon^4}. \quad (\text{C48})$$

Note that all of these algorithms are improper, while for applications, one would prefer proper learning algorithms. We leave it to future work to make these algorithms proper. For example, for proper AUE of Clifford circuits, one approach would be to efficiently find the closest stabilizer Choi state to a given stabilizer state.

3. Pauli strings

In this section, we consider agnostic learning with respect to the class of n -qubit Pauli strings $\mathcal{C}_{\mathcal{P}_n} \triangleq \{I, X, Y, Z\}^{\otimes n}$ and the distance metric d_{avg} . Here, we achieve a proper 1-agnostic unitary estimation algorithm, detailed in Algorithm 1, with an efficient sample complexity and runtime. In fact, the sample complexity and runtime for this class are both independent of the system size n . Moreover, our algorithm only requires access to copies of the Choi state of the unknown unitary U .

Theorem C1 (Agnostic unitary estimation of Pauli strings in d_{avg}). There exists a learning algorithm for proper $(1, \epsilon, \delta)$ -agnostic unitary estimation of the class of

ALGORITHM 1. 1-agnostic learning Pauli strings.

1: Prepare N copies of $|v(U)\rangle$ and measure them in the $\{|v(\sigma_x)\rangle\}_x$ basis to get strings $\{x_i\}_{i=1}^N$
 2: **return** the most frequent string

n -qubit Pauli strings $\mathcal{C}_{\mathcal{P}_n}$ with respect to error in d_{avg} using

$$N = \mathcal{O}\left(\frac{\log(1/\delta)}{\epsilon^4}\right) \quad (\text{C49})$$

queries to the unknown unitary and running in time

$$\mathcal{O}(N). \quad (\text{C50})$$

Proof. The algorithm is detailed in Algorithm 1. First, we prepare N copies of the Choi state $|v(U)\rangle$ (Definition A1). Then, we measure each one in the $\{|v(\sigma_x)\rangle\}_{x \in \mathcal{P}_n}$ basis and output σ_z associated with the most frequent outcome z . Using the fact that

$$|v(U)\rangle = \sum_{x \in \mathcal{P}_n} \hat{U}_x |v(\sigma_x)\rangle, \quad (\text{C51})$$

each measurement returns a string x with probability

$$p(x) \triangleq |\hat{U}_x|^2. \quad (\text{C52})$$

Denote the true most likely Pauli string by

$$x^* \triangleq \arg \max_{x \in \mathcal{P}_n} p(x). \quad (\text{C53})$$

We compute the optimal error for the class of Pauli strings. Recall from Lemma A8 that

$$d_{\text{avg}}(U, \mathcal{V})^2 = \frac{2^n}{2^n + 1} \left(1 - \left| \sum_{x \in \mathcal{P}_n} \hat{U}_x^* \hat{V}_x \right|^2 \right). \quad (\text{C54})$$

Then, we have

$$\text{opt}(U, \mathcal{C}_{\mathcal{P}_n})^2 = \min_{x \in \mathcal{P}_n} d_{\text{avg}}(U, \sigma_x)^2 \quad (\text{C55})$$

$$= \min_{x \in \mathcal{P}_n} \frac{2^n}{2^n + 1} (1 - |\hat{U}_x|^2) \quad (\text{C56})$$

$$= \frac{2^n}{2^n + 1} \left(1 - \max_{x \in \mathcal{P}_n} |\hat{U}_x|^2 \right). \quad (\text{C57})$$

Hence, we see that the hypothesis σ_{x^*} achieves the optimal error for the class $\mathcal{C}_{\mathcal{P}_n}$. Our algorithm aims to estimate the true most likely Pauli string x^* by the one most frequently occurring in our measurements. Let $\{\hat{p}(x)\}_x$ be the

empirical distribution obtained from our measurement outcomes, which estimates p . Then, denote the most frequent outcome out of our measurements by

$$z \triangleq \arg \max_{x \in \{x_i\}_{i \in [N]}} \hat{p}(x). \quad (\text{C58})$$

To bound the error, consider the empirical distribution $\{\hat{p}(x)\}_x$. Using Lemma B1, we can construct estimates $\hat{p}(x)$ of $|\hat{U}_x|^2$ such that with probability at least $1 - \delta$,

$$|\hat{p}(x) - |\hat{U}_x|^2| \leq \epsilon_1, \quad \forall x \in \mathcal{P}_n \quad (\text{C59})$$

for some $\epsilon_1 > 0$ using $\mathcal{O}(\log(1/\delta)/\epsilon_1^2)$ copies of the Choi state $|v(U)\rangle$. Now, since z is the most frequent string observed,

$$\hat{p}(z) \geq \hat{p}(x^*). \quad (\text{C60})$$

Conditioned on obtaining estimates $\hat{p}(x)$ that are ϵ_1 -close to $|\hat{U}_x|^2$ for all $x \in \mathcal{P}_n$, we have

$$|\hat{U}_z|^2 + \epsilon_1 \geq \hat{p}(z) \quad \text{and} \quad \hat{p}(x^*) \geq |\hat{U}_{x^*}|^2 - \epsilon_1. \quad (\text{C61})$$

Thus, we have

$$|\hat{U}_{x^*}|^2 - |\hat{U}_z|^2 \leq 2\epsilon_1. \quad (\text{C62})$$

Now, we can bound the error of the hypothesis as follows:

$$d_{\text{avg}}(U, \sigma_z)^2 = \frac{2^n}{2^n + 1} (1 - |\hat{U}_z|^2) \quad (\text{C63})$$

$$= \frac{2^n}{2^n + 1} (1 - |\hat{U}_{x^*}|^2 + |\hat{U}_{x^*}|^2 - |\hat{U}_z|^2) \quad (\text{C64})$$

$$\leq \text{opt}(U, \mathcal{C}_{\mathcal{P}_n})^2 + \frac{2^n}{2^n + 1} \cdot 2\epsilon_1, \quad (\text{C65})$$

where the first line is true by Lemma A8, and the last line follows by Eqs. (C57) and (C62). Thus, by setting $\epsilon_1 = \epsilon^2/2$, we obtain

$$d_{\text{avg}}(U, \sigma_z)^2 \leq \text{opt}(U, \mathcal{C}_{\mathcal{P}_n})^2 + \epsilon^2 \leq (\text{opt}(U, \mathcal{C}_{\mathcal{P}_n}) + \epsilon)^2, \quad (\text{C66})$$

and the desired bound can be obtained by taking the square root. Now, by our choice of $\epsilon_1 = \epsilon^2/2$, we use

$$N = \mathcal{O}\left(\frac{\log(1/\delta)}{\epsilon_1^2}\right) = \mathcal{O}\left(\frac{\log(1/\delta)}{\epsilon^4}\right) \quad (\text{C67})$$

ALGORITHM 2. 2-agnostic proper learning unitary k -juntas.

-
-
- 1: Prepare $N_1 = \mathcal{O}\left(\frac{k \log(n/\delta)}{\epsilon^4}\right)$ copies of $|v(U)\rangle$ and use them to compute estimates of weights of all subsets of $[n]$ with cardinality k .
 - 2: Select S to be the subset with the greatest estimated weight.
 - 3: Using $N_2 = \mathcal{O}\left(\frac{k 16^k \log(1/\delta)}{\epsilon^2}\right)$ queries to U , construct estimates α_x of the Pauli coefficients for $x \in \mathcal{P}_S$.
 - 4: Compute $V_1 = \sum_{x \in \mathcal{P}_S} \alpha_x \sigma_x$.
 - 5: Project V_1 onto a unitary, $V_2 = \text{Proj}_U(V_1)$.
 - 6: **return** V_2 .
-
-

samples of the Choi state or queries to the unitary U . The time complexity is that of identifying the most frequent string out of N strings, which can be done in time $\mathcal{O}(N)$. ■

4. Unitary k -juntas

Now, we consider agnostic learning with respect to the class of n -qubit unitary k -juntas $\mathcal{J}_{k,n}$ [28]. An n -qubit k -junta is a 2^n -dimensional unitary that acts nontrivially on k qubits. More formally, we have the following definition.

Definition C2 (k -junta). A unitary U on n qubits is a *quantum k -junta* if there exists a set $S \subseteq [n]$ with $|S| = k$ such that $U = V_S \otimes I_{S^c}$, for some k -qubit unitary V_S .

We consider learning with respect to d_F , and obtain an efficient proper 2-agnostic unitary estimation algorithm, detailed in Algorithm 2.

The algorithm is reminiscent of other junta learning algorithms [24, 135] in structure. Namely, we first identify the qubits S that the k -junta acts nontrivially on and then learn the unitary on the set S . In contrast to Ref. [24], we use the Pauli weight (Definition A5) rather than the influence to measure how non-trivially a unitary acts on a given qubit. In the realizable setting, both the weight and influence of a subset are equivalent. However, in the agnostic setting, these notions are quite distinct, as demonstrated by the following example.

Example C1. Consider the unitary $U = \frac{1}{\sqrt{3}}X \otimes Z \otimes I + \sqrt{\frac{2}{3}}Z \otimes Z \otimes Z$. One can show that the subset of size 2 with the most weight is formed by the first two qubits, $\{1, 2\}$, and has weight $w_{\{1,2\}}(U) = 1/3$. However, the most influential subset of size 2 is formed by the last two qubits, $\{2, 3\}$ with influence 1. Note that the corresponding weight $w_{\{2,3\}}(U) = 0$.

As we will show in our proof of Theorem C2, the optimal error depends on the subset with the highest weight, not the most influential subset. So, our algorithm starts by identifying this subset instead of the most influential one.

Moreover, rather than using a tomography algorithm in the second part, which would make proper learning more difficult, we directly estimate the Pauli coefficients corresponding to $\mathcal{P}_S$, where recall we use $\mathcal{P}_S$ to denote

$$\mathcal{P}_S = \{x \in \{0, 1, 2, 3\}^n : \text{supp}(x) \subseteq S\} \quad (\text{C68})$$

for a subset $S \subseteq [n]$. The resulting learned operator $V_1 = \sum_{x \in \mathcal{P}_S} \alpha_x \sigma_x$ may not be unitary, so in order to achieve *proper* agnostic unitary estimation, we need an additional projection which approximates V_1 by a unitary. For this projection, we use Lemma B5, which takes influence from Ref. [23]. This algorithm achieves the following guarantee.

Theorem C2 (Agnostic unitary estimation of k -juntas). Algorithm 2 performs proper $(2, \epsilon, \delta)$ -agnostic unitary estimation of the class of n -qubit unitary k -juntas $\mathcal{J}_{k,n}$ with respect to error in d_F using

$$N = \mathcal{O}\left(\frac{k \log(n/\delta)}{\epsilon^4} + \frac{k 16^k \log(1/\delta)}{\epsilon^2}\right) \quad (\text{C69})$$

queries to the unknown unitary U and running in time

$$\mathcal{O}\left(\frac{n^k k \log(n/\delta)}{\epsilon^4}\right). \quad (\text{C70})$$

As part of the proof of Theorem C2, we first show how to estimate the Pauli weights of all subsets of cardinality k , i.e., $S \subseteq [n]$ with $|S| = k$.

Lemma C1 (Estimating weights of a unitary on many subsets). Given query access to a unitary U and M subsets $\{S_j\}_{j \in [M]}$, $S_j \subseteq [n]$, there exists a method to produce estimates $\hat{w}_{S_j}$ of the Pauli weights $w_{S_j}(U)$ using $N = \mathcal{O}(\log(M/\delta)/\epsilon^2)$ queries to U such that

$$|\hat{w}_{S_j} - w_{S_j}(U)| \leq \epsilon \quad \forall j \in [M] \quad (\text{C71})$$

with probability at least $1 - \delta$. Moreover, the algorithm runs in computational time $\mathcal{O}(NM)$.

Proof. Using N queries to the unitary U , the algorithm first prepares N copies of the Choi state $|v(U)\rangle$ and

measures each copy in the basis $\{|v(\sigma_x)\rangle\}_x$. Suppose the outcomes are $\{x_i\}_{i \in [N]}$. Then, for each subset S , one can construct unbiased estimators

$$\hat{w}_S \triangleq \frac{1}{N} \sum_{i \in [N]} b_{S,i}, \quad b_{S,i} \triangleq \mathbf{1}[x_i \in \mathcal{P}_S]. \quad (\text{C72})$$

Here, $b_{S,i}$ indicates whether x_i is supported entirely in S , i.e., $\text{supp}(x_i) \subseteq S$. We claim that $\hat{w}_S$ has the desired properties given our choice of N . Clearly, since the measurements are independent, then $b_{S,i}$ are i.i.d. random variables so that

$$\mathbf{E}[\hat{w}_S] = \mathbf{E}[b_{S,i}] \quad (\text{C73})$$

$$= \Pr(x_i \in \mathcal{P}_S) \quad (\text{C74})$$

$$= \sum_{x \in \mathcal{P}_S} |\hat{U}_x|^2 \quad (\text{C75})$$

$$= w_S(U). \quad (\text{C76})$$

Here, the third line follows because measuring $|v(U)\rangle$ in the basis $\{|v(\sigma_x)\rangle\}_{x \in \mathcal{P}_n}$ gives outcome x with probability $|\hat{U}_x|^2$. Thus, using

$$N = \mathcal{O}\left(\frac{\log(M/\delta)}{\epsilon^2}\right) \quad (\text{C77})$$

copies of the Choi state, we can see via Hoeffding's inequality and a union bound that

$$|\hat{w}_{S_j} - w_{S_j}(U)| \leq \epsilon \quad \forall j \in [M] \quad (\text{C78})$$

with probability at least $1 - \delta$. For each subset, we use $\mathcal{O}(N)$ computational time to compute $\hat{w}_{S_j}$, giving an overall computational time of $\mathcal{O}(NM)$. ■

With this lemma, we can now prove Theorem C2.

Proof of Theorem C2. First, we start by lower-bounding $\text{opt}(U, \mathcal{J}_{n,k})$.

$$\text{opt}(U, \mathcal{J}_{n,k})^2 = \min_{V \in \mathcal{J}_{n,k}} d_F(U, V)^2 \quad (\text{C79})$$

$$= \min_{V \in \mathcal{J}_{n,k}} \frac{1}{2} \sum_{x \in \mathcal{P}_n} |\hat{U}_x - \hat{V}_x|^2 \quad (\text{C80})$$

$$= 1 - \max_{V \in \mathcal{J}_{n,k}} \sum_{x \in \mathcal{P}_n} \text{Re}(\hat{U}_x^* \hat{V}_x) \quad (\text{C81})$$

$$\geq 1 - \max_{V \in \mathcal{J}_{n,k}} \sum_{x \in \mathcal{P}_n} |\hat{U}_x| |\hat{V}_x| \quad (\text{C82})$$

$$\geq 1 - \max_{S \subseteq [n], |S|=k} \max_{\substack{V \in \mathcal{J}_{n,k} \\ w_S(V)=1}} \sum_{x \in \mathcal{P}_S} |\hat{U}_x| |\hat{V}_x| \quad (\text{C83})$$

$$\geq 1 - \max_{S \subseteq [n], |S|=k} \max_{\substack{V \in \mathcal{J}_{n,k} \\ w_S(V)=1}} \sqrt{w_S(U)} \sqrt{w_S(V)} \quad (\text{C84})$$

$$= 1 - \max_{S \subseteq [n], |S|=k} \sqrt{w_S(U)}, \quad (\text{C85})$$

where the second line follows from Lemma A5 and the second to last line follows from Cauchy-Schwarz. ■

Denote the heaviest cardinality- k set by

$$S^* \triangleq \arg \max_{S \subseteq [n], |S|=k} w_S(U). \quad (\text{C86})$$

Thus, we have

$$\text{opt}(U, \mathcal{J}_{n,k})^2 \geq 1 - \sqrt{w_{S^*}(U)}. \quad (\text{C87})$$

Hence, in order to learn the optimal k -junta, we first identify the subset of cardinality k with the highest Pauli weight. To do this, we estimate the weight of all $\binom{n}{k} = \mathcal{O}(n^k)$ subsets of $[n]$ of cardinality k up to accuracy ϵ_1 with failure probability at most $\delta/2$. In particular, we find estimates $\{\hat{w}_S\}_S$ such that

$$|\hat{w}_S - w_S(U)| \leq \epsilon_1, \quad (\text{C88})$$

with probability at least $1 - \delta/2$, for all subsets $S \subseteq [n]$ of cardinality $|S| = k$. By Lemma C1, this requires

$$N_1 = \mathcal{O}\left(\frac{k \log(n/\delta)}{\epsilon_1^2}\right) \quad (\text{C89})$$

queries to U and computational time $\mathcal{O}(n^k k \log(n/\delta)/\epsilon_1^2)$. We now focus on the subset S' with the heaviest estimated weight, i.e., the largest value of $\hat{w}_S$.

$$S' \triangleq \arg \max_{S \subseteq [n], |S|=k} \hat{w}_S. \quad (\text{C90})$$

In the next step of Algorithm 2, we estimate all 4^k Pauli coefficients for all Pauli strings supported on S' . We use Lemma B2 4^k times, to obtain estimates α_x such that with probability at least $1 - \delta/2$,

$$|\alpha_x - \hat{U}_x| \leq \epsilon_2, \quad \forall x \in \mathcal{P}_{S'}. \quad (\text{C91})$$

Using a union bound over failure probabilities of repeated applications of Lemma B2 for each Pauli coefficient, we

see that this requires

$$N_2 = \mathcal{O}\left(\frac{4^k \log(4^k/\delta)}{\epsilon_2^2}\right) \quad (\text{C92})$$

queries to U and runs in time $\mathcal{O}(4^k \log(4^k/\delta)/\epsilon_2^2)$. Now, we construct the hypothesis

$$V_1 \triangleq \sum_{x \in \mathcal{P}_{S'}} \alpha_x \sigma_x. \quad (\text{C93})$$

This simply aims to approximate U by approximating the Pauli coefficients of U on the subset S' with the heaviest estimated weight. As this hypothesis is not unitary in general, we apply Lemma B5 to project V_1 onto a unitary and obtain the final hypothesis

$$V_2 \triangleq \text{Proj}_U(V_1). \quad (\text{C94})$$

As V_1 acts trivially on all qubits outside of S' , the projection only needs to be applied to an operator of dimension $2^k \times 2^k$, needing time $\mathcal{O}(8^k)$. It remains to bound the error of this hypothesis, i.e., $d_F(U, V_2)$. Denote

$$\tilde{U}_{S'} \triangleq \sum_{x \in \mathcal{P}_{S'}} \hat{U}_x \sigma_x, \quad (\text{C95})$$

which is the operator obtained by truncating the Pauli coefficients of U not supported on S . Now, we have

$$d_F(U, V_2) \leq d_F(U, V_1) + d_F(V_1, V_2) \quad (\text{C96})$$

$$\leq 2d_F(U, V_1) \quad (\text{C97})$$

$$\leq 2d_F(U, \tilde{U}_{S'}) + 2d_F(\tilde{U}_{S'}, V_1), \quad (\text{C98})$$

where the first and last lines follow from the triangle inequality for d_F , and the second line follows from Lemma B5.

For our final bound, we use the following claims on these distances that will be proven later. The claims are conditioned on the estimations for the weights and the Pauli coefficients succeeding. A union bound tells us that the total failure probability is at most $\delta/2 + \delta/2 = \delta$.

Claim C1.

$$d_F(U, \tilde{U}_{S'}) \leq \text{opt}(U, \mathcal{J}_{n,k}) + \sqrt{\epsilon_1}. \quad (\text{C99})$$

Claim C2.

$$d_F(\tilde{U}_{S'}, V_1) \leq \frac{2^k \epsilon_2}{\sqrt{2}}. \quad (\text{C100})$$

Using these claims in the above inequality, we have

$$d_F(U, V_2) \leq 2 \cdot \text{opt}(U, \mathcal{J}_{n,k}) + 2\sqrt{\epsilon_1} + 2^k \sqrt{2} \epsilon_2. \quad (\text{C101})$$

Then, by setting $\epsilon_1 = \epsilon^2/16$, $\epsilon_2 = \epsilon/(2\sqrt{2} \cdot 2^k)$, we obtain

$$d_F(U, V_2) \leq 2 \cdot \text{opt}(U, \mathcal{J}_{n,k}) + \epsilon, \quad (\text{C102})$$

which is the desired result. The stated sample complexity can be obtained by plugging in the values of ϵ_1, ϵ_2 to the previous expressions for estimating weights and Pauli coefficients:

$$\begin{aligned} N &= N_1 + N_2 = \mathcal{O}\left(\frac{k \log(n/\delta)}{\epsilon_1^2} + \frac{4^k \log(4^k/\delta)}{\epsilon_2^2}\right) \\ &= \mathcal{O}\left(\frac{k \log(n/\delta)}{\epsilon^4} + \frac{k 16^k \log(1/\delta)}{\epsilon^2}\right). \end{aligned} \quad (\text{C103})$$

The runtime is dominated by the procedures for finding the heaviest subset and for estimating the 4^k Pauli coefficients:

$$\begin{aligned} &\mathcal{O}\left(\frac{n^k k \log(n/\delta)}{\epsilon_1^2} + \frac{4^k \log(4^k/\delta)}{\epsilon_2^2}\right) \\ &= \mathcal{O}\left(\frac{n^k k \log(n/\delta)}{\epsilon^4} + \frac{k 16^k \log(1/\delta)}{\epsilon^2}\right). \end{aligned} \quad (\text{C104})$$

As the first time dominates the time complexity, we obtain the stated complexity. It remains to prove the two claims above.

Proof of Claim 1. Because the algorithm selects S' , it is the subset with the heaviest estimated weight, i.e.,

$$\hat{w}_{S'} \geq \hat{w}_{S^*}. \quad (\text{C105})$$

Further, by Eq. (C88),

$$w_{S'}(U) + \epsilon_1 \geq \hat{w}_{S'} \quad \text{and} \quad \hat{w}_{S^*} \geq w_{S^*}(U) - \epsilon_1. \quad (\text{C106})$$

Thus,

$$w_{S^*}(U) - w_{S'}(U) \leq 2\epsilon_1 \quad (\text{C107})$$

and $w_{S^*} \geq w_{S'}$. For ease of notation, denote $\mathcal{W} \triangleq \tilde{U}_{S'}$. Then, we have

$$d_F(U, \tilde{U}_{S'})^2 = \frac{1}{2} \sum_{x \in \mathcal{P}_n} |\hat{U}_x - \hat{W}_x|^2 \quad (\text{C108})$$

$$= \frac{1}{2} \left(\sum_{x \in \mathcal{P}_n \setminus \mathcal{P}_{S'}} |\hat{U}_x - 0|^2 + \sum_{x \in \mathcal{P}_{S'}} |\hat{U}_x - \hat{U}_x|^2 \right) \quad (\text{C109})$$

$$= \frac{1 - w_{S'}(U)}{2} \quad (\text{C110})$$

$$= \frac{1 - w_{S^*}(U)}{2} + \frac{w_{S^*}(U) - w_{S'}(U)}{2} \quad (\text{C111})$$

$$\leq \frac{(1 - \sqrt{w_{S^*}(U)})(1 + \sqrt{w_{S^*}(U)})}{2} + \epsilon_1 \quad (\text{C112})$$

$$\leq 1 - \sqrt{w_{S^*}(U)} + \epsilon_1 \quad (\text{C113})$$

$$\leq \text{opt}(U, \mathcal{J}_{n,k})^2 + \epsilon_1 \quad (\text{C114})$$

$$\leq (\text{opt}(U, \mathcal{J}_{n,k}) + \sqrt{\epsilon_1})^2. \quad (\text{C115})$$

In the first line, we use Lemma A5. In the second line, we use that W is only supported on S' , for which its Pauli coefficients are $\hat{U}_x$. In the third line, we use the definition of the Pauli weight (Definition A5). In the fifth line, we use Eq. (C107). In the seventh line, we use Eq. (C87). ■

Proof of Claim 2. As in the proof of Claim 1, for ease of notation, we denote $W \triangleq \tilde{U}_{S'}$.

$$d_F(\tilde{U}_{S'}, V_1)^2 = \frac{1}{2} \sum_{x \in \mathcal{P}_n} |\hat{W}_x - \hat{V}_{1,x}|^2 \quad (\text{C116})$$

$$= \frac{1}{2} \sum_{x \in \mathcal{P}_{S'}} |\hat{U}_x - \hat{\alpha}_x|^2 + \frac{1}{2} \sum_{x \in \mathcal{P}_n \setminus \mathcal{P}_{S'}} |0 - 0|^2 \quad (\text{C117})$$

$$\leq \frac{4^k \epsilon_2^2}{2}. \quad (\text{C118})$$

In the second line, we use that both W and V_1 are only supported on S' , and the Pauli coefficients of W on S' are $\hat{U}_x$. In the last line, we use Eq. (C91) and that $|\mathcal{P}_{S'}| = 4^k$ since $|S'| = k$. ■

Remark C1 (Improper 1-agnostic learner). In the setting of improper agnostic learning, we do not need our final hypothesis to be unitary and can simply output V_1 . This would give an improper $(1, \epsilon, \delta)$ -agnostic learning algorithm with the same sample and time complexity.

5. Low-degree unitaries

Now, we will present two algorithms for agnostic unitary estimation of the class of low-degree unitaries. A low-degree unitary is one whose high-order Pauli coefficients are 0. More formally, we have the following definition.

Definition C3 (Degree- d unitary). A unitary U has degree d if

$$w^{>d}(U) = 0, \quad (\text{C119})$$

where $w^{>d}$ is defined in Definition A4. In other words, the degree of a unitary U is the minimum integer d such that if $|x| > d$, then $\hat{U}_x = 0$.

We use $\mathcal{L}_{n,d}$ to denote the class of n -qubit unitaries of degree d . We note that there are $\binom{n}{d} 3^d \leq n^d 3^d$ Pauli strings of degree exactly d , and thus $\mathcal{O}(n^d 3^d)$ strings of degree at most d .

Our first algorithm given in Algorithm 3 is inspired by the algorithm of Ref. [22]. It involves estimating every low-degree Pauli coefficient, resulting in a query complexity of $\tilde{\mathcal{O}}(n^{2d})$. Our second algorithm given in Algorithm 4 is the low-degree learning algorithm from Ref. [30], which uses the non-commutative Bohnenblust-Hille (BH) inequality of Ref. [116] to obtain a sample complexity independent of n . However, in the agnostic setting, Algorithm 4 does not achieve a sample complexity independent of n . From the guarantees we show in Theorem C4, Algorithm 4 is only efficient when the optimal error is at most inverse-exponential in the system size. In particular, when the unknown unitary is a member of the concept class, we recover the complexity of the realizable setting. In most other regimes for opt , it is better to use Algorithm 3. We discuss Algorithm 4 to illustrate that straightforwardly modifying algorithms from the realizable setting does not always give good agnostic learning algorithms.

Note that unlike in the case of unitary juntas, our method for unitary projection is computationally inefficient when acting on arbitrary low-degree operators, as such operators still have an exponential dimension. As a result, both our algorithms are *improper* learners. We provide our algorithms and their corresponding guarantees below.

Theorem C3 (AUE of low-degree unitaries by learning all low-degree coefficients). Algorithm 3 performs improper $(1, \epsilon, \delta)$ -agnostic unitary estimation of the class of n -qubit unitaries of degree d $\mathcal{L}_{n,d}$ with respect to error in d_F using

$$N = \mathcal{O}\left(\frac{d \cdot 9^d n^{2d} \log(n/\delta)}{\epsilon^2}\right) \quad (\text{C120})$$

queries to the unknown unitary U and running in time

$$\mathcal{O}(N). \quad (\text{C121})$$

ALGORITHM 3. 1-agnostic improper learning low-degree unitaries by learning all coefficients.

- 1: Using N queries to U , construct estimates α_x of the Pauli coefficients for all $x \in \mathcal{P}_n$ s.t. $|x| \leq d$.
 - 2: **return** $V_1 = \sum_{x \in \mathcal{P}_n} \alpha_x \sigma_x$.
-

ALGORITHM 4. 1-agnostic improper learning low-degree unitaries using BH inequality.

-
- 1: Set c as in Equation (C.169).
 - 2: Prepare $N_1 = \mathcal{O}(\log(1/\delta)/c^4)$ copies of $|v(U)\rangle$ and use them to construct estimates $|\hat{U}'_x|^2$ of $|\hat{U}_x|^2$.
 - 3: Construct $\chi_c = \{x \in \mathcal{P}_n : |\hat{U}'_x| \geq c, |x| \leq d\}$.
 - 4: Using $N_2 = \mathcal{O}\left(\frac{\log(1/c^2\delta)}{c^4\epsilon^2}\right)$ queries to U , construct estimates α_x of the Pauli coefficients for all $x \in \chi_c$.
 - 5: **return** $V_1 = \sum_{x \in \chi_c} \alpha_x \sigma_x$.
-

Theorem C4 (AUE of low-degree unitaries using BH inequality). Algorithm 4 performs improper $(1, \epsilon, \delta)$ -agnostic unitary estimation of the class of n -qubit unitaries of degree d $\mathcal{L}_{n,d}$ with respect to error in d_F using

$$N = \exp\left(\tilde{\mathcal{O}}\left(d^2 + d \log(1/\epsilon)\right) + \min\left[d^2 \log n, d \log(1 + 2^{(n+1)/2} \cdot \text{opt})\right]\right) \cdot \log(1/\delta) \quad (\text{C122})$$

queries to the unknown unitary U and running in time

$$\mathcal{O}(N). \quad (\text{C123})$$

In order to prove Theorems C3 and C4, we will use the following lemma on the optimal error of the class of low-degree unitaries.

Lemma C2 (Optimal error for low-degree unitaries). The optimal error in d_F of the class of low-degree unitaries $\mathcal{L}_{n,d}$ is lower-bounded by

$$\text{opt}(U, \mathcal{L}_{n,d}) \geq \sqrt{1 - \sqrt{w^{\leq d}(U)}}. \quad (\text{C124})$$

Proof. From the definition of opt , we have

$$\text{opt}(U, \mathcal{L}_{n,d})^2 = \min_{V \in \mathcal{L}_{n,d}} d_F^2(U, V) \quad (\text{C125})$$

$$= \frac{1}{2} \min_{V \in \mathcal{L}_{n,d}} \sum_{x \in \mathcal{P}_n} |\hat{U}_x - \hat{V}_x|^2 \quad (\text{C126})$$

$$= 1 - \max_{V \in \mathcal{L}_{n,d}} \sum_{x \in \mathcal{P}_n} \text{Re}(\hat{U}_x \hat{V}_x^*) \quad (\text{C127})$$

$$= 1 - \max_{V \in \mathcal{L}_{n,d}} \sum_{|x| \leq d} \text{Re}(\hat{U}_x \hat{V}_x^*) \quad (\text{C128})$$

$$\geq 1 - \max_{V \in \mathcal{L}_{n,d}} \sum_{|x| \leq d} |\hat{U}_x \hat{V}_x^*| \quad (\text{C129})$$

$$\geq 1 - \max_{V \in \mathcal{L}_{n,d}} \sqrt{\sum_{|x| \leq d} |\hat{U}_x|^2} \sqrt{\sum_{|x| \leq d} |\hat{V}_x|^2} \quad (\text{C130})$$

$$= 1 - \max_{V \in \mathcal{L}_{n,d}} \sqrt{w^{\leq d}(U)} \sqrt{w^{\leq d}(V)} \quad (\text{C131})$$

$$= 1 - \sqrt{w^{\leq d}(U)}, \quad (\text{C132})$$

where the second line follows from Lemma A5, the fourth line uses that V has degree at most d , the first inequality uses the fact that the real component of a complex number is smaller than its absolute value, the second inequality uses Cauchy-Schwarz, and the final equality follows from the fact that V is a degree- d unitary. This concludes the proof. ■

With this, we can prove Theorem C3.

Proof of Theorem C3. We start by estimating all $\mathcal{O}(n^d 3^d)$ Pauli coefficients of degree d . We use Lemma B2 $\mathcal{O}(n^d 3^d)$ times, to obtain estimates α_x such that with probability at least $1 - \delta$,

$$|\alpha_x - \hat{U}_x| \leq \epsilon_1, \quad \forall x \in \mathcal{P}_n \text{ s.t. } |x| \leq d. \quad (\text{C133})$$

Using a union bound over failure probabilities of repeated applications of Lemma B2 for each Pauli coefficient, we see that this requires

$$N = \mathcal{O}\left(\frac{d \cdot 3^d n^d \log(n/\delta)}{\epsilon_1^2}\right) \quad (\text{C134})$$

queries to U and runs in time $\mathcal{O}(N)$. Now, we construct our hypothesis

$$V \triangleq \sum_{|x| \leq d} \alpha_x \sigma_x. \quad (\text{C135})$$

This simply aims to approximate U by estimating all low-degree coefficients. It remains to bound the error of this

hypothesis, i.e., $d_F(U, V)$. Denote

$$\tilde{U}_d = \sum_{|x| \leq d} \hat{U}_x \sigma_x, \quad (\text{C136})$$

which is the low-degree Pauli truncation of U . Now, using the triangle inequality for d_F , we have

$$d_F(U, V) \leq d_F(U, \tilde{U}_d) + d_F(\tilde{U}_d, V). \quad (\text{C137})$$

For ease of notation, let $\hat{W}_x$ denote the Pauli coefficients of $\tilde{U}_d$. We upper bound $d_F(U, \tilde{U}_d)$ as follows:

$$d_F(U, \tilde{U}_d)^2 = \frac{1}{2} \sum_{x \in \mathcal{P}_n} |\hat{U}_x - \hat{W}_x|^2 \quad (\text{C138})$$

$$= \frac{1}{2} \left(\sum_{|x| \leq d} |\hat{U}_x - \hat{U}_x|^2 + \sum_{|x| > d} |\hat{U}_x - 0|^2 \right) \quad (\text{C139})$$

$$= \frac{w^{>d}(U)}{2} \quad (\text{C140})$$

$$= \frac{1 - w^{\leq d}(U)}{2} \quad (\text{C141})$$

$$= \frac{(1 - \sqrt{w^{\leq d}(U)})(1 + \sqrt{w^{\leq d}(U)})}{2} \quad (\text{C142})$$

$$\leq 1 - \sqrt{w^{\leq d}(U)} \quad (\text{C143})$$

$$\leq \text{opt}(U, \mathcal{L}_{n,d})^2, \quad (\text{C144})$$

where in the last line, we use Lemma C2. Finally, conditioned on the successful estimation of α_x , we can also bound $d_F(\tilde{U}_d, V)$.

$$d_F(\tilde{U}_d, V)^2 = \frac{1}{2} \sum_{x \in \mathcal{P}_n} |\hat{W}_x - \hat{V}_x|^2 \quad (\text{C145})$$

$$= \frac{1}{2} \sum_{|x| \leq d} |\hat{U}_x - \alpha_x|^2 \quad (\text{C146})$$

$$\leq \mathcal{O}(n^d 3^d \epsilon_1^2). \quad (\text{C147})$$

The inequality follows because there are at most $\mathcal{O}(n^d 3^d)$ Pauli strings x such that $|x| \leq d$. Thus, with probability at least $1 - \delta$, we have

$$d_F(U, V) \leq \text{opt}(U, \mathcal{L}_{n,d}) + \mathcal{O}(\sqrt{n^d 3^d \epsilon_1}). \quad (\text{C148})$$

To obtain the 1-agnostic learning guarantee, we choose $\epsilon_1 = \mathcal{O}(\epsilon / \sqrt{n^d 3^d})$. This gives us a query complexity

$$N = \mathcal{O}\left(\frac{d \cdot 9^d n^{2d} \log(n/\delta)}{\epsilon^2}\right) \quad (\text{C149})$$

and time complexity

$$\mathcal{O}(N) \quad (\text{C150})$$

as stated. $\blacksquare$

Now, we will use the non-commutative BH inequality of Ref. [116] to prove Theorem C4.

Lemma C3 (Non-commutative BH inequality; Theorem 1.2 in Ref. [116]). For any $d \geq 1$, there exists $C > 0$, such that for all $n \geq 1$ and all $A \in \mathcal{B}(\mathcal{H})$ of degree at most d , we have

$$\left(\sum_{|x| \leq d} |\hat{A}_x|^{\frac{2d}{d+1}} \right)^{\frac{d+1}{2d}} \leq C^d \|A\|_\infty \quad (\text{C151})$$

for some constant C .

Now, we can prove Theorem C4.

Proof of Theorem C4. Here, we consider the learning algorithm of Ref. [30], and our analysis is similar to theirs. The algorithm is sketched in Algorithm 4. $\blacksquare$

Let c be a constant which we will specify later. In Step 1 of Algorithm 4, we prepare N_1 copies of the Choi state $|v(U)\rangle$ and measure each of them in the basis $\{|v(\sigma_x)\rangle\}_x$, where again we specify N_1 later. This allows us to construct empirical estimates $\{|\hat{U}'_x|^2\}_x$ by sampling from the distribution $\{|\hat{U}_x|^2\}_x$. We wish to construct estimates such that, with probability at least $1 - \delta/2$,

$$||\hat{U}'_x|^2 - |\hat{U}_x|^2| \leq c^2, \quad \forall x \in \mathcal{P}_n. \quad (\text{C152})$$

Using Lemma B1, we see that this can be done using

$$N_1 = \mathcal{O}\left(\frac{\log(1/\delta)}{c^4}\right) \quad (\text{C153})$$

copies of the Choi state and $\mathcal{O}(N_1)$ time. In what follows, we condition on successful estimation. Then, we have

$$||\hat{U}'_x| - |\hat{U}_x|| \leq c, \quad \forall x \in \mathcal{P}_n. \quad (\text{C154})$$

In Step 2 of Algorithm 4, we denote the set of low-degree strings with high empirical weight by χ_c , i.e.,

$$\chi_c \triangleq \{x \in \mathcal{P}_n : |\hat{U}'_x| \geq c, |x| \leq d\}. \quad (\text{C155})$$

From Parseval's identity (Lemma A2), we see that

$$|\chi_c| \leq 1/c^2. \quad (\text{C156})$$

Note that for any string x with $|x| \leq d$, if $x \notin \chi_c$,

$$|\hat{U}_x| \leq |\hat{U}'_x| + ||\hat{U}_x| - |\hat{U}'_x|| \leq 2c, \quad (\text{C157})$$

where we use Eq. (C154) and the definition of χ_c . In Step 3 of Algorithm 4, we want to construct estimates α_x for

the Pauli coefficients of all strings $x \in \chi_c$, such that with probability at least $1 - \delta/2$,

$$|\alpha_x - \hat{U}_x| \leq c\epsilon, \quad \forall x \in \chi_c. \quad (\text{C158})$$

This requires $|\chi_c|$ applications of Lemma B2, using

$$N_2 = \mathcal{O}\left(\frac{|\chi_c| \log(|\chi_c|/\delta)}{c^2 \epsilon^2}\right) = \mathcal{O}\left(\frac{\log(1/c^2 \delta)}{c^4 \epsilon^2}\right) \quad (\text{C159})$$

queries to the unknown unitary and $\mathcal{O}(N_2)$ time. Our hypothesis will then be

$$V \triangleq \sum_{x \in \chi_c} \alpha_x \sigma_x. \quad (\text{C160})$$

It remains to bound the error of this hypothesis, i.e., $d_F(U, V)$. Using the triangle inequality for d_F , we have

$$d_F(U, V) \leq d_F(U, \tilde{U}_d) + d_F(\tilde{U}_d, V) \quad (\text{C161})$$

where

$$\tilde{U}_d = \sum_{|x| \leq d} \hat{U}_x \sigma_x, \quad (\text{C162})$$

as in the proof of Theorem C3. As we showed in the proof of Theorem C3,

$$d_F(U, \tilde{U}_d) \leq \text{opt}(U, \mathcal{L}_{n,d}). \quad (\text{C163})$$

It remains to bound $d_F(\tilde{U}_d, V)$, for which we follow the arguments of Ref. [30]. Denote by $\hat{W}_x$ the Pauli coefficients of $\tilde{U}_d$. Then, we have

$$d_F(\tilde{U}_d, V)^2 = \frac{1}{2} \sum_{x \in \mathcal{P}_n} |\hat{W}_x - \hat{V}_x|^2 \quad (\text{C164})$$

$$= \frac{\sum_{x \in \chi_c} |\hat{U}_x - \hat{V}_x|^2}{2} + \frac{\sum_{x \notin \chi_c, |x| \leq d} |\hat{U}_x|^2}{2} + \frac{\sum_{|x| > d} 0}{2} \quad (\text{C165})$$

$$\leq \frac{|\chi_c| c^2 \epsilon^2}{2} + \frac{\sum_{x \notin \chi_c, |x| \leq d} (2c)^{\frac{2}{d+1}} |\hat{U}_x|^{\frac{2d}{d+1}}}{2} \quad (\text{C166})$$

$$\leq \frac{\epsilon^2}{2} + \frac{(2c)^{\frac{2}{d+1}} (C^d \|\tilde{U}_d\|_\infty)^{\frac{2d}{d+1}}}{2}. \quad (\text{C167})$$

The first line follows by Lemma A5. The third line follows from Eqs. (C157) and (C158). The last line follows from Lemma C3.

In the realizable setting, $\tilde{U}_d$ is a unitary and thus has operator norm 1. However, this need not be the case in the agnostic setting, so we need to upper bound $\|\tilde{U}_d\|_\infty$. We state these upper bounds in the following claim, which we will prove later.

Claim C3.

$$\|\tilde{U}_d\|_\infty \leq \min[n^{d/2} 3^{d/2}, 1 + 2^{(n+1)/2} \cdot \text{opt}(U, \mathcal{L}_{n,d})]. \quad (\text{C168})$$

Let u be the tightest upper bound from Claim 3. Hence, by choosing

$$c = \frac{\epsilon^{d+1}}{2C^{d^2} u^d}, \quad (\text{C169})$$

we obtain

$$d_F(\tilde{U}_d, V) \leq \epsilon, \quad (\text{C170})$$

giving us

$$d_F(U, V) \leq \text{opt}(U, \mathcal{L}_{n,k}) + \epsilon \quad (\text{C171})$$

as desired. The query complexity is dominated by N_2 , giving us

$$N = \exp\left(\tilde{\mathcal{O}}\left(d^2 + d \log(1/\epsilon)\right) + \min\left[d^2 \log n, d \log(1 + 2^{(n+1)/2} \cdot \text{opt})\right]\right) \cdot \log(1/\delta). \quad (\text{C172})$$

The running time is given by

$$\mathcal{O}(N_1 + N_2) = \mathcal{O}(N). \quad (\text{C173})$$

It remains to prove Claim 3.

Proof of Claim 3. We start by proving the first bound.

$$\|\tilde{U}_d\|_\infty = \left\| \sum_{|x| \leq d} \hat{U}_x \sigma_x \right\|_\infty \quad (\text{C174})$$

$$\leq \sum_{|x| \leq d} \|\hat{U}_x \sigma_x\|_\infty \quad (\text{C175})$$

$$= \sum_{|x| \leq d} |\hat{U}_x| \quad (\text{C176})$$

$$\leq \sqrt{n^d 3^d} \sqrt{\sum_{|x| \leq d} |\hat{U}_x|^2} \quad (\text{C177})$$

$$= n^{d/2} 3^{d/2} \sqrt{w^{\leq d}(U)} \quad (\text{C178})$$

$$\leq n^{d/2} 3^{d/2}, \quad (\text{C179})$$

where in the second line we use the triangle inequality of $\|\cdot\|_\infty$. In the third line we use the fact that $\|\sigma_x\|_\infty = 1$ for

all $x \in \mathcal{P}_n$. In the fourth line we use the Cauchy-Schwarz inequality. Finally, in the last line we use the fact that the weight of a unitary up to any degree is at most one.

We will now prove the second bound which is dependent on opt . Let W be the closest low-degree unitary to U , i.e.,

$$w^{\leq d}(W) = 1 \quad \text{and} \quad d_F(U, W) = \text{opt}(U, \mathcal{L}_{n,d}). \quad (\text{C180})$$

We start by lower bounding the normalized Frobenius distance between W and U to obtain a bound on the normalized Frobenius distance between W and $\tilde{U}_d$.

$$\text{opt}(U, \mathcal{L}_{n,d})^2 = d_F(U, W)^2 \quad (\text{C181})$$

$$= \frac{1}{2} \sum_{|x| \leq d} |\hat{U}_x - \hat{W}_x|^2 + \frac{1}{2} \sum_{|x| > d} |\hat{U}_x|^2 \quad (\text{C182})$$

$$= d_F(\tilde{U}_d, W)^2 + w^{>d}(U) \quad (\text{C183})$$

$$\geq d_F(\tilde{U}_d, W)^2 \quad (\text{C184})$$

$$= \frac{\|\tilde{U}_d - W\|_F^2}{2^{n+1}}. \quad (\text{C185})$$

In the second and third lines, we have used Lemma A5 and the fact that W has degree at most d . In the fourth line, we use the fact that the weight above a degree is non-negative, and in the last line we use the definition of d_F . Now, we can directly bound the operator norm of $\tilde{U}_d$.

$$\|\tilde{U}_d\|_\infty \leq \|\tilde{U}_d - W\|_\infty + \|W\|_\infty \quad (\text{C186})$$

$$\leq \|\tilde{U}_d - W\|_F + 1 \quad (\text{C187})$$

$$\leq \sqrt{2^{n+1}} \cdot \text{opt}(U, \mathcal{L}_{n,d}) + 1. \quad (\text{C188})$$

In the first line, we use the triangle inequality of $\|\cdot\|_\infty$. In the second line, we use the fact that $\|\cdot\|_\infty \leq \|\cdot\|_F$ and that W is a unitary. Finally, in the last line, we use Eq. (C185). This concludes the proof. ■

APPENDIX D: AGNOSTIC PROCESS TOMOGRAPHY (APT)

In this section, we generalize our agnostic unitary estimation algorithms from Appendix C to obtain algorithms for agnostic process tomography. Recall that in agnostic process tomography, we have access to an unknown quantum channel rather than a unitary. This often makes the analysis more complicated.

1. Bounded-gate circuits

In Appendix C 1, we showed that agnostic unitary estimation can be easily performed efficiently with respect to the number of samples. However, it can be inefficient with respect to the computational complexity. This motivates

why we should care more about the runtime of algorithms in agnostic unitary estimation. In this section, we aim to show a similar result for agnostic process tomography.

As in Appendix C 1, we consider APT with respect to any class $\mathcal{C}$ of unitaries with finite size $|\mathcal{C}| < +\infty$. Here, the error bound we prove is not 1-agnostic but depends on the fourth root of the optimal error. Nevertheless, this qualitatively gives the same result that APT is easy with respect to sample complexity but can be more difficult with respect to runtime.

Proposition 1 (Agnostic process tomography of finite-sized unitary classes). Let $\epsilon, \delta > 0$. Given query access to an unknown channel $\mathcal{E}$, for any finite-sized class $\mathcal{C}$ of unitaries, there exists a proper learning algorithm outputting a hypothesis $U \in \mathcal{C}$, such that

$$d_{\text{avg}}(\mathcal{U}, \mathcal{E}) \leq \text{opt}(\mathcal{E}, \mathcal{C})^{1/4} + \epsilon \quad (\text{D1})$$

with probability at least $1 - \delta$. Here, opt is measured with respect to d_{avg} , and $\mathcal{U}$ denotes the unitary quantum channel defined via $\mathcal{U}(\rho) = U\rho U^\dagger$. The algorithm makes

$$N = \mathcal{O}\left(\frac{\log(|\mathcal{C}|/\delta)}{\epsilon^4}\right) \quad (\text{D2})$$

queries to the unknown channel $\mathcal{E}$ and runs in time

$$\mathcal{O}\left(T \cdot \frac{|\mathcal{C}| \log(|\mathcal{C}|/\delta)}{\epsilon^4}\right), \quad (\text{D3})$$

where T is the time complexity of computing the inner product of two states given their exact description.

As in Appendix C 1, the time complexity in the proposition depends on the complexity of computing the inner products of pairs of states. This is again due to the use of Clifford classical shadows [111], which is not computationally efficient for arbitrary observables (see Lemma B4). We do not analyze the time complexity in detail because the scaling with $|\mathcal{C}|$ is already computationally inefficient for most classes of interest.

Proof. The algorithm and analysis are similar to the proof of Proposition 1. Namely, we use Clifford classical shadows (Item 1 of Lemma B4) to estimate the distance $d_{\text{avg}}(\mathcal{U}_i, \mathcal{E})$ between the unknown channel $\mathcal{E}$ and all unitaries $U_i \in \mathcal{C}$, where again $\mathcal{U}_i(\rho) = U_i \rho U_i^\dagger$. Then, our algorithm outputs the one with the minimal estimated distance.

Again, we prepare random input states from the ensemble

$$P = \text{Uniform}(\{|\psi\rangle = U|0\rangle, U \in \text{Cl}(n)\}). \quad (\text{D4})$$

Note that because our unknown channel may not be a unitary, the quantity estimated is not the root mean squared

trace distance but instead the average infidelity. This is clear from the proof in Sec. C.2(a) of Ref. [27]. For any measure μ over input states, we denote the average infidelity by

$$\bar{F}_\mu(\mathcal{E}_1, \mathcal{E}_2) \triangleq \mathbf{E}_{\rho \sim \mu} [1 - F(\mathcal{E}_1(\rho), \mathcal{E}_2(\rho))]. \quad (\text{D5})$$

Now, by Eq. (C.8) of Ref. [27], using

$$N = 204 \frac{\log(2|\mathcal{C}|/\delta)}{\epsilon'^2} = \mathcal{O}\left(\frac{\log(|\mathcal{C}|/\delta)}{\epsilon'^2}\right) \quad (\text{D6})$$

queries, we see that we can construct unbiased estimators $\hat{\delta}_i$ such that $1 - \hat{\delta}_i$ estimates the average infidelity with respect to P

$$|\bar{F}_P(\mathcal{U}_i, \mathcal{E}) - (1 - \hat{\delta}_i)| \leq \epsilon', \quad \forall i \in [|\mathcal{C}|] \quad (\text{D7})$$

with probability at least $1 - \delta$ for some $\epsilon' > 0$. Then, we can select $i^* = \arg \min_i (1 - \hat{\delta}_i)$ and output U_{i^*} , i.e., the $U_i \in \mathcal{C}$ with the smallest estimated distance from $\mathcal{E}$ with respect to $\bar{F}_P$. To relate this back to d_{avg} , we recall that P forms a 3-design [130–132], implying that the average infidelity over P is the same as that over the Haar measure. This gives us the following relation between the average distances:

$$d_P(\mathcal{U}_i, \mathcal{E})^2 = \mathbf{E}_{|\psi\rangle \sim P} [d_{\text{tr}}(\mathcal{U}_i(|\psi\rangle\langle\psi|), \mathcal{E}(|\psi\rangle\langle\psi|))^2] \quad (\text{D8})$$

$$\leq \mathbf{E}_{|\psi\rangle \sim P} [1 - F(\mathcal{U}_i(|\psi\rangle\langle\psi|), \mathcal{E}(|\psi\rangle\langle\psi|))] \quad (\text{D9})$$

$$= \mathbf{E}_{|\psi\rangle \sim \mu_S} [1 - F(\mathcal{U}_i(|\psi\rangle\langle\psi|), \mathcal{E}(|\psi\rangle\langle\psi|))] \quad (\text{D10})$$

$$\leq \mathbf{E}_{|\psi\rangle \sim \mu_S} [d_{\text{tr}}(\mathcal{U}_i(|\psi\rangle\langle\psi|), \mathcal{E}(|\psi\rangle\langle\psi|))] \quad (\text{D11})$$

$$\leq \sqrt{\mathbf{E}_{|\psi\rangle \sim \mu_S} [d_{\text{tr}}(\mathcal{U}_i(|\psi\rangle\langle\psi|), \mathcal{E}(|\psi\rangle\langle\psi|))^2]} \quad (\text{D12})$$

$$= d_{\text{avg}}(\mathcal{U}_i, \mathcal{E}), \quad (\text{D13})$$

where in the second line, we use Definition A10. In the third line, we use the fact that P is a 2-design. In the fourth line, we use Definition A10 and the fact that $\mathcal{U}_i(\psi)$ is a pure state. In the fifth line we use Jensen's inequality. We can now bound the error of our hypothesis with respect to

d_{avg} .

$$d_{\text{avg}}(\mathcal{U}_{i^*}, \mathcal{E})^2 = \mathbf{E}_{|\psi\rangle \sim \mu_S} [d_{\text{tr}}(\mathcal{U}_{i^*}(|\psi\rangle\langle\psi|), \mathcal{E}(|\psi\rangle\langle\psi|))^2] \quad (\text{D14})$$

$$\leq \mathbf{E}_{|\psi\rangle \sim \mu_S} [1 - F(\mathcal{U}_{i^*}(|\psi\rangle\langle\psi|), \mathcal{E}(|\psi\rangle\langle\psi|))] \quad (\text{D15})$$

$$= \mathbf{E}_{|\psi\rangle \sim P} [1 - F(\mathcal{U}_{i^*}(|\psi\rangle\langle\psi|), \mathcal{E}(|\psi\rangle\langle\psi|))] \quad (\text{D16})$$

$$\leq (1 - \hat{\delta}_{i^*}) + \epsilon' \quad (\text{D17})$$

$$= \min_i (1 - \hat{\delta}_i) + \epsilon' \quad (\text{D18})$$

$$\leq \min_i (\bar{F}_P(\mathcal{U}_i, \mathcal{E}) + \epsilon') + \epsilon' \quad (\text{D19})$$

$$\leq \min_i \left(\mathbf{E}_{|\psi\rangle \sim P} [d_{\text{tr}}(\mathcal{U}_i(|\psi\rangle\langle\psi|), \mathcal{E}(|\psi\rangle\langle\psi|))] + 2\epsilon' \right) \quad (\text{D20})$$

$$\leq \min_i d_P(\mathcal{U}_i, \mathcal{E}) + 2\epsilon' \quad (\text{D21})$$

$$\leq \min_i \sqrt{d_{\text{avg}}(\mathcal{U}_i, \mathcal{E})} + 2\epsilon' \quad (\text{D22})$$

$$= \sqrt{\text{opt}(\mathcal{E}, \mathcal{C})} + 2\epsilon' \quad (\text{D23})$$

$$\leq \left(\text{opt}(\mathcal{E}, \mathcal{C})^{1/4} + \sqrt{2\epsilon'} \right)^2. \quad (\text{D24})$$

In the second line, we use Definition A10. In the third line, we use the fact that P is a 2-design. In the fourth and sixth lines, we use Eq. (D7). In the seventh line, we use Definition A10 and the definition of $\bar{F}_P$. In the eighth line, we use Jensen's inequality. In the ninth line, we use Eq. (D13). Finally, in the second to last line, we use the definition of opt . Thus, by setting $\epsilon' = \epsilon^2/2$, we obtain

$$d_{\text{avg}}(\mathcal{U}_{i^*}, \mathcal{E}) \leq \text{opt}(\mathcal{E}, \mathcal{C})^{1/4} + \epsilon, \quad (\text{D25})$$

giving us the desired bound. The claimed sample complexity can be obtained from Eq. (D6) for our choice of ϵ' , namely,

$$N = \mathcal{O}\left(\frac{\log(|\mathcal{C}|/\delta)}{\epsilon^4}\right). \quad (\text{D26})$$

The time complexity is due to the computation of the $\hat{\delta}_i$ for all elements of $\mathcal{C}$, each of which requires computing inner products between $\mathcal{O}(N)$ pairs of states. ■

While Proposition 3 only holds for concept classes of finite size, we now show that this can be easily extended to any class as long as it has bounded metric entropy (see Definition A12).

Corollary D1. Let $\epsilon, \epsilon', \delta > 0$. Given query access to an unknown channel $\mathcal{E}$, for any class of unitaries $\mathcal{C}$ with covering number $\mathcal{N}(\mathcal{C}, d_{\text{avg}}, \epsilon')$, there exists a proper learning algorithm outputting a hypothesis U such that

$$d_{\text{avg}}(\mathcal{U}, \mathcal{E}) \leq \text{opt}(\mathcal{E}, \mathcal{C})^{1/4} + \epsilon + (\epsilon')^{1/4} \quad (\text{D27})$$

with probability at least $1 - \delta$. Here, opt is measured with respect to d_{avg} , and $\mathcal{U}$ denotes the unitary quantum channel defined via $\mathcal{U}(\rho) = U\rho U^\dagger$. The algorithm makes

$$N = \mathcal{O}\left(\frac{\log(\mathcal{N}(\mathcal{C}, d_{\text{avg}}, \epsilon')/\delta)}{\epsilon^4}\right) \quad (\text{D28})$$

queries to the unknown channel $\mathcal{E}$ and runs in time

$$\mathcal{O}(T \cdot \mathcal{N}(\mathcal{C}, d_{\text{avg}}, \epsilon') \cdot N), \quad (\text{D29})$$

where T is the time complexity of computing the inner product of two states given their exact description.

Proof. Because $\mathcal{C}$ has covering number $\mathcal{N}(\mathcal{C}, d_{\text{avg}}, \epsilon')$, there exists some ϵ' -covering net M with cardinality $\mathcal{N}(\mathcal{C}, d_{\text{avg}}, \epsilon')$. Consider applying Proposition 3 to the ϵ' -cover M . By Eq. (D23) for our choice of $\epsilon' = \epsilon^2/2$, this gives us a hypothesis $U \in M$ satisfying

$$d_{\text{avg}}(\mathcal{E}, \mathcal{U})^2 \leq \min_{V \in M} \sqrt{d_{\text{avg}}(\mathcal{E}, \mathcal{V})} + \epsilon^2, \quad (\text{D30})$$

where again we use $\mathcal{U}(\rho) = U\rho U^\dagger$ and similarly for $\mathcal{V}$. Let V^* be the optimal hypothesis from $\mathcal{C}$, i.e.,

$$V^* \triangleq \arg \min_{V \in \mathcal{C}} d_{\text{avg}}(\mathcal{E}, \mathcal{V}) \quad \text{and} \quad \text{opt}(\mathcal{E}, \mathcal{C}) = d_{\text{avg}}(\mathcal{E}, \mathcal{V}^*). \quad (\text{D31})$$

Then,

$$d_{\text{avg}}(\mathcal{E}, \mathcal{U})^2 \leq \min_{V \in M} \sqrt{d_{\text{avg}}(\mathcal{E}, \mathcal{V})} + \epsilon^2 \quad (\text{D32})$$

$$\leq \min_{V \in M} \sqrt{d_{\text{avg}}(\mathcal{E}, \mathcal{V}^*) + d_{\text{avg}}(\mathcal{V}^*, \mathcal{V})} + \epsilon^2 \quad (\text{D33})$$

$$\leq \sqrt{\text{opt}(\mathcal{E}, \mathcal{C})} + \min_{V \in M} \sqrt{d_{\text{avg}}(\mathcal{V}^*, \mathcal{V})} + \epsilon^2 \quad (\text{D34})$$

$$\leq \sqrt{\text{opt}(\mathcal{E}, \mathcal{C})} + \sqrt{\epsilon'} + \epsilon^2 \quad (\text{D35})$$

$$\leq (\text{opt}(\mathcal{E}, \mathcal{C})^{1/4} + \epsilon + (\epsilon')^{1/4})^2 \quad (\text{D36})$$

where the first inequality follows from Eq. (D23) for our choice of $\epsilon' = \epsilon^2/2$; the second inequality follows from the triangle inequality for d_{avg} ; the third inequality follows from the fact that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for all $a, b \geq 0$. The

fourth inequality uses that M ϵ' -covers $\mathcal{C}$. The sample complexity and runtime follow from Proposition 3 because $|\mathcal{C}| = \mathcal{N}(\mathcal{C}, d_{\text{avg}}, \epsilon')$. ■

From Lemma A9, we see that the class of bounded-gate unitaries has bounded metric entropy, which gives us a query-efficient algorithm for agnostic process tomography of this class. However, as the running time is of the order of the covering number, the algorithm is computationally inefficient. From Corollary D1 and Lemma A9, we have the following corollary.

Corollary D2 (Agnostic process tomography of bounded-gate unitaries). Let $\epsilon, \delta > 0$. Given query access to an unknown channel $\mathcal{E}$, for the class $\mathcal{C}$ of n -qubit unitaries of G two-qubit gates, there exists a (computationally inefficient) proper learning algorithm outputting a hypothesis $U \in \mathcal{C}$, such that with probability at least $1 - \delta$,

$$d_{\text{avg}}(\mathcal{U}, \mathcal{E}) \leq \text{opt}(\mathcal{E}, \mathcal{C})^{1/4} + \epsilon. \quad (\text{D37})$$

The algorithm makes

$$N = \tilde{\mathcal{O}}\left(\frac{G \log(n/\delta)}{\epsilon^4}\right) \quad (\text{D38})$$

queries to the unknown channel $\mathcal{E}$.

2. APT from agnostic state tomography

In Appendix C2, we showed that we can convert any agnostic state tomography algorithm (Definition C1) into one for improper agnostic unitary estimation using ancilla qubits, as long as the concept classes are compatible. This introduces a quadratic overhead in the sample and time complexity with respect to the error ϵ .

In this section, we prove a similar result for improper agnostic process tomography. The error bound we prove is not 1-agnostic but depends on $2^{1/4}$ times the square root of the optimal error, when we measure error with respect to the normalized Frobenius distance for superoperators (Definition A9). We state this result formally in the following proposition.

Proposition 2 (Improper APT from agnostic state tomography). Let $1 > \epsilon, \delta > 0$. Suppose there exists a learning algorithm for (proper) $(1, \epsilon^2, \delta)$ -agnostic state tomography with respect to fidelity F for some class $\mathcal{D}$ of pure states. Let $\mathcal{C}$ be a concept class of unitaries. Suppose that $|v(U)\rangle \in \mathcal{D}$ for all $U \in \mathcal{C}$. Given query access to an unknown channel Φ , there exists an improper learning algorithm outputting a hypothesis Φ' such that

$$d_F(\Phi, \Phi') \leq 2^{1/4} \text{opt}(\Phi, \mathcal{C})^{1/2} + \epsilon \quad (\text{D39})$$

with probability at least $1 - \delta$. Here, opt is measured with respect to d_F . Moreover, this requires the same sample and

time complexity as the $(1, \epsilon^2, \delta)$ -agnostic state tomography algorithm.

As in Appendix C 2, our algorithm simply applies agnostic state tomography to the Choi state of the unknown process, inspired by Ref. [87]. We also note that we can trivially obtain $(1, \epsilon, \delta)$ -agnostic process tomography with respect to the entanglement infidelity from $(1, \epsilon, \delta)$ -agnostic state tomography (notably, without the quadratic blowup with respect to ϵ) under the same compatibility condition. Recall that we defined the entanglement infidelity in Eq. (C30), and this metric has been widely considered in (realizable) unitary estimation [6, 87–93].

Proof. The proof is similar to that of Proposition 2. Let Φ denote the unknown unitary. The algorithm is simple: first, prepare the Choi state $v(\Phi)$ of Φ and then apply the $(1, \epsilon^2, \delta)$ -agnostic state tomography algorithm to $v(\Phi)$. Note that the known agnostic state tomography algorithms work for mixed state inputs [42, 43]. The output of our algorithm is the density matrix of the pure state output by the agnostic state tomography algorithm.

The agnostic state tomography algorithm outputs some state $|\psi\rangle$ such that

$$F(v(\Phi), |\psi\rangle\langle\psi|) \geq \max_{|\phi\rangle \in \mathcal{D}} F(v(\Phi), |\phi\rangle\langle\phi|) - \epsilon^2, \quad (\text{D40})$$

where $v(\Phi)$ denotes the Choi state of Φ . We can equivalently write the above guarantee in terms of the infidelity:

$$1 - F(v(\Phi), |\psi\rangle\langle\psi|) \leq \min_{|\phi\rangle \in \mathcal{D}} (1 - F(v(\Phi), |\phi\rangle\langle\phi|)) + \epsilon^2. \quad (\text{D41})$$

Using the inequalities relating trace distance and fidelity when one of the inputs is pure (see after Definition A10), we have

$$d_{\text{tr}}(v(\Phi), |\psi\rangle\langle\psi|)^2 \leq \min_{|\phi\rangle \in \mathcal{D}} d_{\text{tr}}(v(\Phi), |\phi\rangle\langle\phi|) + \epsilon^2. \quad (\text{D42})$$

We can relate this to the Frobenius norm using a result from Ref. [136]. Namely, they prove that

$$\frac{1}{2} \|\rho - \sigma\|_F^2 \leq d_{\text{tr}}(\rho, \sigma)^2 \leq \frac{\text{rank}(\rho)\text{rank}(\sigma)}{\text{rank}(\rho) + \text{rank}(\sigma)} \|\rho - \sigma\|_F^2. \quad (\text{D43})$$

Using this in the previous inequality, we have

$$\begin{aligned} & \frac{1}{2} \|v(\Phi) - |\psi\rangle\langle\psi|\|_F^2 \\ & \leq \min_{|\phi\rangle \in \mathcal{D}} \sqrt{\frac{\text{rank}(v(\Phi))\text{rank}(|\phi\rangle\langle\phi|)}{\text{rank}(v(\Phi)) + \text{rank}(|\phi\rangle\langle\phi|)}} \|v(\Phi) - |\phi\rangle\langle\phi|\|_F \\ & + \epsilon^2. \end{aligned} \quad (\text{D44})$$

Since $|\phi\rangle\langle\phi|$ is pure, then $\text{rank}(|\phi\rangle\langle\phi|) = 1$ so that the term in the square root can be bounded by 1. Moreover,

by definition of the normalized Frobenius distance for operators (Definition A8),

$$d_F(v(\Phi), |\psi\rangle\langle\psi|) = \frac{1}{2^n \sqrt{2}} \|v(\Phi) - |\psi\rangle\langle\psi|\|_F \quad (\text{D45})$$

since $v(\Phi), |\psi\rangle\langle\psi|$ are on $2n$ qubits. Thus, rewriting the above inequality in terms of the normalized Frobenius distance, we have

$$4^n d_F(v(\Phi), |\psi\rangle\langle\psi|)^2 \leq \sqrt{2} \cdot 2^n \min_{|\phi\rangle \in \mathcal{D}} d_F(v(\Phi), |\phi\rangle\langle\phi|) + \epsilon^2. \quad (\text{D46})$$

Taking the square root of both sides, we have

$$2^n d_F(v(\Phi), |\psi\rangle\langle\psi|) \leq 2^{1/4} \sqrt{2^n \min_{|\phi\rangle \in \mathcal{D}} d_F(v(\Phi), |\phi\rangle\langle\phi|)} + \epsilon. \quad (\text{D47})$$

Now, by the same calculation as in Appendix C 2, agnostically learning the Choi state $v(\Phi)$ with respect to $2^n d_F$ is equivalent to agnostically learning the channel Φ with respect to d_F . Note that the first d_F is the normalized Frobenius distance for operators (Definition A8) while the second d_F is the normalized Frobenius norm for superoperators (Definition A9). We repeat the calculation here since it is short:

$$d_F(\Phi, \Psi) = \frac{1}{2^n \sqrt{2}} \|\Phi - \Psi\|_F \quad (\text{D48})$$

$$= \frac{1}{2^n \sqrt{2}} \|J(\Phi) - J(\Psi)\|_F \quad (\text{D49})$$

$$= \frac{1}{2} \|v(\Phi) - v(\Psi)\|_F \quad (\text{D50})$$

$$= 2^n d_F(v(\Phi), v(\Psi)). \quad (\text{D51})$$

In the first and second lines, we use the definition of d_F for superoperators (Definition A9). In the third line, we use that the Choi representation is simply the unnormalized Choi state. Finally, in the last line, we use the definition of the normalized Frobenius distance for operators (Definition A8). This completes the proof. ■

This gives a fairly general result for converting an agnostic state tomography algorithm with respect to $\mathcal{D}$ into an improper agnostic process tomography algorithm with respect to $\mathcal{C}$ when $|v(U)\rangle \in \mathcal{D}$ for all $U \in \mathcal{C}$. However, our result is weaker than that of Proposition 2, as our algorithm is no longer 1-agnostic but has an error bound of $2^{1/4} \sqrt{\text{opt}}$ instead.

As remarked earlier, the compatibility condition holds for many classes of unitaries for which agnostic state tomography algorithms already exist with respect to a suitable class. In particular, there exist agnostic state tomography algorithms for stabilizer states, states with high stabilizer dimension, and discrete product states [42, 43]. Note

that the corresponding classes of unitaries, i.e., Clifford circuits [134], Clifford + T circuits, and unitaries consisting of a tensor product of single-qubit gates, respectively, all satisfy the required property of Proposition 4. Thus, similarly to Appendix C 2, we can directly apply Proposition 2 with the state-of-the-art sample and time complexity bounds for agnostic state tomography to obtain improper APT algorithms for the respective unitary classes. These are each direct corollaries of Corollary 6.3, Theorem 7.1, Corollary 8.3, and in Ref. [43].

Corollary D3 (Improper APT of Clifford circuits; via Corollary 6.3 in Ref. [43]). Let $1 > \epsilon, \delta > 0$. Let $\mathcal{C}$ be the class of n -qubit Clifford unitaries. Given query access to an unknown channel Φ , there exists an improper learning algorithm outputting a hypothesis Φ' such that

$$d_F(\Phi, \Phi') \leq 2^{1/4} \text{opt}(\Phi, \mathcal{C})^{1/2} + \epsilon, \quad (\text{D52})$$

with probability at least $1 - \delta$, using

$$N = 2n \log(1/\delta) \left(\frac{1}{\epsilon} \right)^{\mathcal{O}(\log 1/\epsilon)} + \mathcal{O} \left(\frac{\log(1/\delta) + \log^2(1/\epsilon)}{\epsilon^4} \right) \quad (\text{D53})$$

queries to the unknown unitary and running in time

$$\mathcal{O} \left(n^2 \log(1/\delta) \left(n + \frac{\log(1/\delta)}{\epsilon^4} \right) \right) \cdot \left(\frac{1}{\epsilon} \right)^{\mathcal{O}(\log 1/\epsilon)}. \quad (\text{D54})$$

Corollary D4 (Improper APT of Clifford + T circuits; via Theorem 7.1 in Ref. [43]). Let $1 > \epsilon, \delta > 0$ and $t \geq 0$. Let $\mathcal{C}$ be the class of n -qubit unitaries consisting of Clifford gates and t T gates. Given query access to an unknown channel Φ , there exists an improper learning algorithm outputting a hypothesis Φ' such that

$$d_F(\Phi, \Phi') \leq 2^{1/4} \text{opt}(\Phi, \mathcal{C})^{1/2} + \epsilon, \quad (\text{D55})$$

with probability at least $1 - \delta$, using

$$N = 2n \log(1/\delta) \left(\frac{2^t}{\epsilon^2} \right)^{\mathcal{O}(\log 1/\epsilon)} \quad (\text{D56})$$

queries to the unknown unitary and running in time

$$\mathcal{O}(n^2 \log(1/\delta) N). \quad (\text{D57})$$

Corollary D5 (Improper APT of product circuits; via Corollary 8.3 in Ref. [43]). Let $\mathcal{K}$ be a set of single-qubit gates for which $|\langle v(U_1) | v(U_2) \rangle|^2 \leq 1 - \mu$ for any

distinct $U_1, U_2 \in \mathcal{K}$. Let $\mathcal{C} = \mathcal{K}^{\otimes n}$. Given query access to an unknown channel Φ , there exists an improper learning algorithm outputting a hypothesis Φ' such that

$$d_F(\Phi, \Phi') \leq 2^{1/4} \text{opt}(\Phi, \mathcal{C})^{1/2} + \epsilon, \quad (\text{D58})$$

with probability at least $1 - \delta$, using

$$N = \frac{\log^2(1/\delta) (2n|\mathcal{K}|)^{\mathcal{O}(\log(1/\epsilon)/\mu)}}{\epsilon^4} \quad (\text{D59})$$

queries to the unknown unitary and running in time $\mathcal{O}(N)$.

Corollary D6 (Improper APT of Clifford product circuits; via Corollary 9.3 in Ref. [43]). Let $1 > \epsilon, \delta > 0$. Let $\mathcal{C}$ be the class of tensor products of single-qubit Clifford gates. Given query access to an unknown channel Φ , there exists an improper learning algorithm outputting a hypothesis Φ' such that

$$d_F(\Phi, \Phi') \leq 2^{1/4} \text{opt}(\Phi, \mathcal{C})^{1/2} + \epsilon, \quad (\text{D60})$$

with probability at least $1 - \delta$, using

$$N = \log(2n) \log(1/\delta) \left(\frac{1}{\epsilon} \right)^{\mathcal{O}(\log 1/\epsilon)} + \mathcal{O} \left(\frac{\log^2(1/\epsilon) + \log(1/\delta)}{\epsilon^4} \right) \quad (\text{D61})$$

and running in time

$$\frac{4n^2 \log^2(1/\delta) (1/\epsilon)^{\mathcal{O}(\log 1/\epsilon)}}{\epsilon^4}. \quad (\text{D62})$$

Note that all of these algorithms are improper, while for applications, one would prefer proper learning algorithms. We leave it to future work to make these algorithms proper.

3. Pauli strings and Pauli channels

In this section, we consider agnostic process tomography of the class of n -qubit Pauli strings $\mathcal{C}_{\mathcal{P}_n} = \{I, X, Y, Z\}^{\otimes n}$ as well as the class of n -qubit Pauli channels $\mathcal{C}_{\mathcal{PC}}$. A Pauli channel is defined as follows.

Definition D1 (Pauli channel). A quantum channel Φ on $\mathcal{B}(\mathcal{H})$ is a *Pauli channel* if

$$\Phi(\rho) = \sum_{x \in \mathcal{P}_n} p_x \sigma_x \rho \sigma_x, \quad (\text{D63})$$

where $\{p(x)\}_{x \in \mathcal{P}_n}$ forms a probability distribution.

Pauli channels form a naturally occurring class of noise channels, where the probabilities p_x denote the chance of

a particular Pauli error occurring. Clearly, Pauli channels have Fourier coefficients of the following form:

$$\hat{\Phi}(x, x) = p_x, \quad \forall x \in \mathcal{P}_n \quad (\text{D64})$$

and

$$\hat{\Phi}(x, y) = 0, \quad x \neq y. \quad (\text{D65})$$

For both classes, our results are with respect to d_F , the normalized Frobenius distance for superoperators (Definition A9). Our proof techniques are similar to those used in Appendix C3, but we use the more general Fourier spectrum of superoperators, as introduced in Appendix A3. We start with the class of Pauli strings.

Theorem D1 (Agnostic process tomography of Pauli strings in d_F). There exists a learning algorithm for proper $(1, \epsilon, \delta)$ -agnostic process tomography of the class of n -qubit Pauli strings $\mathcal{C}_{\mathcal{P}_n}$ with respect to error in d_F using

$$N = \mathcal{O}\left(\frac{\log(1/\delta)}{\epsilon^4}\right) \quad (\text{D66})$$

samples and running in time

$$\mathcal{O}(N). \quad (\text{D67})$$

Proof. The algorithm is the same as Algorithm 1, but we instead prepare the Choi state $v(\Phi)$ of the unknown quantum channel Φ . Recall that we can write the Choi state as

$$v(\Phi) = \sum_{x, y \in \mathcal{P}_n} \hat{\Phi}(x, y) |v(\sigma_x)\rangle \langle v(\sigma_y)|. \quad (\text{D68})$$

Then, by measuring $v(\Phi)$ in the basis $\{|v(\sigma_x)\rangle\}_x$, we measure x with probability

$$p(x) \triangleq \langle v(\sigma_x) | v(\Phi) | v(\sigma_x) \rangle = \hat{\Phi}(x, x). \quad (\text{D69})$$

Now, our analysis proceeds similarly to the proof of Theorem C1. Denote the true most likely Pauli string by

$$z^* \triangleq \arg \max_{z \in \mathcal{P}_n} p(z). \quad (\text{D70})$$

We compute the optimal error for the class of Pauli strings. Recall from Lemma A6 that for quantum channels Φ, Ψ ,

$$d_F(\Phi, \Psi)^2 = \frac{1}{2} \sum_{x, y \in \mathcal{P}_n} |\hat{\Phi}(x, y) - \hat{\Psi}(x, y)|^2. \quad (\text{D71})$$

Then, we have

$$\text{opt}(\Phi, \mathcal{C}_{\mathcal{P}_n})^2 = \min_{z \in \mathcal{P}_n} d_F(\Phi, \Phi_{z,z})^2 \quad (\text{D72})$$

$$= \frac{1}{2} \left(\min_{z \in \mathcal{P}_n} \sum_{\substack{x, y \in \mathcal{P}_n \\ x \neq z \text{ or } y \neq z}} |\hat{\Phi}(x, y)|^2 + |\hat{\Phi}(z, z) - 1|^2 \right) \quad (\text{D73})$$

$$= \frac{1 + \|\hat{\Phi}\|_F^2}{2} - \max_{z \in \mathcal{P}_n} \hat{\Phi}(z, z) \quad (\text{D74})$$

$$= \frac{1 + \|\hat{\Phi}\|_F^2}{2} - \hat{\Phi}(z^*, z^*), \quad (\text{D75})$$

where we use $\Phi_{z,z}$ to denote the channel $\Phi_{z,z}(\rho) = \sigma_z \rho \sigma_z$ so that $\hat{\Phi}_{z,z}(z, z) = 1$ and zero otherwise. The second equality follows from Lemma A6. In the third line, note that $\|\hat{\Phi}\|_F^2$ denotes $\|\hat{\Phi}\|_F^2 \triangleq \sum_{x, y \in \mathcal{P}_n} |\hat{\Phi}(x, y)|^2$, and we use the fact that the diagonal elements of $\hat{\Phi}$ are all real and non-negative. In the last line, we use the definition of z^* .

Hence, we see that the hypothesis σ_{z^*} achieves the optimal error for the class $\mathcal{C}_{\mathcal{P}_n}$. Our algorithm aims to estimate the true most likely Pauli string z^* by the one most frequently occurring in our measurements. Let $\{\hat{p}(x)\}_x$ be the empirical distribution obtained from our measurement outcomes, which estimates p . Then, denote the most frequent outcome out of our measurements by

$$x^* \triangleq \arg \max_{x \in \{x_i\}_{i \in [N]}} \hat{p}(x). \quad (\text{D76})$$

Using Lemma B1, for some $\epsilon_1 > 0$, the empirical estimates $\hat{p}(x)$ of $p(x) = \hat{\Phi}(x, x)$ satisfy, with probability at least $1 - \delta$,

$$|\hat{p}(x) - \hat{\Phi}(x, x)| \leq \epsilon_1, \quad \forall x \in \mathcal{P}_n, \quad (\text{D77})$$

using $\mathcal{O}(\log(1/\delta)/\epsilon_1^2)$ copies of the Choi state $|v(U)\rangle$. Now, since x^* is the most frequent string observed,

$$\hat{p}(x^*) \geq \hat{p}(z^*). \quad (\text{D78})$$

Conditioned on obtaining estimates $\hat{p}(x)$ that are ϵ_1 -close to $|\hat{U}_x|^2$ for all $x \in \mathcal{P}_n$, we have

$$\hat{\Phi}(x^*, x^*) + \epsilon_1 \geq \hat{p}(x^*) \quad \text{and} \quad \hat{p}(z^*) \geq \hat{\Phi}(z^*, z^*) - \epsilon_1. \quad (\text{D79})$$

Thus, we have

$$\hat{\Phi}(z^*, z^*) - \hat{\Phi}(x^*, x^*) \leq 2\epsilon_1. \quad (\text{D80})$$

ALGORITHM 5. 1-agnostic proper learning Pauli channels.

-
-
- 1: Prepare $N = \mathcal{O}(\log(1/\delta)/\epsilon^2)$ copies of $v(\Phi)$ and measure them in the $\{|v(\sigma_x)\rangle\}_x$ basis to construct estimates $\hat{p}(x)$ of $\hat{\Phi}(x, x)$.
 - 2: **return** $\Phi = \sum_{x \in \mathcal{P}_n} \hat{p}(x) \sigma_x(\cdot) \sigma_x$.
-
-

Now, we can bound the error of the hypothesis as follows:

$$d_F(\Phi, \sigma_{x^*})^2 = \frac{1 + \|\hat{\Phi}\|_F^2}{2} - \hat{\Phi}(x^*, x^*) \quad (\text{D81})$$

$$= \frac{1 + \|\hat{\Phi}\|_F^2}{2} - \hat{\Phi}(z^*, z^*) + \hat{\Phi}(z^*, z^*) - \hat{\Phi}(x^*, x^*) \quad (\text{D82})$$

$$\leq \text{opt}(\Phi, \mathcal{C}_{\mathcal{P}_n})^2 + 2\epsilon_1, \quad (\text{D83})$$

where the first line is true by Lemma A6 (a similar calculation as before for opt), and the last line follows by Eqs. (D75) and (D80). Thus, by setting $\epsilon_1 = \epsilon^2/2$, we obtain

$$d_F(\Phi, \sigma_{x^*})^2 \leq \text{opt}(\Phi, \mathcal{C}_{\mathcal{P}_n})^2 + \epsilon^2 \leq (\text{opt}(\Phi, \mathcal{C}_{\mathcal{P}_n}) + \epsilon)^2, \quad (\text{D84})$$

which is the desired bound. Now, by our choice of $\epsilon_1 = \epsilon^2/2$, we use

$$N = \mathcal{O}\left(\frac{\log(1/\delta)}{\epsilon_1^2}\right) = \mathcal{O}\left(\frac{\log(1/\delta)}{\epsilon^4}\right) \quad (\text{D85})$$

samples of the Choi state or queries to the channel Φ . The time complexity is that of identifying the most frequent string out of N strings, which can be done in time $\mathcal{O}(N)$. ■

Now, we present a proper 1-agnostic learning algorithm for the class of Pauli channels.

Theorem D2 (Agnostic process tomography of Pauli channels in d_F). Algorithm 5 performs proper $(1, \epsilon, \delta)$ -agnostic process tomography of the class of n -qubit Pauli channels $\mathcal{C}_{\mathcal{P}_n}$ with respect to error in d_F using

$$N = \mathcal{O}\left(\frac{\log(1/\delta)}{\epsilon^2}\right) \quad (\text{D86})$$

samples and running in time

$$\mathcal{O}(N \log(N)). \quad (\text{D87})$$

Proof. The algorithm is detailed in Algorithm 5. We start by computing the optimal error for the class $\mathcal{C}_{\mathcal{P}_n}$.

$$\text{opt}(\Phi, \mathcal{C}_{\mathcal{P}_n})^2 = \min_{\Psi \in \mathcal{C}_{\mathcal{P}_n}} d_F(\Phi, \Psi)^2 \quad (\text{D88})$$

$$= \frac{1}{2} \left(\min_{\Psi \in \mathcal{C}_{\mathcal{P}_n}} \sum_{x \in \mathcal{P}_n} |\hat{\Phi}(x, x) - \hat{\Psi}(x, x)|^2 + \sum_{\substack{x, y \in \mathcal{P}_n \\ x \neq y}} |\hat{\Phi}(x, y) - 0|^2 \right) \quad (\text{D89})$$

$$\geq \frac{1}{2} \sum_{\substack{x, y \in \mathcal{P}_n \\ x \neq y}} |\hat{\Phi}(x, y)|^2, \quad (\text{D90})$$

where the second line follows from Lemma A6 and the fact that Ψ is a Pauli channel. As $\{\hat{\Phi}(x, x)\}_{x \in \mathcal{P}_n}$ form a probability distribution, the superoperator

$$\Phi' = \sum_{x \in \mathcal{P}_n} \hat{\Phi}(x, x) \sigma_x(\cdot) \sigma_x \quad (\text{D91})$$

is a Pauli channel. Clearly,

$$d_F(\Phi, \Phi')^2 = \frac{1}{2} \sum_{\substack{x, y \in \mathcal{P}_n \\ x \neq y}} |\hat{\Phi}(x, y)|^2. \quad (\text{D92})$$

Thus, from the lower bound on opt from the above discussion, we have

$$\text{opt}(\Phi, \mathcal{C}_{\mathcal{P}_C})^2 = \frac{1}{2} \sum_{\substack{x, y \in \mathcal{P}_n \\ x \neq y}} |\hat{\Phi}(x, y)|^2. \quad (\text{D93})$$

Since measuring $v(\Phi)$ in the basis $\{|v(\sigma_x)\rangle\}_{x \in \mathcal{P}_n}$ returns samples from the distribution $\{\hat{\Phi}(x, x)\}_x$, agnostic process tomography of Pauli channels reduces to the task of learning this distribution. From Lemma B1, we see that $\mathcal{O}(\log(1/\delta)/\epsilon^2)$ samples allow us to empirically construct a distribution $\{\hat{p}(x)\}_x$, that satisfies, with probability at least $1 - \delta$,

$$\sum_{x \in \mathcal{P}_n} |\hat{\Phi}(x, x) - \hat{p}(x)|^2 \leq 2\epsilon^2. \quad (\text{D94})$$

Now, let our hypothesis channel be defined by

$$\Phi_{\hat{p}} \triangleq \sum_{x \in \mathcal{P}_n} \hat{p}(x) \sigma_x(\cdot) \sigma_x. \quad (\text{D95})$$

The error of this hypothesis is bounded by

$$\begin{aligned} d_F(\Phi, \Phi_{\hat{p}})^2 &= \frac{1}{2} \sum_{\substack{x, y \in \mathcal{P}_n \\ x \neq y}} |\hat{\Phi}(x, y)|^2 \\ &\quad + \frac{1}{2} \sum_{x \in \mathcal{P}_n} |\hat{\Phi}(x, x) - \hat{p}(x)|^2 \end{aligned} \quad (\text{D96})$$

$$\leq \text{opt}(\Phi, \mathcal{C}_{\mathcal{P}_C})^2 + \epsilon^2 \quad (\text{D97})$$

$$\leq (\text{opt}(\Phi, \mathcal{C}_{\mathcal{P}_C}) + \epsilon)^2, \quad (\text{D98})$$

where the second line follows from Eqs. (D93) and (D94), giving us the desired 1-agnostic learning guarantee. The running time is that of counting the number of occurrences of each measurement outcome, which can be done in time $\mathcal{O}(N \log(N))$. ■

4. Junta channels

In this section, we present an algorithm for agnostic process tomography of junta channels. Similar to the definition of unitary juntas, k -junta channels on n -qubits are defined as quantum channels that act non-trivially only on k qubits.

Definition D2 (k -junta channel). A quantum channel Φ on n qubits is a k -junta channel if there exists a set $S \subseteq [n]$ with $|S| = k$ such that $\Phi = \Psi_S \otimes \mathcal{I}_{S^c}$, for some k -qubit quantum channel Ψ_S .

In Appendix C4, for the problem of agnostic unitary estimation of unitary k -juntas, we identified the subset of qubits with the highest Pauli weight and then learned all Pauli coefficients restricted to this subset. We were able to identify this subset efficiently because, for an n -qubit unitary U , we could efficiently sample from the distribution $\{|\hat{U}_x^2\rangle\}_{x \in \mathcal{P}_n}$.

However, in the case of a quantum channel Φ , a similar method allows us to sample from the distribution $\{\hat{\Phi}(x, x)\}_{x \in \mathcal{P}_n}$. This does not allow us to estimate the weight of a subset S , which is given by $w_S(\Phi) = \sum_{x, y \in \mathcal{P}_S} |\hat{\Phi}(x, y)|^2$. To overcome this issue, we learn the Fourier coefficients associated with each subset and use them to estimate the weight. Unsurprisingly, this results in a worse sample complexity than the unitary junta case (Theorem C2). Once we have estimated the weights, we can construct an improper hypothesis by restricting to the Fourier coefficients corresponding to the heaviest estimated subset. The algorithm is sketched in Algorithm 6.

Theorem D3 (Agnostic process tomography of k -junta channels in d_F). Algorithm 6 performs improper $(1, \epsilon, \delta)$ -agnostic process tomography of the class of quantum k -junta channels on $\mathcal{B}(\mathcal{H})$, $\mathcal{JC}_{n,k}$, with respect to error in d_F using

$$N = \mathcal{O}\left(\frac{k 2^{12k} n^k \log(n/\delta)}{\epsilon^4}\right) \quad (\text{D99})$$

queries to the unknown channel Φ and running in time

$$\mathcal{O}(N). \quad (\text{D100})$$

Proof. The algorithm is detailed in Algorithm 6. We start by lower bounding the optimal error for the class of junta channels. For a junta channel $\Psi_S \in \mathcal{JC}_{n,k}$ acting only

ALGORITHM 6. 1-agnostic improper learning k -junta channels.

- 1: Use N queries to Φ to construct estimates $\alpha_{x,y}$ of the Fourier coefficients with $x, y \in \mathcal{P}_S$, for all subsets $S \subseteq [n]$, $|S| = k$.
 - 2: Construct estimates of weights of all subsets of $[n]$ of cardinality k using $\alpha_{x,y}$.
 - 3: Select S' to be the subset with the greatest estimated weight.
 - 4: **return** $\Psi = \sum_{x, y \in \mathcal{P}_{S'}} \alpha_{x,y} \sigma_x(\cdot) \sigma_y$.
-

on a set S with $|S| = k$, we have

$$d_F(\Phi, \Psi_S)^2 = \frac{1}{2} \sum_{x,y \in \mathcal{P}_n} |\hat{\Phi}(x,y) - \hat{\Psi}_S(x,y)|^2 \quad (\text{D101})$$

$$= \frac{1}{2} \sum_{x,y \in \mathcal{P}_S} |\hat{\Phi}(x,y) - \hat{\Psi}_S(x,y)|^2 + \frac{1}{2} \sum_{\substack{x,y \in \mathcal{P}_n \\ x \notin \mathcal{P}_S \text{ or } y \notin \mathcal{P}_S}} |\hat{\Phi}(x,y)|^2 \quad (\text{D102})$$

$$\geq 0 + \frac{1}{2} \sum_{\substack{x,y \in \mathcal{P}_n \\ x \notin \mathcal{P}_S \text{ or } y \notin \mathcal{P}_S}} |\hat{\Phi}(x,y)|^2 \quad (\text{D103})$$

$$= \frac{\|\hat{\Phi}\|_F^2 - w_S(\Phi)}{2}, \quad (\text{D104})$$

where the first line follows from Lemma A6. The second line uses that Ψ_S is a k -junta on the set S . The final equality uses the definition of weight and $\|\hat{\Phi}\|_F^2 = \sum_{x,y \in \mathcal{P}_n} |\hat{\Phi}(x,y)|^2$. Then, we can obtain a lower bound on opt by minimizing over the choice of set S and junta channel Ψ_S .

$$\text{opt}(\Phi, \mathcal{J}_{C_{n,k}})^2 = \min_{S \subseteq [n], |S|=k} \min_{\Psi_S} d_F(\Phi, \Psi_S)^2 \quad (\text{D105})$$

$$\geq \min_{S \subseteq [n], |S|=k} \min_{\Psi_S} \frac{\|\hat{\Phi}\|_F^2 - w_S(\Phi)}{2} \quad (\text{D106})$$

$$= \min_{S \subseteq [n], |S|=k} \frac{\|\hat{\Phi}\|_F^2 - w_S(\Phi)}{2}. \quad (\text{D107})$$

Let $S^* \triangleq \arg \max_{S \subseteq [n], |S|=k} w_S(\Phi)$ be the set of cardinality k with the highest Fourier weight. Then, we have

$$\text{opt}(\Phi, \mathcal{J}_{C_{n,k}})^2 \geq \frac{\|\hat{\Phi}\|_F^2 - w_{S^*}(\Phi)}{2}. \quad (\text{D108})$$

Now that we have lower bounded opt , we analyze our algorithm (Algorithm 6), which starts by estimating all Fourier coefficients associated with subsets of size k . Note that there are at most $n^k 16^k$ such Fourier coefficients. Thus, we use Lemma B3 $\mathcal{O}(n^k 16^k)$ times, to obtain estimates $\alpha_{x,y}$, such that with probability at least $1 - \delta$,

$$|\alpha_{x,y} - \hat{\Phi}(x,y)| \leq \epsilon_1, \quad \forall x,y \in \mathcal{P}_S, S \subseteq [n], |S| = k. \quad (\text{D109})$$

Using a union bound over failure probabilities of repeated applications of Lemma B3 for each Fourier coefficient, we

see that this requires

$$N = \mathcal{O}\left(\frac{k \cdot 16^k n^k \log(n/\delta)}{\epsilon_1^2}\right) \quad (\text{D110})$$

queries to Φ and runs in time $\mathcal{O}(N)$. In the following, we condition on successful estimation. Now, we use the estimated coefficients to compute estimates $\hat{w}_S$ for all subsets of cardinality k .

$$\hat{w}_S \triangleq \sum_{x,y \in \mathcal{P}_S} |\alpha_{x,y}|^2, \quad \forall S \subseteq [n], |S| = k. \quad (\text{D111})$$

As we have already estimated all Fourier coefficients associated with subsets of cardinality k by $\alpha_{x,y}$, we can now compute $\hat{w}_S$. As there are $\mathcal{O}(n^k)$ subsets of cardinality k , and computing each weight requires adding weights of 16^k coefficients, finding all the weights takes time $\mathcal{O}(n^k \cdot 16^k)$. We now bound the error of these estimates.

$$|\hat{w}_S - w_S(\Phi)| = \left| \sum_{x,y \in \mathcal{P}_S} |\alpha_{x,y}|^2 - |\hat{\Phi}(x,y)|^2 \right| \quad (\text{D112})$$

$$\leq \sum_{x,y \in \mathcal{P}_S} ||\alpha_{x,y}|^2 - |\hat{\Phi}(x,y)|^2| \quad (\text{D113})$$

$$= \sum_{x,y \in \mathcal{P}_S} ||\alpha_{x,y}| - |\hat{\Phi}(x,y)|| \cdot (|\alpha_{x,y}| + |\hat{\Phi}(x,y)|) \quad (\text{D114})$$

$$\leq \sum_{x,y \in \mathcal{P}_S} |\alpha_{x,y} - \hat{\Phi}(x,y)| \cdot (|\alpha_{x,y}| + |\hat{\Phi}(x,y)|) \quad (\text{D115})$$

$$\leq \sum_{x,y \in \mathcal{P}_S} \epsilon_1 \cdot 2 \quad (\text{D116})$$

$$= 2 \cdot 16^k \epsilon_1. \quad (\text{D117})$$

The first line follows from the definitions of the weights. The second line follows by the triangle inequality. The fourth line follows by the reverse triangle inequality. In the fifth line, we use Eq. (D109) and the fact that the absolute value of a Fourier coefficient is at most 1. In the last line, we sum over all 16^k Fourier coefficients with support in S .

Let the subset with the highest estimated weight be S' .

$$S' \triangleq \arg \max_{S \subseteq [n], |S|=k} \hat{w}_S. \quad (\text{D118})$$

Then, our hypothesis will be

$$\Psi \triangleq \sum_{x,y \in \mathcal{P}_{S'}} \alpha_{x,y} \Phi_{x,y}, \quad (\text{D119})$$

where recall we denote $\Phi_{x,y} = \sigma_x(\cdot) \sigma_y$. This simply aims to approximate Φ by estimating all Fourier coefficients of

the highest weight subset. It remains to bound the error of this hypothesis, i.e., $d_F(\Phi, \Psi)$. Denote by

$$\Phi_S \triangleq \sum_{x,y \in \mathcal{P}_S} \hat{\Phi}(x,y) \Phi_{x,y}, \quad (\text{D120})$$

the Fourier truncation of Φ restricted to the subset S , for some subset $S \subseteq [n]$. Now, using the triangle inequality for d_F , we have

$$d_F(\Phi, \Psi) \leq d_F(\Phi, \Phi_{S'}) + d_F(\Phi_{S'}, \Psi). \quad (\text{D121})$$

We can upper bound $d_F(\Phi, \Phi_{S'})$ as

$$d_F(\Phi, \Phi_{S'})^2 = \frac{1}{2} \sum_{x,y \in \mathcal{P}_n} |\hat{\Phi}(x,y) - \hat{\Phi}_{S'}(x,y)|^2 \quad (\text{D122})$$

$$= \frac{1}{2} \sum_{x,y \in \mathcal{P}_{S'}} |\hat{\Phi}(x,y) - \hat{\Phi}(x,y)|^2 + \frac{1}{2} \sum_{\substack{x,y \in \mathcal{P}_n \\ x \notin \mathcal{P}_{S'} \text{ or } y \notin \mathcal{P}_{S'}}} |\hat{\Phi}(x,y)|^2 \quad (\text{D123})$$

$$= \frac{\|\hat{\Phi}\|_F^2 - w_{S'}(\Phi)}{2} \quad (\text{D124})$$

$$= \frac{\|\hat{\Phi}\|_F^2 - w_{S^*}(\Phi)}{2} + \frac{w_{S^*}(\Phi) - w_{S'}(\Phi)}{2} \quad (\text{D125})$$

$$\leq \text{opt}(\Phi, \mathcal{JC}_{n,k})^2 + \frac{w_{S^*}(\Phi) - w_{S'}(\Phi)}{2}, \quad (\text{D126})$$

where the last inequality follows from Eq. (D108). To bound the differences in weights, observe that from Eq. (D117),

$$w_{S'}(\Phi) + 2 \cdot 16^k \epsilon_1 \geq \hat{w}_{S'} \quad \text{and} \quad \hat{w}_{S^*} \geq w_{S^*}(\Phi) - 2 \cdot 16^k \epsilon_1. \quad (\text{D127})$$

As S' is the set with the heaviest estimated weight,

$$\hat{w}_{S'} \geq \hat{w}_{S^*}. \quad (\text{D128})$$

Thus, we finally obtain

$$\begin{aligned} d_F(\Phi, \Phi_{S'})^2 &\leq \text{opt}(\Phi, \mathcal{JC}_{n,k})^2 + 2 \cdot 16^k \epsilon_1 \\ &\leq \left(\text{opt}(\Phi, \mathcal{JC}_{n,k}) + 4^k \sqrt{2\epsilon_1} \right)^2. \end{aligned} \quad (\text{D129})$$

Now, it remains to bound $d_F(\Phi_{S'}, \Psi)$.

$$d_F(\Phi_{S'}, \Psi)^2 = \frac{1}{2} \sum_{x,y \in \mathcal{P}_{S'}} |\hat{\Phi}(x,y) - \alpha_{x,y}|^2 \leq \frac{16^k \epsilon_1^2}{2}, \quad (\text{D130})$$

where the inequality follows by Eq. (D109) and the fact that there are at most 16^k Fourier coefficients for $x, y \in \mathcal{P}_{S'}$

since $|S'| = k$. Thus, with probability at least $1 - \delta$, we have

$$d_F(\Phi, \Psi) \leq \text{opt}(\Phi, \mathcal{JC}_{n,k}) + 4^k \sqrt{2\epsilon_1} + \frac{4^k \epsilon_1}{\sqrt{2}}. \quad (\text{D131})$$

To obtain the 1-agnostic learning guarantee, we choose $\epsilon_1 = \mathcal{O}(\epsilon^2/16^k)$. This gives us the stated query complexity. To show the stated time complexity, we observe that it is dominated by that of learning the Fourier coefficients. ■

Remark D1. When considering improper learning, we remark that APT with respect to the class of junta channels $\mathcal{JC}_{n,k}$ immediately implies APT with respect to the class of unitary juntas $\mathcal{J}_{n,k}$. To see this, we note that the class $\mathcal{J}_{n,k}$ is contained within $\mathcal{JC}_{n,k}$, and thus must have a higher optimal error.

We discuss a potential method for taking this improper hypothesis and obtaining a proper one in Appendix D 7.

5. Low-degree channels

In this section, we present an algorithm for agnostic process tomography of low-degree channels. We define a low-degree channel similarly to the definition of a low-degree unitary.

Definition D3 (Low-degree channels). A quantum channel Φ is said to have *degree* d if

$$w^{>d}(\Phi) = 0, \quad (\text{D132})$$

where $w^{>d}$ is defined in Definition A7. In other words, the degree of a channel Φ is the minimum integer d such that if $|x|, |y| > d$, then $\hat{\Phi}(x,y) = 0$.

We note a difference between our definition of degree- d channels and that of Ref. [30]. Our definition requires that for all $\hat{\Phi}(x,y) \neq 0$, both $|x|$ and $|y|$ are at most d . On the other hand, the definition of Ref. [30] requires that for all non-zero $\hat{\Phi}(x,y)$, $|x| + |y| \leq d$. With our chosen definition, the channels corresponding to degree- d unitaries (Definition C3) are also channels of degree- d . Thus, we find this to be a more natural definition, and we also use this fact in Remark D2. However, the definitions do not differ significantly, and one can show that an approach similar to Algorithm 7 allows for agnostic tomography of low-degree channels adhering to the definition in Ref. [30].

Now, we will show that by learning all low-degree Fourier coefficients of the channel, we obtain an improper 1-agnostic hypothesis. This is similar to the first step of the junta channel learning algorithm (Theorem D3). The algorithm is detailed in Algorithm 7.

ALGORITHM 7. 1-agnostic improper learning degree- d channels.

-
- 1: Use N queries to Φ to construct estimates $\alpha_{x,y}$ of the Fourier coefficients for all $x, y \in \mathcal{P}_n$ such that $|x|, |y| \leq d$.
 2: **return** $\Psi = \sum_{\substack{x,y \in \mathcal{P}_n \\ |x|, |y| \leq d}} \alpha_{x,y} \sigma_x(\cdot) \sigma_y$.
-

Theorem D4 (Agnostic process tomography of low-degree channels in d_F). Algorithm 7 performs improper $(1, \epsilon, \delta)$ -agnostic process tomography of the class of quantum channels on $\mathcal{B}(\mathcal{H})$ of degree- d , $\mathcal{LC}_{n,d}$, with respect to error in d_F using

$$N = \mathcal{O}\left(\frac{d \cdot 81^d n^{4d} \log(n/\delta)}{\epsilon^2}\right) \quad (\text{D133})$$

queries to the unknown channel Φ and running in time

$$\mathcal{O}(N). \quad (\text{D134})$$

Proof. The algorithm is detailed in Algorithm 7, and the analysis is similar to that of Theorem D3. We start by lower bounding the optimal error for the class of low-degree channels. We have

$$\text{opt}(\Phi, \mathcal{LC}_{n,d})^2 = \min_{\Psi \in \mathcal{LC}_{n,d}} d_F^2(\Phi, \Psi) \quad (\text{D135})$$

$$= \frac{1}{2} \min_{\Psi \in \mathcal{LC}_{n,d}} \sum_{x,y \in \mathcal{P}_n} |\hat{\Phi}(x,y) - \hat{\Psi}(x,y)|^2 \quad (\text{D136})$$

$$= \frac{1}{2} \min_{\Psi \in \mathcal{LC}_{n,d}} \sum_{\substack{x,y \in \mathcal{P}_n \\ |x|, |y| \leq d}} |\hat{\Phi}(x,y) - \hat{\Psi}(x,y)|^2 + \frac{1}{2} \sum_{\substack{x,y \in \mathcal{P}_n \\ |x| > d \text{ or } |y| > d}} |\hat{\Phi}(x,y)|^2 \quad (\text{D137})$$

$$\geq 0 + \frac{1}{2} \sum_{\substack{x,y \in \mathcal{P}_n \\ |x| > d \text{ or } |y| > d}} |\hat{\Phi}(x,y)|^2 \quad (\text{D138})$$

$$= \frac{\|\hat{\Phi}\|_F^2 - w^{\leq d}(\Phi)}{2}, \quad (\text{D139})$$

where the second line follows from Lemma A6. In the third line, we use that Ψ has degree d . The final equality uses the definition of weight (Definition A7) and $\|\hat{\Phi}\|_F^2 = \sum_{x,y \in \mathcal{P}_n} |\hat{\Phi}(x,y)|^2$.

Our algorithm (Algorithm 7) estimates all Fourier coefficients of degree at most d . Note that there are at most $\mathcal{O}(n^{2d} 9^d)$ of them. We use Lemma B3 $\mathcal{O}(n^{2d} 9^d)$ times, to

obtain estimates $\alpha_{x,y}$ such that

$$|\alpha_{x,y} - \hat{\Phi}(x,y)| \leq \epsilon_1, \quad \forall x, y \in \mathcal{P}_n \text{ s.t. } |x|, |y| \leq d, \quad (\text{D140})$$

with probability at least $1 - \delta$. Using a union bound over failure probabilities of repeated applications of Lemma B3 for each Fourier coefficient, we see that this requires

$$N = \mathcal{O}\left(\frac{d \cdot 9^d n^{2d} \log(n/\delta)}{\epsilon_1^2}\right) \quad (\text{D141})$$

queries to Φ and runs in time $\mathcal{O}(N)$. Now, we construct our hypothesis

$$\Psi \triangleq \sum_{\substack{x,y \in \mathcal{P}_n \\ |x|, |y| \leq d}} \alpha_{x,y} \Phi_{x,y}. \quad (\text{D142})$$

This simply aims to approximate Φ by estimating all low-degree coefficients. It remains to bound the error of this hypothesis, i.e., $d_F(\Phi, \Psi)$. Denote the low-degree Fourier truncation of Φ by

$$\Phi_d \triangleq \sum_{\substack{x,y \in \mathcal{P}_n \\ |x|, |y| \leq d}} \hat{\Phi}(x,y) \Phi_{x,y}. \quad (\text{D143})$$

Now, using the triangle inequality for d_F , we have

$$d_F(\Phi, \Psi) \leq d_F(\Phi, \Phi_d) + d_F(\Phi_d, \Psi). \quad (\text{D144})$$

We can upper bound $d_F(\Phi, \Phi_d)$ as

$$d_F(\Phi, \Phi_d)^2 = \frac{1}{2} \sum_{x,y \in \mathcal{P}_n} |\hat{\Phi}(x,y) - \hat{\Phi}_d(x,y)|^2 \quad (\text{D145})$$

$$= \frac{1}{2} \sum_{\substack{x,y \in \mathcal{P}_n \\ |x|, |y| \leq d}} |\hat{\Phi}(x,y) - \hat{\Phi}(x,y)|^2$$

$$+ \frac{1}{2} \sum_{\substack{x,y \in \mathcal{P}_n \\ |x| > d \text{ or } |y| > d}} |\hat{\Phi}(x,y)|^2 \quad (\text{D146})$$

$$= \frac{\|\hat{\Phi}\|_F^2 - w^{\leq d}(\Phi)}{2} \quad (\text{D147})$$

$$\leq \text{opt}(\Phi, \mathcal{LC}_{n,d})^2. \quad (\text{D148})$$

The first line follows by Lemma A6. The second line uses the definition of Φ_d . The third line uses the definition

of weight (Definition A7). Finally, the last line uses Eq. (D139). Now, conditioned on the successful estimation of $\alpha_{x,y}$, we can bound $d_F(\Phi_d, \Psi)$.

$$d_F(\Phi_d, \Psi)^2 = \frac{1}{2} \sum_{x,y \in \mathcal{P}_n} |\hat{\Phi}_d(x,y) - \alpha_{x,y}|^2 \quad (\text{D149})$$

$$= \frac{1}{2} \sum_{\substack{x,y \in \mathcal{P}_n \\ |x|, |y| \leq d}} |\hat{\Phi}(x,y) - \alpha_{x,y}|^2 \quad (\text{D150})$$

$$\leq \mathcal{O}(n^{2d} 9^d \epsilon_1^2). \quad (\text{D151})$$

In the second line, we use the definition of Φ_d . In the third line, we use Eq. (D140) and that there are at most $\mathcal{O}(n^{2d} 9^d)$ Fourier coefficients of degree at most d . Thus, with probability at least $1 - \delta$, we have

$$d_F(\Phi, \Psi) \leq \text{opt}(\Phi, \mathcal{L}_{n,d}) + \mathcal{O}(n^{d/2} 3^d \epsilon_1). \quad (\text{D152})$$

To obtain the 1-agnostic learning guarantee, we choose $\epsilon_1 = \mathcal{O}(\epsilon/n^{d/2} 3^d)$. This gives us the stated query and time complexities. ■

Remark D2. Using a similar argument to Remark D1, we note that improper APT with respect to the class of low-degree channels $\mathcal{L}_{n,d}$ implies improper APT with respect to the class of low-degree unitaries $\mathcal{L}_{n,d}$.

Similar to the case of junta channels, we will discuss a potential method for taking this improper hypothesis and obtaining a proper one in Appendix D7.

6. n -to-1 QAC^0 channels

In this section, we consider the class of channels QAC^0 formed by constant-depth circuits with arbitrary single-qubit gates and controlled phase-flip gates of arbitrary width [86]. Following standard terminology, we refer to the *size* of a quantum circuit as the number of multi-qubit gates, and the *depth* as the number of layers of multi-qubit gates in the circuit. We consider the algorithm of Ref. [22] and show that it is an improper 1-agnostic learning algorithm. We follow the definition of Ref. [22] for the circuit class QAC^0 .

Definition D4 (QAC^0 circuits; Definition 16 in Ref. [22]). The class QAC^0 contains quantum circuit families consisting of multiple-qubit CZ gates and arbitrary single-qubit gates, with depth bounded below a constant. That is, the circuit family $\{C_n\}_{n \in \mathbb{N}}$ is in QAC^0 if there exists a $d \in \mathbb{N}$ such that for each $n \in \mathbb{N}$, C_n is a quantum circuit on n qubits consisting only of multiple-qubit CZ gates and arbitrary single-qubit gates, with depth at most d .

The algorithm of Ref. [22] focuses on learning a specific class of channels formed by QAC^0 circuits. In particular,

they consider the class of channels formed by tracing out all except a single qubit.

Definition D5 (n -to-1 QAC^0 channels; Definition 12 in Ref. [22]). Let $n, a \in \mathbb{N}$. Let U be an $(n+a)$ -qubit unitary implemented by a QAC^0 circuit, and let ψ be an a -qubit quantum state. We define the n -to-1 QAC^0 channel $\mathcal{E}_{U,\psi}$ to be the following channel mapping n qubits to 1 qubit

$$\mathcal{E}_{U,\psi}(\rho) \triangleq \text{tr}_{[n+a-1]}(U(\rho \otimes \psi)U^\dagger), \quad (\text{D153})$$

where $\text{tr}_{[n+a-1]}$ denotes the partial trace over all except the last qubit.

We consider the Pauli spectrum expansion (see Definition A3) of the Choi representation of such channels. Namely, for an n -to-1 QAC^0 channel $\mathcal{E}_{U,\psi}$, the Choi representation $J(\mathcal{E}_{U,\psi}) \in \mathcal{B}((\mathbb{C}^2)^{\otimes n+1})$ can be expressed as

$$J(\mathcal{E}_{U,\psi}) = \sum_{x \in \mathcal{P}_{n+1}} \alpha_x \sigma_x. \quad (\text{D154})$$

One of the key technical results of Ref. [22] is that when expressing the Choi representation of n -to-1 QAC^0 channels in the Pauli spectrum, the weight is concentrated on the low-degree terms. Formally, we have the following lemma.

Lemma D1 (Low-degree concentration of QAC^0 channels; Theorem 18 in Ref. [22]). Let U be a unitary implemented by an $(n+a)$ -qubit QAC^0 circuit of depth d and size s . Let $\psi \in \mathcal{B}((\mathbb{C}^2)^{\otimes a})$ be an a -qubit quantum state. For each $k \in [n+1]$,

$$w^{>k}(J(\mathcal{E}_{U,\psi})) \triangleq \sum_{\substack{x \in \mathcal{P}_{n+1} \\ |x| > k}} |\alpha_x|^2 \leq \mathcal{O}\left(s^2 2^{-k/d}\right) \cdot 2^a \|\psi\|_F^2. \quad (\text{D155})$$

Our distance metric for learning QAC^0 channels is the normalized Frobenius distance, as defined in Definition A9, which is equal to the normalized Frobenius distance between Choi representations. From Lemma D1, we see that for improper learning of the Choi representation within normalized Frobenius distance, one simply needs to learn all of its low-degree Pauli coefficients. We describe the algorithm in Algorithm 8.

Theorem D5 (Improper 1-APT of n -to-1 QAC^0 channels in d_F). There exists a learning algorithm for improper $(1, \epsilon, \delta)$ -agnostic process tomography of $\text{QAC}^0(n, a, s, d)$, the class of n -to-1 QAC^0 channels $\{\mathcal{E}_{U,|0^a\rangle}\}$, where the unitaries are implemented by $(n+a)$ -qubit QAC^0 circuits

ALGORITHM 8. 1-agnostic improper learning n -to-1 QAC⁰ channels.

-
- 1: Prepare N copies of $v(\Phi)$.
 - 2: Use classical shadows for $v(\Phi)$ to obtain estimates $\hat{\alpha}_x$ of the Pauli coefficients for $x \in \mathcal{P}_{n+1}$ s.t. $|x| \leq \tilde{k} = \log^d(s^4/\epsilon^4)$.
 - 3: **return** $A = \sum_{\substack{x \in \mathcal{P}_{n+1} \\ |x| \leq \tilde{k}}} \hat{\alpha}_x \sigma_x$.
-

with size s and depth d , and $a \leq 2 \log(s/\epsilon)$, with respect to error in d_F , using

$$N = n^{\mathcal{O}(\log^d(s^4/\epsilon^4))} \cdot \log(1/\delta) \quad (\text{D156})$$

queries to the unknown channel Φ and running in time

$$n^{\mathcal{O}(\log^d(s^4/\epsilon^4))} \cdot \log(1/\delta). \quad (\text{D157})$$

We define the following classes of low-degree operators, which we agnostically learn as part of our proof of Theorem D5.

Definition D6 (Low-degree operators). Define the class of k -degree, $n+1$ -qubit bounded linear operators $\mathcal{O}_k \subseteq \mathcal{B}((\mathbb{C}^2)^{\otimes n+1})$ as

$$\mathcal{O}_k \triangleq \{A \in \mathcal{B}((\mathbb{C}^2)^{\otimes n+1}) : w^{>k}(A) = 0\}. \quad (\text{D158})$$

In order to prove Theorem D5, we will make use of the following lemma on learning the closest low-degree operator to the Choi representation of an unknown channel.

Lemma D2 (Agnostic-learning a low-degree operator). Let $\epsilon, \delta > 0$. Given copies of the Choi-state $v(\Phi)$ of an unknown quantum channel Φ from $\mathcal{B}((\mathbb{C}^2)^{\otimes n})$ to $\mathcal{B}(\mathbb{C}^2)$, there exists an algorithm that outputs a hypothesis $A \in \mathcal{O}_k$, such that

$$d_F(J(\Phi), A) \leq \min_{B \in \mathcal{O}_k} d_F(J(\Phi), B) + \epsilon, \quad (\text{D159})$$

with probability at least $1 - \delta$. Moreover, this algorithm uses

$$N = \frac{n^{\mathcal{O}(k)} \log(1/\delta)}{\epsilon^2} \quad (\text{D160})$$

copies of $v(\Phi)$ and runs in time

$$kn^{\mathcal{O}(k)}N. \quad (\text{D161})$$

Proof. Recall from Eq. (D154) that we can write $J(\Phi)$ in the Pauli basis as

$$J(\Phi) = \sum_{x \in \mathcal{P}_{n+1}} \alpha_x \sigma_x. \quad (\text{D162})$$

Now, we compute the optimal error over the class of low-degree operators $\mathcal{O}_k$.

$$\min_{B \in \mathcal{O}_k} d_F(J(\Phi), B)^2 = \min_{B \in \mathcal{O}_k} \sum_{x \in \mathcal{P}_{n+1}} \frac{|\alpha_x - \hat{B}_x|^2}{2} \quad (\text{D163})$$

$$= \min_{B \in \mathcal{O}_k} \sum_{\substack{x \in \mathcal{P}_{n+1} \\ |x| \leq k}} \frac{|\alpha_x - \hat{B}_x|^2}{2} + \sum_{\substack{x \in \mathcal{P}_{n+1} \\ |x| > k}} \frac{|\alpha_x|^2}{2} \quad (\text{D164})$$

$$\geq \min_{B \in \mathcal{O}_k} \sum_{\substack{x \in \mathcal{P}_{n+1} \\ |x| > k}} \frac{|\alpha_x|^2}{2} \quad (\text{D165})$$

$$= \frac{w^{>k}(J(\Phi))}{2}. \quad (\text{D166})$$

In the first line, we use Lemma A5. In the second line, we use that B is low-degree. One can check that this lower bound is achieved by the low-degree truncation of $J(\Phi)$, i.e., the operator

$$B^* \triangleq \sum_{\substack{x \in \mathcal{P}_{n+1} \\ |x| \leq k}} \alpha_x \sigma_x \quad (\text{D167})$$

has error

$$\min_{B \in \mathcal{O}_k} d_F(J(\Phi), B) = d_F(J(\Phi), B^*) = \sqrt{\frac{w^{>k}(J(\Phi))}{2}}. \quad (\text{D168})$$

Thus, it suffices to learn the low-degree truncation of $J(\Phi)$. First, we observe that because Φ is trace-preserving, $\text{tr}(J(\Phi)) = 2^n$, the dimension of the input system. Thus, we have

$$v(\Phi) = \frac{J(\Phi)}{2^n} = \frac{1}{2^n} \sum_{x \in \mathcal{P}_{n+1}} \alpha_x \sigma_x. \quad (\text{D169})$$

To estimate the Pauli coefficients α_x , we use classical shadows (Lemma B4) to estimate the expectation values of the

corresponding Pauli σ_x with respect to the Choi state $v(\Phi)$. These expectation values indeed tell us information about the Pauli coefficients: for any $y \in \mathcal{P}_{n+1}$,

$$\text{tr}(v(\Phi)\sigma_y) = \sum_{x \in \mathcal{P}_{n+1}} \frac{\alpha_x}{2^n} \text{tr}(\sigma_x \sigma_y) \quad (\text{D170})$$

$$= \sum_{x \in \mathcal{P}_{n+1}} \frac{\alpha_x}{2^n} \delta_{x,y} \text{tr}(I^{\otimes n+1}) \quad (\text{D171})$$

$$= 2\alpha_y. \quad (\text{D172})$$

Thus, by estimating the expectation values of low-degree Paulis with respect to the Choi state $v(\Phi)$, we can obtain estimates of the Pauli coefficients of $J(\Phi)$. We need to estimate $n^{\mathcal{O}(k)}$ Paulis of degree at most k . Let the desired error in each estimate be at most $\epsilon_1 > 0$ and the overall failure probability be at most $\delta > 0$. Then, from Lemma B4, we see that this can be done using

$$N = \mathcal{O}\left(\frac{k \cdot 3^k \log(n/\delta)}{\epsilon_1^2}\right) \quad (\text{D173})$$

copies of $v(\Phi)$ and in time

$$kn^{\mathcal{O}(k)}N. \quad (\text{D174})$$

Suppose this procedure gives us estimates $\hat{\beta}_x$, for all strings $x \in \mathcal{P}_{n+1}$ of degree at most k . By Lemma B4, we have

$$|\hat{\beta}_x - \text{tr}(v(\Phi)\sigma_x)| \leq \epsilon_1, \quad \forall x \in \mathcal{P}_{n+1} \text{ s.t. } |x| \leq k \quad (\text{D175})$$

with probability at least $1 - \delta$. By Eq. (D172), we then have

$$|\hat{\beta}_x - 2\alpha_x| \leq \epsilon_1, \quad \forall x \in \mathcal{P}_{n+1} \text{ s.t. } |x| \leq k. \quad (\text{D176})$$

Then, define our hypothesis by

$$A \triangleq \sum_{\substack{x \in \mathcal{P}_{n+1} \\ |x| \leq k}} \frac{\hat{\beta}_x}{2} \sigma_x. \quad (\text{D177})$$

It remains to bound the error of this hypothesis.

$$d_F(J(\Phi), A) \leq d_F(J(\Phi), B^*) + d_F(B^*, A) \quad (\text{D178})$$

$$\leq \min_{B \in \mathcal{O}_k} d_F(J(\Phi), B) + \frac{1}{\sqrt{2}} \sqrt{\sum_{\substack{x \in \mathcal{P}_{n+1} \\ |x| \leq k}} \left| \alpha_x - \frac{\hat{\beta}_x}{2} \right|^2} \quad (\text{D179})$$

$$\leq \min_{B \in \mathcal{O}_k} d_F(J(\Phi), B) + n^{\mathcal{O}(k)} \frac{\epsilon_1}{\sqrt{2}}. \quad (\text{D180})$$

In the first line, we use the triangle inequality for d_F . In the second line, we use the definition of B^* and A and Lemma

A5. In the third line, we use Eq. (D176) and the fact that there are at most $n^{\mathcal{O}(k)}$ Pauli coefficients of degree at most k . Then, setting $\epsilon_1 = \epsilon/n^{\mathcal{O}(k)}$, we obtain the desired bound

$$d_F(J(\Phi), A) \leq \min_{B \in \mathcal{O}_k} d_F(J(\Phi), B) + \epsilon. \quad (\text{D181})$$

Our choice of ϵ_1 can be used to derive the sample and time complexities. ■

With this, we can prove Theorem D5.

Proof of Theorem D5. The algorithm is described in Algorithm 8. For the unknown quantum channel Φ , let Φ^* be the closest channel from the class of channels $\text{QAC}^0(n, a, s, d)$, as defined in the theorem statement. In other words, Φ^* achieves $\text{opt}(\Phi, \text{QAC}^0(n, a, s, d))$. Then, from Lemma D1, we know that the Choi representation of this channel satisfies

$$w^{>k}(J(\Phi^*)) \leq \mathcal{O}\left(s^2 2^{-k^{1/d}+a}\right). \quad (\text{D182})$$

Then, setting $\tilde{k} \triangleq \log^d(s^4/\epsilon^4)$, and using $a \leq 2 \log(s/\epsilon) = \frac{1}{2}\tilde{k}^{1/d}$, we have

$$w^{>\tilde{k}}(J(\Phi^*)) \leq \mathcal{O}(\epsilon^2). \quad (\text{D183})$$

Our algorithm (Algorithm 8) is simply to apply Lemma D2 to the Choi state $v(\Phi)$ of the channel Φ . This outputs a low-degree operator $A \in \mathcal{O}_{\tilde{k}}$, which is also our final hypothesis. We can bound the error $d_F(J(\Phi), A)$ as

$$d_F(J(\Phi), A) \leq \min_{B \in \mathcal{O}_{\tilde{k}}} d_F(J(\Phi), B) + \epsilon \quad (\text{D184})$$

$$\leq \min_{B \in \mathcal{O}_{\tilde{k}}} d_F(J(\Phi), J(\Phi^*)) + d_F(J(\Phi^*), B) + \epsilon \quad (\text{D185})$$

$$\leq d_F(J(\Phi), J(\Phi^*)) + \min_{B \in \mathcal{O}_{\tilde{k}}} d_F(J(\Phi^*), B) + \epsilon \quad (\text{D186})$$

$$= d_F(J(\Phi), J(\Phi^*)) + \sqrt{\frac{w^{>\tilde{k}}(J(\Phi^*))}{2}} + \epsilon \quad (\text{D187})$$

$$\leq d_F(J(\Phi), J(\Phi^*)) + \mathcal{O}(\epsilon) \quad (\text{D188})$$

$$= \min_{\Psi \in \text{QAC}^0(n, a, s, d)} d_F(J(\Phi), J(\Psi)) + \mathcal{O}(\epsilon), \quad (\text{D189})$$

where the second line is by triangle inequality. The fourth line follows from Eq. (D168), and the last line is due to Eq. (D183). Due to the equivalence of d_F between channels and their Choi representations, agnostic learning the Choi representation is the same as agnostic learning a description of the channel. Thus, by rescaling ϵ by

an appropriate constant, we obtain the required improper 1-agnostic learning bound. To obtain the stated sample and time complexities, we simply plug in the values of k and ϵ in the bounds of Lemma D2. ■

Unlike the case for APT of junta channels and low-degree channels, the method discussed in Appendix D7 does not help us obtain a proper hypothesis for QAC^0 channels. As discussed in Ref. [22], one can map the problem of obtaining the nearest quantum channel from the improper hypothesis to a convex optimization problem. However, it is not clear how to obtain the nearest quantum channel that can be implemented using a QAC^0 circuit. As a result, we only achieve improper agnostic process tomography of QAC^0 channels.

7. Proper APT

In this section, we discuss an approach towards proper learning for APT based on convex optimization. In particular, we aim to devise a projection algorithm which takes as input an improper hypothesis output by one of our algorithms (e.g., Appendices D4 and D5) and returns a valid quantum channel. We are unable to fully solve this problem due to a difficult-to-verify condition from Ref. [109], which we rephrase in our setting in as Conjecture 1. Nonetheless, we consider this as a significant first step towards resolving this open question, which is also recurrent in much of the realizable process tomography literature that uses Pauli spectrum analysis [22,24,25]. Note that ideas from the quasiprobability decomposition and probabilistic error cancellation which approximate linear superoperators by linear combinations of quantum channels [81,137–144] do not help because we need to find exactly a quantum channel.

Suppose that we are given an improper hypothesis Ψ_1 , from an agnostic process tomography algorithm of a concept class $\mathcal{C}$ and an unknown channel Φ , with respect to the normalized Frobenius distance d_F . Denote the Fourier coefficient matrix of Φ by $\hat{\Phi}$. Moreover, suppose that our hypothesis is given to us as a Hermitian matrix of Fourier coefficients, $\hat{\Psi}_1$. This aligns with our learning algorithms from Appendices D4 and D5, where we learn the Fourier coefficients corresponding to a principal submatrix of $\hat{\Phi}$ and simply set the rest to 0. To ensure that $\hat{\Psi}_1$ is Hermitian, we can instead only learn the upper triangular coefficients of this principal submatrix, and set the corresponding coefficients in the lower triangle as their complex conjugates. From the improper APT guarantee, we have

$$d_F(\Phi, \Psi_1) \leq \text{opt}(\Phi, \mathcal{C}) + \epsilon_1, \quad (\text{D190})$$

for some $\epsilon_1 > 0$. Our goal in the rest of this section will be to construct a *proper* hypothesis satisfying a similar error guarantee. We break this up into several steps in the following discussion, which we outline below.

- (1) First, we derive constraints that the Fourier coefficient matrix must satisfy in order to be a valid quantum channel. Thus, we can formulate our projection task as a (complex) convex optimization problem.
- (2) Second, we map this problem to another convex optimization problem over real, symmetric matrices to allow us to leverage powerful tools in the optimization literature to solve it.
- (3) Next, we phrase our problem as a well-studied convex optimization problem with known algorithms and discuss the key condition needed for fast convergence.
- (4) Although we may have fast convergence, in a finite number of steps, this does not guarantee *exact* convergence to a quantum channel. By mixing the output of the optimization algorithm with the maximally mixed state, we can guarantee that the resulting matrix corresponds to a quantum channel.
- (5) Finally, we complete the rigorous error analysis for this proper hypothesis.

Note that Step 3 is the only step where our proof is incomplete. Namely, we identify this key mathematical condition (see Conjecture 1) but are unable to prove (or disprove) that it holds for our problem.

a. Optimization constraints for quantum channels

First, we derive the constraints that the Fourier coefficients of any quantum channel must satisfy. Recall that the Choi representation of a quantum channel can be written as

$$J(\Phi) = \sum_{x,y} \hat{\Phi}(x,y) J(\Phi_{x,y}), \quad (\text{D191})$$

where $\Phi_{x,y} \triangleq \sigma_x(\cdot)\sigma_y$. From Lemma A1, we have that

$$\text{tr}_1(J(\Phi)) = I. \quad (\text{D192})$$

We first show the effect of tracing out the first register on the Choi representations $J(\Phi_{x,y})$:

$$\text{tr}_1(J(\Phi_{x,y})) = \text{tr}_1 \left(\sum_{a,b} \sigma_x|a\rangle\langle b|\sigma_y \otimes |a\rangle\langle b| \right) \quad (\text{D193})$$

$$= \sum_{a,b} \text{tr}(\sigma_x|a\rangle\langle b|\sigma_y)|a\rangle\langle b| \quad (\text{D194})$$

$$= \sum_{a,b} (\langle b|\sigma_y\sigma_x|a\rangle)|a\rangle\langle b| \quad (\text{D195})$$

$$= (\sigma_y\sigma_x)^\top \quad (\text{D196})$$

$$= \sigma_x^\top \sigma_y^\top \quad (\text{D197})$$

$$= \sigma_x^* \sigma_y^*, \quad (\text{D198})$$

where in the last line we use the fact that all Pauli strings are Hermitian. It is easy to verify that complex conjugation leaves the Paulis I, X , and Z invariant, while $Y^* = -Y$. Thus, we see that $\sigma_x^* \sigma_y^*$ is equal to $\sigma_x \sigma_y$ up to a phase. We define the phase $\theta(x, y)$ by the following equation:

$$\sigma_x^* \sigma_y^* = e^{i\theta(x, y)} \sigma_x \sigma_y. \quad (\text{D199})$$

Further, we define the phase $\phi(x, y)$ by the following equation:

$$\sigma_x \sigma_y = e^{i\phi(x, y)} \sigma_z, \quad (\text{D200})$$

which is associated with a unique Pauli string z for all pairs x, y . Now, we can rewrite the partial trace of the Choi representation $J(\Phi)$ as follows:

$$I = \text{tr}_1(J(\Phi)) \quad (\text{D201})$$

$$= \sum_{x, y} \hat{\Phi}(x, y) \text{tr}_1(J(\Phi_{x, y})) \quad (\text{D202})$$

$$= \sum_{x, y} \hat{\Phi}(x, y) e^{i\theta(x, y)} \sigma_x \sigma_y. \quad (\text{D203})$$

By summing over the terms with $x = y$, we obtain the following constraint:

$$\sum_x \hat{\Phi}(x, x) = 1, \quad (\text{D204})$$

which can also be seen from Lemma A3. Now, for each non-identity Pauli string z , we obtain the following constraint:

$$\sum_{x, y} e^{i\theta(x, y)} e^{i\phi(x, y)} \hat{\Phi}(x, y) = 0. \quad (\text{D205})$$

Finally, from Lemma A1, we have

$$\hat{\Phi} \succcurlyeq 0. \quad (\text{D206})$$

Now, the original task of projecting $\hat{\Psi}$ reduces to the following optimization problem with respect to the matrix $\hat{\Phi}$.

$$\min_{\hat{\Phi}} \|\hat{\Psi} - \hat{\Phi}\|_F^2 \quad (\text{D207})$$

subject to the constraints given by Eqs. (D204) to (D206). For ease of notation, we will denote the dimensions of $\hat{\Psi}$ as $D \times D$ for the rest of this discussion and specify them for the concept classes we consider later.

b. Mapping to real, symmetric matrices

This problem can be seen as the complex, Hermitian generalization of a corresponding problem in convex optimization over real, symmetric matrices. We wish to map the problem to one with respect to real, symmetric matrices to leverage tools in convex optimization literature, which typically only deal with this case. To this end, note that any complex matrix A can be decomposed as a sum of its real and complex parts, $A \triangleq X + iY$, where X and Y are both real. Then, we can map A to the following real operator:

$$M = \begin{pmatrix} X & -Y \\ Y & X \end{pmatrix}. \quad (\text{D208})$$

This transformation is well-known, first noted in Ref. [145], and can be used to transform any complex semidefinite programming problem into its real analogue. It can be shown that this transformation preserves the properties of the original operator. For example, a Hermitian operator is mapped to a symmetric one and a Hermitian positive semi-definite operator is mapped to a symmetric positive semi-definite one. Using this transformation, we see that it suffices to solve the following convex optimization problem over reals:

$$\min_Q \|P - Q\|_F^2 \quad \text{s.t.} \quad Q \succcurlyeq 0, \text{tr}(A_i^\top Q) = b_i, \forall i \in [2D], \quad (\text{D209})$$

where $P, Q, A_i \in \mathbb{R}^{2D \times 2D}$ and $b \in \mathbb{R}^{2D}$, and P is a symmetric operator. Here, the A_i matrices and the b vector need to be chosen suitably to encompass our constraints in Eqs. (D204) to (D206) after the real mapping. It is clear that our problem can be written this way for a suitable choice of A_i, b . This is the problem of finding the nearest projection (with respect to the Frobenius norm) of a symmetric operator onto the intersection of the positive semidefinite cone and an affine subspace.

c. Convex optimization algorithms and convergence.

In particular, we wish to project onto the intersection of two convex sets, which we denote as C_1 and C_2 . Explicitly,

$$\begin{aligned} C_1 &\triangleq S_{2D}^+, \\ C_2 &\triangleq \{Q \in \mathbb{R}^{2D \times 2D} : \text{tr}(A_i^\top Q) = b_i, \forall i \in [2D]\}, \end{aligned} \quad (\text{D210})$$

where S_{2D}^+ denotes the positive semidefinite (PSD) cone of $2D \times 2D$ matrices. We can equivalently consider the vectorizations of these sets, where the vectorization of a matrix is defined by combining all columns of the matrix into a single vector. Vectorization can be inverted when the dimensions of the original matrix are known. We denote

the vectorization operation as vec and its inverse as vec^{-1} . Then, the sets C_1 and C_2 can be written as

$$\begin{aligned} C_1 &= \{x \in \mathbb{R}^{4D^2} : \text{vec}^{-1}(x) \in S_{2D}^+\}, \\ C_2 &= \{x \in \mathbb{R}^{4D^2} : \tilde{A}x = \tilde{b}\}, \end{aligned} \quad (\text{D211})$$

where we use $\tilde{A}$ and $\tilde{b}$ to denote the constraints for the vectorized version of Q [i.e., x should be thought of as $\text{vec}(Q)$ for C_2].

Note that projecting onto either C_1 or C_2 individually is easy [running in time $\text{poly}(D)$] since they are simply the PSD cone and a polyhedral set (see, e.g., Ref. [146]). Note that the runtime being in terms of D is not a problem for our applications, as we discuss at the end of this section. Thus, the problem is as follows: how can we use projections onto C_1 and C_2 individually to project onto $C_1 \cap C_2$? This problem is well-studied in convex optimization, where researchers have developed methods such as alternating projections [147,148] and Dykstra's algorithm [149]. As the name suggests, at a high level, these algorithms work by projecting the input onto C_1 and C_2 in turn. Classical results in optimization theory show that Dykstra's algorithm converges to the desired projection [149–153], and its variants are also widely studied [154–158]; however, researchers have only recently rigorously analyzed its convergence *rate* for non-polyhedral sets [109]. Convergence rates have previously been analyzed for the basic alternating projections algorithm [159–161] and Dykstra's algorithm for projection onto the intersection of polyhedral sets [155,162,163]. However, the former only guarantees convergence to a point in $C_1 \cap C_2$, rather than the closest one, which would not help for our purposes since we require an error bound. Also, the latter does not apply to our setting, as $C_1 = S_{2D}^+$ is clearly not polyhedral. In our setting, using Dykstra's algorithm, one would hope for linear convergence, as this leads to a computationally efficient (in the size of the input matrix) projection algorithm. In fact, Ref. [109] shows that Dykstra's algorithm achieves a linear convergence rate under certain conditions. Namely, they prove the following theorem.

Theorem D6 (Theorem 5.3(i) in Ref. [109]; simplified). Consider the optimization problem

$$\min_{x \in \mathbb{R}^n} \frac{1}{2} \|x - x_0\|_2^2 \quad \text{s.t.} \quad x \in \bigcap_{i=1}^{\ell} C_i, \quad (\text{D212})$$

where ℓ is the number of convex sets C_i and $x_0 \in \mathbb{R}^n$ is the initial point to be projected. Suppose the following three conditions hold.

- (1) Each C_i is a $C^{1,1}$ -cone reducible closed convex set.
- (2) $\bigcap_{i=1}^{\ell} \text{ri}(C_i) \neq \emptyset$, where ri denotes the relative interior.

- (3) $0 \in x^* - x_0 + \text{ri}\left(\partial\left(\sum_{i=1}^{\ell} \delta_{C_i}\right)(x^*)\right)$, where x^* is the optimal projected point, ∂ denotes the subdifferential, and δ_{C_i} is the indicator function of the set C_i (0 for inputs inside C_i and $+\infty$ otherwise).

Then, Dykstra's algorithm produces iterates that linearly converge to the optimal point x^* .

Note that Ref. [109] considers a more general problem, in which they wish to project onto a point such that a linear transformation of this point lies in the intersection of the convex sets. For our setting, we can take this transformation to be the identity, in which case the algorithm from Ref. [109] reduces to Dykstra's algorithm. Thus, this theorem would give the convergence result we need if its conditions hold.

Now, let us discuss each of the conditions. First, for Condition 1, the exact definition of $C^{1,\alpha}$ -cone reducibility, for $\alpha \in (0, 1]$, is outside of the scope of this work, and we refer to Definition 4.1 in Ref. [109] for the definition. Importantly, this is an extension of the more common notion of C^2 -cone reducibility [164,165]. Specifically, Ref. [109] states that any C^2 -cone reducible set is also $C^{1,\alpha}$ -cone reducible for all $\alpha \in (0, 1]$. Moreover, Ref. [164] shows that all polyhedrons and the cone of positive semidefinite matrices are both C^2 -cone reducible. Hence, our convex sets C_1 and C_2 are both C^2 -cone reducible and hence $C^{1,\alpha}$ -cone reducible for all $\alpha \in (0, 1]$. In particular, C_1 and C_2 are both $C^{1,1}$ -cone reducible, satisfying Condition 1.

For Condition 2, this is a standard assumption in optimization theory, which simply says that there is a strictly feasible point. This is satisfied for our setting.

Finally, we turn to Condition 3. Let us define some of the notation more properly. Here, δ_{C_i} is the indicator function of the set C_i :

$$\delta_{C_i}(x) \triangleq \begin{cases} 0 & x \in C_i \\ +\infty & x \notin C_i \end{cases}. \quad (\text{D213})$$

Also, ∂ denotes the subdifferential. For a function $f : \mathbb{R}^{4D^2} \rightarrow \mathbb{R}$, the subdifferential is defined as

$$\begin{aligned} \partial f(x) &\triangleq \{z \in \mathbb{R}^{4D^2} : f(y) \\ &\geq f(x) + (y - x)^T z \quad \forall y \in \mathbb{R}^{4D^2}\}. \end{aligned} \quad (\text{D214})$$

Finally, the relative interior for a set $S \subseteq \mathbb{R}^{4D^2}$ is defined as

$$\text{ri}(S) \triangleq \{x \in S : \forall y \in S, \exists \lambda > 1 \text{ s.t. } \lambda x + (1 - \lambda)y \in S\}. \quad (\text{D215})$$

Because of Condition 2, standard results in subdifferential calculus (see, e.g., Ref. [166]) state that

$$\partial(\delta_{C_1} + \delta_{C_2}) = \partial\delta_{C_1} + \partial\delta_{C_2}. \quad (\text{D216})$$

Thus, we need to compute the subdifferentials of the indicator functions of our convex sets C_1 and C_2 . It is well-known in convex optimization that the subdifferential of an indicator function of a set S is the normal cone of $S \subseteq \mathbb{R}^{4D^2}$, where the normal cone is defined as

$$N_S(y) \triangleq \{a \in \mathbb{R}^{4D^2} : a^\top y \geq a^\top x, \quad \forall x \in S\}. \quad (\text{D217})$$

The normal cone of the cone of positive semidefinite matrices is known to be

$$\delta_{C_1}(X) = N_{C_1}(X) = \{C \in \mathbb{R}^{2D \times 2D} : C \preceq 0, \text{tr}(CX) = 0\}, \quad (\text{D218})$$

where here $X \in \mathbb{R}^{2D \times 2D}$ are matrices. Writing in terms of vectors using vectorization, we can write

$$\delta_{C_1}(x) = \{c \in \mathbb{R}^{4D^2} : \text{vec}^{-1}(c) \preceq 0, \text{tr}(\text{vec}^{-1}(c)\text{vec}^{-1}(x)) = 0\}, \quad (\text{D219})$$

where $x \in \mathbb{R}^{4D^2}$. Moreover, we can compute the normal cone of the polyhedral set $C_2 = \{x \in \mathbb{R}^{4D^2} : \tilde{A}x = \tilde{b}\}$:

$$\delta_{C_2}(x) = N_{C_2}(x) \quad (\text{D220})$$

$$= \{g : g^\top(x - y) \geq 0, \quad \forall y \text{ s.t. } \tilde{A}y = \tilde{b}\} \quad (\text{D221})$$

$$= \{g : g^\top z \geq 0, \quad \forall z \in \ker(\tilde{A})\} \quad (\text{D222})$$

$$= (\ker(\tilde{A}))^\perp \quad (\text{D223})$$

$$= \text{im}(\tilde{A}^\top). \quad (\text{D224})$$

Here, the second line is the definition of the normal cone. The third line follows because $x - y$ is in the kernel of the matrix $\tilde{A}$ since both $x, y \in C_2$. The fifth line follows by the standard result in linear algebra that the orthogonal complement of the kernel space is the image of the transpose.

Now, putting everything together, Condition 3 is as follows.

Conjecture 1. Let $x_0 \in \mathbb{R}^{4D^2}$ be the vectorization of the real-mapped matrix of Fourier coefficients that we wish to project onto that of a valid quantum channel. Let $x^* \in \mathbb{R}^{4D^2}$ denote the vectorization of the real-mapped matrix of Fourier coefficients of the closest valid quantum channel to x_0 in 2-norm. Then, it remains to show the following

condition holds.

$$x_0 - x^* \in \text{ri} \left(\left\{ c + \tilde{A}^\top y : c, y \in \mathbb{R}^{4D^2}, \text{vec}^{-1}(c) \preceq 0, \text{tr}(\text{vec}^{-1}(c)\text{vec}^{-1}(x^*)) = 0 \right\} \right), \quad (\text{D225})$$

where ri denotes the relative interior.

We are unsure if this condition holds (or does not hold) in our setting, but suppose for the rest of this discussion that it does. Now, by Theorem D6, Dykstra's algorithm converges linearly to the point x^* , which is the vectorization of the matrix of Fourier coefficients (after being mapped to a real, symmetric matrix) of the closest valid quantum channel to x_0 . Recall that Dykstra's algorithm works by repeatedly projecting alternately onto C_1 and then C_2 (with a correction term). By linear convergence, this converges to some point ϵ_2 -close in 2-norm to x^* in $\mathcal{O}(\log(1/\epsilon_2))$ iterations, for some $\epsilon_2 > 0$. Moreover, suppose without loss of generality that the final projection applied is the one projecting onto C_2 . We will make use of this property in the next step.

Mapping the output of the algorithm back to our original problem for Fourier coefficients of channels gives us a Fourier coefficient matrix $\hat{\Psi}_2$. Then, we obtain the following guarantee for $\hat{\Psi}_2$.

$$\frac{\|\hat{\Psi}_2 - \hat{\Psi}^*\|_F}{2} \leq \epsilon_2, \quad (\text{D226})$$

where

$$\hat{\Psi}^* = \arg \min_{\Psi \in \text{CPTP}} \|\hat{\Psi}_1 - \hat{\Psi}\|_F. \quad (\text{D227})$$

An immediate application of the triangle inequality for the Frobenius norm then gives us

$$\frac{\|\hat{\Psi}_1 - \hat{\Psi}_2\|_F}{2} \leq \frac{\|\hat{\Psi}_1 - \hat{\Psi}\|_F}{2} + \epsilon_2, \quad \forall \text{ CPTP maps } \Psi. \quad (\text{D228})$$

d. Mixing to guarantee that the output is a quantum channel

Even though $\hat{\Psi}_2$ has desirable distance properties, as it does not achieve *exact* convergence to the projection, it may not correspond to a CPTP map. However, as our final projection was onto the set C_2 , $\hat{\Psi}_2$ is Hermitian and satisfies the affine constraints. Thus, the only remaining condition necessary to obtain a proper hypothesis is to obtain a positive semidefinite hypothesis. To ensure that our final hypothesis is PSD while still retaining the other properties, we take inspiration from Ref. [167] and mix it

with the maximally mixed state. In particular, for a suitably chosen $p > 0$, our output is

$$\hat{\Psi}_3 = (1 - p)\hat{\Psi}_2 + p\frac{I}{D}, \quad (\text{D229})$$

where I is the $D \times D$ identity operator. To understand the intuition behind mixing, we first note that the maximally mixed state can be viewed as the Fourier coefficient matrix of a CPTP map that applies a Pauli operator uniformly at random. Thus, it must satisfy all required constraints. Then, we note that if p is taken large enough, $\hat{\Psi}_3$ also satisfies the constraints. This can be seen by noting that the convex combination of two points satisfying the affine constraints will also satisfy them. For the positive semidefiniteness, we note that $\hat{\Psi}_2$ is Hermitian, and we only need to take p large enough to make the most negative eigenvalue of $\hat{\Psi}_2$ map to 0 in the eigenspectrum of $\hat{\Psi}_3$. Using the fact that $\hat{\Psi}_2$ is not too far from a PSD matrix [Eq. (D226)], we will then show that p does not need to be too large.

First, suppose that $\hat{\Psi}_2$ does have some negative eigenvalue. Otherwise, our constraints are already satisfied, and we do not need to proceed further. Let the negative eigenvalue of $\hat{\Psi}_2$ with the largest absolute value be λ^* . Clearly, by Eq. (D229), we need to take p large enough such that

$$(1 - p)\lambda^* + p/D \geq 0. \quad (\text{D230})$$

This can be rewritten as

$$p \geq \frac{D|\lambda^*|}{1 + D|\lambda^*|}. \quad (\text{D231})$$

We can now upper bound $|\lambda^*|$ as follows. First, recall that projection of a Hermitian operator onto the positive semidefinite cone is well-known (see, e.g., Ref. [146]). In particular, this projection works by taking the positive component of the Hermitian operator's eigenspectrum and setting the negative eigenvalues to 0. For a Hermitian A , we denote this by $(A)^+$. Then, we have

$$\epsilon_2 \geq \frac{\|\hat{\Psi}_2 - \hat{\Psi}^*\|_F}{2} \quad (\text{D232})$$

$$\geq \frac{\|\hat{\Psi}_2 - (\hat{\Psi}_2)^+\|_F}{2} \quad (\text{D233})$$

$$= \frac{\sqrt{\sum_{\lambda_i < 0} \lambda_i^2}}{2} \quad (\text{D234})$$

$$\geq \frac{|\lambda^*|}{2} \quad (\text{D235})$$

where the first line follows from Eq. (D226). The second line follows from the fact that $\hat{\Psi}^*$ is PSD and $(\hat{\Psi}_2)^+$ is the nearest PSD operator to $\hat{\Psi}_2$. In the third line, we use

λ_i to denote the eigenvalues of $\hat{\Psi}_2$ and only sum over the negative ones. In the final line, we ignore all except the negative eigenvalue with the largest absolute value. Thus, we set

$$p = \frac{2D\epsilon_2}{1 + 2D\epsilon_2}, \quad (\text{D236})$$

which satisfies Eq. (D231).

e. Error analysis

Finally, it remains to bound the error between $\hat{\Psi}_2$ and $\hat{\Psi}_3$.

$$\frac{\|\hat{\Psi}_2 - \hat{\Psi}_3\|_F}{2} = \frac{\|p(\hat{\Psi}_2 - I/D)\|_F}{2} \quad (\text{D237})$$

$$= \frac{p}{2} \left\| \hat{\Psi}_2 - \frac{I}{D} \right\|_F \quad (\text{D238})$$

$$\leq \frac{p}{2} \left(\|\hat{\Psi}_2\|_F + \left\| \frac{I}{D} \right\|_F \right) \quad (\text{D239})$$

$$\leq \frac{p}{2} \left(\|\hat{\Psi}_2 - \hat{\Psi}^*\|_F + \|\hat{\Psi}^*\|_F + \left\| \frac{I}{D} \right\|_F \right) \quad (\text{D240})$$

$$\leq \frac{p}{2} \left(2\epsilon_2 + 1 + \frac{1}{\sqrt{D}} \right) \quad (\text{D241})$$

$$\leq 2p \quad (\text{D242})$$

$$= \frac{4D\epsilon_2}{1 + 2D\epsilon_2}, \quad (\text{D243})$$

where in the first line we use Eq. (D229). In the third and fourth lines, we use the triangle inequality for the Frobenius norm. In the fifth line, we use Eq. (D226), the fact that $\|\hat{\Psi}^*\|_F \leq 1$ for a CPTP map 1, and we compute the Frobenius norm of the identity operator. In the sixth line, we take a loose upper bound on these quantities, and finally, in the last line we use Eq. (D236). Finally, we need this error to be at most $\mathcal{O}(\epsilon_1)$. Then, by taking $\epsilon_2 \leq \epsilon_1/4D$, we obtain

$$\frac{\|\hat{\Psi}_2 - \hat{\Psi}_3\|_F}{2} \leq \epsilon_1. \quad (\text{D244})$$

Then, using the triangle inequality for d_F and combining Eqs. (D190), (D228), and (D244), we obtain

$$d_F(\Phi, \Psi_3) \leq 2 \cdot \text{opt}(\Phi, \mathcal{C}) + 3\epsilon_1, \quad (\text{D245})$$

which is the desired bound, up to constant factors for ϵ_1 . Here, we implicitly make use of the fact that $d_F(\Phi, \Psi) = (\|\hat{\Phi} - \hat{\Psi}\|_F)/2$ for any quantum channels Φ, Ψ , which can be seen from Lemma A4.

We can now bound the time complexity of this method. In order to mix the Fourier coefficient matrix, we need

to compute the smallest eigenvalue of $\hat{\Psi}_2$. We also require making alternating projections. Each of these steps can be done in time $\text{poly}(D)$. Moreover, due to linear convergence, each projection needs to be carried out $\mathcal{O}(\log(1/\epsilon_2)) = \mathcal{O}(\log(D/\epsilon_1))$ times. This gives us a final time complexity of $\text{poly}(D) \log(1/\epsilon_1)$. Again, this analysis only holds assuming that Conjecture 1 holds.

For k -junta channels, the principle submatrix that we optimize over would be of dimension $4^k \times 4^k$, giving us a time complexity $\text{poly}(4^k) \log(1/\epsilon_1)$. Similarly, the dimension for low-degree channels would be at most $n^k 4^k \times n^k 4^k$, giving us a time complexity $\text{poly}(n^k 4^k) \log(1/\epsilon_1)$.

Thus, assuming Conjecture 1 holds, we can efficiently obtain proper 2-agnostic process tomography of junta channels and low-degree channels by applying this method to the outputs of Theorems D3 and D4 respectively.

-
- [1] M. Mohseni, A. T. Rezakhani, and D. A. Lidar, Quantum-process tomography: Resource analysis of different strategies, *Phys. Rev. A—At., Mol. Opt. Phys.* **77**, 032322 (2008).
 - [2] J. L. O’Brien, G. J. Pryde, A. Gilchrist, D. F. V. James, N. K. Langford, T. C. Ralph, and A. G. White, Quantum process tomography of a controlled-not gate, *Phys. Rev. Lett.* **93**, 080502 (2004).
 - [3] A. J. Scott, Optimizing quantum process tomography with unitary 2-designs, *J. Phys. A: Math. Theor.* **41**, 055308 (2008).
 - [4] I. L. Chuang, and M. A. Nielsen, Prescription for experimental determination of the dynamics of a quantum black box, *J. Mod. Opt.* **44**, 2455 (1997).
 - [5] G. M. D’Ariano, and P. Lo Presti, Quantum tomography for measuring experimentally the matrix elements of an arbitrary quantum operation, *Phys. Rev. Lett.* **86**, 4195 (2001).
 - [6] J. Haah, R. Kothari, R. O’Donnell, and E. Tang, in *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, 2023), p. 363.
 - [7] J. F. Poyatos, J. I. Cirac, and P. Zoller, Complete characterization of a quantum process: The two-bit quantum gate, *Phys. Rev. Lett.* **78**, 390 (1997).
 - [8] R. Levy, D. Luo, and B. K. Clark, Classical shadows for quantum process tomography on near-term quantum computers, *Phys. Rev. Res.* **6**, 013029 (2024).
 - [9] H.-Y. Huang, S. T. Flammia, and J. Preskill, Foundations for learning from noisy quantum experiments, *ArXiv:2204.13691*.
 - [10] 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—At., Mol. Opt. Phys.* **87**, 062119 (2013).
 - [11] 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).
 - [12] S. Endo, Z. Cai, S. C. Benjamin, and X. Yuan, Hybrid quantum-classical algorithms and quantum error mitigation, *J. Phys. Soc. Jpn.* **90**, 032001 (2021).
 - [13] A. Seif, Z.-P. Ciani, S. Zhou, S. Chen, and L. Jiang, Shadow distillation: Quantum error mitigation with classical shadows for near-term quantum processors, *PRX Quantum* **4**, 010303 (2023).
 - [14] 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, *Rev. Mod. Phys.* **95**, 045005 (2023).
 - [15] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum metrology, *Phys. Rev. Lett.* **96**, 010401 (2006).
 - [16] V. Giovannetti, S. Lloyd, and L. Maccone, Advances in quantum metrology, *Nat. Photonics* **5**, 222 (2011).
 - [17] D. Kim, Z. Shi, C. Simmons, D. Ward, J. Prance, T. S. Koh, J. K. Gamble, D. Savage, M. Lagally, M. Friesen *et al.*, Quantum control and process tomography of a semiconductor quantum dot hybrid qubit, *Nature* **511**, 70 (2014).
 - [18] M. Arapinis, M. Delavar, M. Doosti, and E. Kashefi, Quantum physical unclonable functions: Possibilities and impossibilities, *Quantum* **5**, 475 (2021).
 - [19] T. Schuster, J. Haferkamp, and H.-Y. Huang, Random unitaries in extremely low depth, *Science* **389**, 92 (2025).
 - [20] R. A. Low, Learning and testing algorithms for the Clifford group, *Phys. Rev. A—At., Mol. Opt. Phys.* **80**, 052314 (2009).
 - [21] C.-Y. Lai, and H.-C. Cheng, Learning quantum circuits of some T gates, *IEEE Trans. Inf. Theory* **68**, 3951 (2022).
 - [22] S. Nadimpalli, N. Parham, F. Vasconcelos, and H. Yuen, in *Proceedings of the 56th Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2024), p. 1498.
 - [23] H.-Y. Huang, Y. Liu, M. Broughton, I. Kim, A. Anshu, Z. Landau, and J. R. McClean, Learning shallow quantum circuits, in *Proceedings of the 56th Annual ACM Symposium on Theory of Computing* (STOC 2024).
 - [24] T. Chen, S. Nadimpalli, and H. Yuen, in *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (SIAM, 2023), p. 1163.
 - [25] Z. Bao, and P. Yao, in *The Thirty Sixth Annual Conference on Learning Theory* (PMLR, 2023), p. 1064.
 - [26] 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.* **14**, 2858 (2023).
 - [27] H. Zhao, L. Lewis, I. Kannan, Y. Quek, H.-Y. Huang, and M. C. Caro, Learning quantum states and unitaries of bounded gate complexity, *PRX Quantum* **5**, 040306 (2024).
 - [28] A. Montanaro, and T. J. Osborne, Quantum boolean functions, *Chicago Journal OF Theoretical Computer Science* **16**, 1 (2010).
 - [29] C. Rouzé, M. Wirth, and H. Zhang, Quantum Talagrand, KKL and Friedgut’s theorems and the learnability of quantum boolean functions, *Commun. Math. Phys.* **405**, 95 (2024).
 - [30] S. Arunachalam, A. Dutt, F. Escudero Gutiérrez, and C. Palazuelos, in *51st International Colloquium on Automata, Languages, and Programming (ICALP*

- 2024) (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2024), p. 13.
- [31] R. Harper, S. T. Flammia, and J. J. Wallman, Efficient learning of quantum noise, *Nat. Phys.* **16**, 1184 (2020).
 - [32] R. Harper, W. Yu, and S. T. Flammia, Fast estimation of sparse quantum noise, *PRX Quantum* **2**, 010322 (2021).
 - [33] S. T. Flammia, and J. J. Wallman, Efficient estimation of Pauli channels, *ACM Trans. Quantum Comput.* **1**, 1 (2020).
 - [34] S. Chen, S. Zhou, A. Seif, and L. Jiang, Quantum advantages for Pauli channel estimation, *Phys. Rev. A* **105**, 032435 (2022).
 - [35] S. T. Flammia, and R. O’Donnell, Pauli error estimation via population recovery, *Quantum* **5**, 549 (2021).
 - [36] O. Fawzi, A. Oufkir, and D. S. França, Lower bounds on learning Pauli channels, *IEEE Trans. Inf. Theory* **71**, 2642 (2025).
 - [37] S. Chen, C. Oh, S. Zhou, H.-Y. Huang, and L. Jiang, Tight bounds on Pauli channel learning without entanglement, *Phys. Rev. Lett.* **132**, 180805 (2024).
 - [38] J. Kunjummen, M. C. Tran, D. Carney, and J. M. Taylor, Shadow process tomography of quantum channels, *Phys. Rev. A* **107**, 042403 (2023).
 - [39] H.-Y. Huang, S. Chen, and J. Preskill, Learning to predict arbitrary quantum processes, *PRX Quantum* **4**, 040337 (2023).
 - [40] S. Chen, J. de Dios Pont, J. Hsieh, H. Huang, J. Lange, and J. Li, in *Proceedings of Thirty Eighth Conference on Learning Theory* (Proceedings of Machine Learning Research, 2025), Vol. 291 p. 986, <https://proceedings.mlr.press/v291/chen25c.html>.
 - [41] K.-M. Chung, and H.-H. Lin, in *16th Conference on the Theory of Quantum Computation, Communication and Cryptography* (TQC 2021), Leibniz International Proceedings in Informatics (LIPIcs), Vol. 197, pp. 3:1–3:22, Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2021).
 - [42] S. Grewal, V. Iyer, W. Kretschmer, and D. Liang, Agnostic tomography of stabilizer product states, *ArXiv:2404.03813*.
 - [43] S. Chen, W. Gong, Q. Ye, and Z. Zhang, in *Proceedings of the 57th Annual ACM Symposium on Theory of Computing* (2025), p. 429.
 - [44] A. Anshu, and S. Arunachalam, A survey on the complexity of learning quantum states, *Nat. Rev. Phys.* **6**, 59 (2024).
 - [45] C. Bădescu, and R. O’Donnell, in *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2021), p. 1398.
 - [46] S. Grewal, V. Iyer, W. Kretschmer, and D. Liang, in *Proceedings of the 56th Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2024), p. 1352.
 - [47] M. J. Kearns, R. E. Schapire, and L. M. Sellie, in *Proceedings of the fifth annual workshop on Computational learning theory* (Association for Computing Machinery, New York, NY, USA, 1992), p. 341.
 - [48] D. Haussler, Decision theoretic generalizations of the PAC model for neural net and other learning applications, *Inf. Comput.* **100**, 78 (1992).
 - [49] M. C. Caro, M. Hinsche, M. Ioannou, A. Nietner, and R. Sweke, in *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)* (Schloss-Dagstuhl-Leibniz Zentrum für Informatik, Dagstuhl, Germany, 2024).
 - [50] A. Angrisani, Learning unitaries with quantum statistical queries, *Quantum* **9**, 1817 (2025).
 - [51] C. Wadhwa, and M. Doosti, Learning quantum processes with quantum statistical queries, *Quantum* **9**, 1739 (2025).
 - [52] C. Wadhwa, and M. Doosti, Noise-tolerant learnability of shallow quantum circuits from statistics and the cost of quantum pseudorandomness, *ArXiv:2405.12085*.
 - [53] M. Hinsche, M. Ioannou, A. Nietner, J. Haferkamp, Y. Quek, D. Hangleiter, J. P. Seifert, J. Eisert, and R. Sweke, One T gate makes distribution learning hard, *Phys. Rev. Lett.* **130**, 240602 (2023).
 - [54] A. Nietner, M. Ioannou, R. Sweke, R. Kueng, J. Eisert, M. Hinsche, and J. Haferkamp, On the average-case complexity of learning output distributions of quantum circuits, *Quantum* **9**, 1883 (2025).
 - [55] A. Nietner, Unifying (quantum) statistical and parametrized (quantum) algorithms, *ArXiv:2310.17716*.
 - [56] A. Montanaro, and R. de Wolf, A survey of quantum property testing, *Theory Comput. Library* **7**, 1 (2016).
 - [57] G. Wang, Property testing of unitary operators, *Phys. Rev. A–At., Mol. Opt. Phys.* **84**, 052328 (2011).
 - [58] S. Aaronson, and D. Gottesman, Improved simulation of stabilizer circuits, *Phys. Rev. A–At., Mol. Opt. Phys.* **70**, 052328 (2004).
 - [59] D. Gottesman, Theory of fault-tolerant quantum computation, *Phys. Rev. A* **57**, 127 (1998).
 - [60] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio *et al.*, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
 - [61] A. M. Childs, and N. Wiebe, Hamiltonian simulation using linear combinations of unitary operations, *Quantum Inf. Comput.* **12**, 901 (2012).
 - [62] J. Heredge, M. West, L. Hollenberg, and M. Sevir, Non-unitary quantum machine learning, *Phys. Rev. Appl.* **23**, 044046 (2025).
 - [63] B. Coyle, S. Raj, N. Mathur, S. Kazdaghi, and I. Kerendi, Training-efficient density quantum machine learning, *npj Quantum Inf.* **11**, 172 (2025).
 - [64] S. Arunachalam, A. B. Grilo, and H. Yuen, Quantum statistical query learning, *ArXiv:2002.08240*.
 - [65] S. Arunachalam, V. Havlicek, and L. Schatzki, On the role of entanglement and statistics in learning, *Adv. Neural Inf. Process. Syst.* **36**, 55064 (2024).
 - [66] W. Górecki, A. Riccardi, and L. Maccone, Quantum metrology of noisy spreading channels, *Phys. Rev. Lett.* **129**, 240503 (2022).
 - [67] R. Demkowicz-Dobrzański, J. Kołodyński, and M. Guţă, The elusive Heisenberg limit in quantum-enhanced metrology, *Nat. Commun.* **3**, 1063 (2012).
 - [68] B. Escher, R. L. de Matos Filho, and L. Davidovich, General framework for estimating the ultimate precision limit in noisy quantum-enhanced metrology, *Nat. Phys.* **7**, 406 (2011).

- [69] Z. Ji, G. Wang, R. Duan, Y. Feng, and M. Ying, Parameter estimation of quantum channels, *IEEE Trans. Inf. Theory* **54**, 5172 (2008).
- [70] S. Zhou, M. Zhang, J. Preskill, and L. Jiang, Achieving the Heisenberg limit in quantum metrology using quantum error correction, *Nat. Commun.* **9**, 78 (2018).
- [71] R. Demkowicz-Dobrzański, J. Czakowski, and P. Sekatski, Adaptive quantum metrology under general Markovian noise, *Phys. Rev. X* **7**, 041009 (2017).
- [72] S. Zhou, A. G. Manes, and L. Jiang, Achieving metrological limits using ancilla-free quantum error-correcting codes, *Phys. Rev. A* **109**, 042406 (2024).
- [73] T. Ono, and H. F. Hofmann, Effects of photon losses on phase estimation near the Heisenberg limit using coherent light and squeezed vacuum, *Phys. Rev. A-At., Mol. Opt. Phys.* **81**, 033819 (2010).
- [74] G. Gilbert, M. Hamrick, and Y. S. Weinstein, Use of maximally entangled n-photon states for practical quantum interferometry, *JOSA B* **25**, 1336 (2008).
- [75] M. A. Rubin, and S. Kaushik, Loss-induced limits to phase measurement precision with maximally entangled states, *Phys. Rev. A-At., Mol. Opt. Phys.* **75**, 053805 (2007).
- [76] N. Shettell, E. Kashefi, and D. Markham, Cryptographic approach to quantum metrology, *Phys. Rev. A* **105**, L010401 (2022).
- [77] Z. Huang, C. Macchiavello, and L. Maccone, Cryptographic quantum metrology, *Phys. Rev. A* **99**, 022314 (2019).
- [78] S. Zhou, Limits of noisy quantum metrology with restricted quantum controls, *Phys. Rev. Lett.* **133**, 170801 (2024).
- [79] P. Czarnik, A. Arrasmith, P. J. Coles, and L. Cincio, Error mitigation with Clifford quantum-circuit data, *Quantum* **5**, 592 (2021).
- [80] A. Strikis, D. Qin, Y. Chen, S. C. Benjamin, and Y. Li, Learning-based quantum error mitigation, *PRX Quantum* **2**, 040330 (2021).
- [81] E. Van Den Berg, Z. K. Mineev, A. Kandala, and K. Temme, Probabilistic error cancellation with sparse Pauli-Lindblad models on noisy quantum processors, *Nat. Phys.* **19**, 1116 (2023).
- [82] R. S. Gupta, E. Van Den Berg, M. Takita, D. Ristè, K. Temme, and A. Kandala, Probabilistic error cancellation for dynamic quantum circuits, *Phys. Rev. A* **109**, 062617 (2024).
- [83] M. R. Geller, and Z. Zhou, Efficient error models for fault-tolerant architectures and the Pauli twirling approximation, *Phys. Rev. A-At., Mol. Opt. Phys.* **88**, 012314 (2013).
- [84] J. J. Wallman, and J. Emerson, Noise tailoring for scalable quantum computation via randomized compiling, *Phys. Rev. A* **94**, 052325 (2016).
- [85] S. Ferracin, A. Hashim, J.-L. Ville, R. Naik, A. Carignan-Dugas, H. Qassim, A. Morvan, D. I. Santiago, I. Siddiqi, and J. J. Wallman, Efficiently improving the performance of noisy quantum computers, *Quantum* **8**, 1410 (2024).
- [86] C. Moore, Quantum circuits: Fanout, parity, and counting, *arXiv:quant-ph/9903046*.
- [87] D. W. C. Leung, *Towards Robust Quantum Computation* (Stanford University, 2000).
- [88] A. Acín, E. Jané, and G. Vidal, Optimal estimation of quantum dynamics, *Phys. Rev. A* **64**, 050302(R) (2001).
- [89] A. Peres, and P. Scudo, Covariant quantum measurements may not be optimal, *J. Mod. Opt.* **49**, 1235 (2002).
- [90] M. Hayashi, Parallel treatment of estimation of SU (2) and phase estimation, *Phys. Lett. A* **354**, 183 (2006).
- [91] G. Chiribella, G. M. D’ariano, and M. F. Sacchi, Optimal estimation of group transformations using entanglement, *Phys. Rev. A-At., Mol. Opt. Phys.* **72**, 042338 (2005).
- [92] J. Kahn, Fast rate estimation of a unitary operation in SU (d), *Phys. Rev. A-At., Mol. Opt. Phys.* **75**, 022326 (2007).
- [93] Y. Yang, R. Renner, and G. Chiribella, Optimal universal programming of unitary gates, *Phys. Rev. Lett.* **125**, 210501 (2020).
- [94] S. Aaronson, in *Proceedings of the 50th annual ACM SIGACT symposium on theory of computing* (Association for Computing Machinery, New York, NY, USA, 2018), p. 325.
- [95] L. Clinton, J. Bausch, and T. Cubitt, Hamiltonian simulation algorithms for near-term quantum hardware, *Nat. Commun.* **12**, 4989 (2021).
- [96] M. Schmitt, M. M. Rams, J. Dziarmaga, M. Heyl, and W. H. Zurek, Quantum phase transition dynamics in the two-dimensional transverse-field Ising model, *Sci. Adv.* **8**, eabl6850 (2022).
- [97] Y. Yanay, J. Braumüller, S. Gustavsson, W. D. Oliver, and C. Tahan, Two-dimensional hard-core Bose-Hubbard model with superconducting qubits, *Npj Quantum Inf.* **6**, 58 (2020).
- [98] A. Chatterjee, S. Rappaport, A. Giri, S. Johri, T. Proctor, D. E. B. Neira, P. Sathe, and T. Lubinski, A comprehensive cross-model framework for benchmarking the performance of quantum hamiltonian simulations, *IEEE Trans. Quantum Eng.* (2025).
- [99] J. Romero, J. P. Olson, and A. Aspuru-Guzik, Quantum autoencoders for efficient compression of quantum data, *Quantum Sci. Technol.* **2**, 045001 (2017).
- [100] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature* **567**, 209 (2019).
- [101] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, The power of quantum neural networks, *Nat. Comput. Sci.* **1**, 403 (2021).
- [102] R. O’Donnell, *Analysis of Boolean Functions* (Cambridge University Press, USA, 2014).
- [103] O. Goldreich, and L. A. Levin, in *Proceedings of the twenty-first annual ACM symposium on Theory of computing* (Association for Computing Machinery, New York, NY, USA, 1989), p. 25.
- [104] E. Kushilevitz, and Y. Mansour, in *Proceedings of the twenty-third annual ACM symposium on Theory of computing* (Association for Computing Machinery, New York, NY, USA, 1991), p. 455.
- [105] N. Linial, Y. Mansour, and N. Nisan, Constant depth circuits, Fourier transform, and learnability, *J. ACM (JACM)* **40**, 607 (1993).
- [106] M. Heidari, and W. Szpankowski, in *International Conference on Artificial Intelligence and Statistics* (PMLR, 2023), p. 2922.

- [107] S. Goldwasser, G. N. Rothblum, J. Shafer, and A. Yehudayoff, in *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)* (Schloss-Dagstuhl-Leibniz Zentrum für Informatik, Dagstuhl, Germany, 2021).
- [108] T. Gur, M. M. Jahanara, M. M. Khodabandeh, N. Rajgopal, B. Salamatian, and I. Shinkar, in *Proceedings of the 56th Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2024), p. 1063.
- [109] X. Wang, and T. K. Pong, Convergence rate analysis of a Dykstra-type projection algorithm, *SIAM J. Optim.* **34**, 563 (2024).
- [110] Note that ideas from the quasiprobability decomposition and probabilistic error cancellation which approximate linear superoperators by linear combinations of quantum channels [81,137–144] do not help because we need to exactly output a quantum channel.
- [111] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [112] A. Erhard, J. J. Wallman, L. Postler, M. Meth, R. Stricker, E. A. Martinez, P. Schindler, T. Monz, J. Emerson, and R. Blatt, Characterizing large-scale quantum computers via cycle benchmarking, *Nat. Commun.* **10**, 5347 (2019).
- [113] A. Carignan-Dugas, D. Dahlen, I. Hincks, E. Ospadov, S. J. Beale, S. Ferracin, J. Skanes-Norman, J. Emerson, and J. J. Wallman, The error reconstruction and compiled calibration of quantum computing cycles, *ArXiv:2303.17714*.
- [114] D. K. Tuckett, S. D. Bartlett, and S. T. Flammia, Ultrahigh error threshold for surface codes with biased noise, *Phys. Rev. Lett.* **120**, 050505 (2018).
- [115] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. Van Den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme *et al.*, Evidence for the utility of quantum computing before fault tolerance, *Nature* **618**, 500 (2023).
- [116] A. Volberg, and H. Zhang, Noncommutative Bohnenblust–Hille inequalities, *Math. Ann.* **389**, 1657 (2024).
- [117] M.-D. Choi, Completely positive linear maps on complex matrices, *Linear Algebra Appl.* **10**, 285 (1975).
- [118] A. Jamiołkowski, Linear transformations which preserve trace and positive semidefiniteness of operators, *Rep. Math. Phys.* **3**, 275 (1972).
- [119] J. Watrous, *The Theory of Quantum Information* (Cambridge University Press, USA, 2018).
- [120] M. A. Nielsen, and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, USA, 2010).
- [121] W.-T. Kuo, A. A. Akhtar, D. P. Arovas, and Y.-Z. You, Markovian entanglement dynamics under locally scrambled quantum evolution, *Phys. Rev. B* **101**, 224202 (2020).
- [122] H.-Y. Hu, S. Choi, and Y.-Z. You, Classical shadow tomography with locally scrambled quantum dynamics, *Phys. Rev. Res.* **5**, 023027 (2023).
- [123] A. A. Mele, Introduction to Haar measure tools in quantum information: A beginner’s tutorial, *Quantum* **8**, 1340 (2024).
- [124] C. L. Canonne, A short note on learning discrete distributions, *ArXiv:2002.11457*.
- [125] H. Buhrman, R. Cleve, J. Watrous, and R. De Wolf, Quantum fingerprinting, *Phys. Rev. Lett.* **87**, 167902 (2001).
- [126] A. Barenco, A. Berthiaume, D. Deutsch, A. Ekert, R. Jozsa, and C. Macchiavello, Stabilization of quantum computations by symmetrization, *SIAM J. Comput.* **26**, 1541 (1997).
- [127] H. Kobayashi, K. Matsumoto, and T. Yamakami, in *Algorithms and Computation: 14th International Symposium, ISAAC 2003, Kyoto, Japan, December 15–17, 2003. Proceedings 14* (Springer, Berlin, Heidelberg, 2003), p. 189.
- [128] L. N. Trefethen, and D. Bau, *Numerical Linear Algebra* (SIAM, Philadelphia, PA, 2022).
- [129] G. Golub, and W. Kahan, Calculating the singular values and pseudo-inverse of a matrix, *J. Soc. Ind. Appl. Math., Ser. B: Numer. Anal.* **2**, 205 (1965).
- [130] R. Kueng, and D. Gross, Qubit stabilizer states are complex projective 3-designs, *ArXiv:1510.02767*.
- [131] Z. Webb, The Clifford group forms a unitary 3-design, *Quantum Inf. Comput.* **16**, 1379 (2016).
- [132] H. Zhu, Multiqubit Clifford groups are unitary 3-designs, *Phys. Rev. A* **96**, 062336 (2017).
- [133] V. Lyubashevsky, C. Peikert, and O. Regev, in *Advances in Cryptology—EUROCRYPT 2010: 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, French Riviera, May 30–June 3, 2010. Proceedings 29* (Springer, Berlin, Heidelberg, 2010), p. 1.
- [134] D. Gottesman, *Stabilizer Codes and Quantum Error Correction* (California Institute of Technology, Pasadena, California, 1997).
- [135] A. Atıcı, and R. A. Servedio, Quantum algorithms for learning and testing juntas, *Quantum Inf. Process.* **6**, 323 (2007).
- [136] P. J. Coles, M. Cerezo, and L. Cincio, Strong bound between trace distance and Hilbert-Schmidt distance for low-rank states, *Phys. Rev. A* **100**, 022103 (2019).
- [137] J. Jiang, K. Wang, and X. Wang, Physical implementability of linear maps and its application in error mitigation, *Quantum* **5**, 600 (2021).
- [138] B. Regula, R. Takagi, and M. Gu, Operational applications of the diamond norm and related measures in quantifying the non-physicality of quantum maps, *Quantum* **5**, 522 (2021).
- [139] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, *Phys. Rev. Lett.* **119**, 180509 (2017).
- [140] X. Zhao, L. Zhang, B. Zhao, and X. Wang, Power of quantum measurement in simulating unphysical operations, *Phys. Rev. Res.* **7**, 013334 (2025).
- [141] P. Horodecki, and A. Ekert, Method for direct detection of quantum entanglement, *Phys. Rev. Lett.* **89**, 127902 (2002).
- [142] P. Horodecki, From limits of quantum operations to multicopy entanglement witnesses and state-spectrum estimation, *Phys. Rev. A* **68**, 052101 (2003).
- [143] J. Fiurášek, Structural physical approximations of unphysical maps and generalized quantum measurements, *Phys. Rev. A* **66**, 052315 (2002).
- [144] J. K. Korbicz, M. L. Almeida, J. Bae, M. Lewenstein, and A. Acin, Structural approximations to positive maps and

- entanglement-breaking channels, *Phys. Rev. A–At., Mol. Opt. Phys.* **78**, 062105 (2008).
- [145] M. X. Goemans, and D. Williamson, Approximation algorithms for max-3-cut and other problems via complex semidefinite programming, in *Proceedings of the thirty-third annual ACM symposium on Theory of computing* (Association for Computing Machinery, New York, NY, USA, 2001), p. 443.
- [146] D. Henrion, and J. Malick, Projection methods in conic optimization, *Handbook on Semidefinite, Conic and Polynomial Optimization*, (2012), p. 565.
- [147] J. Von Neumann, On rings of operators. Reduction theory, *Ann. Math.* **50**, 401 (1949).
- [148] J. von Neumann, Functional operators, Vol. II. Number 22 in *annals of mathematics studies* (1950).
- [149] R. L. Dykstra, An algorithm for restricted least squares regression, *J. Am. Stat. Assoc.* **78**, 837 (1983).
- [150] J. P. Boyle, and R. L. Dykstra, in *Advances in Order Restricted Statistical Inference: Proceedings of the Symposium on Order Restricted Statistical Inference held in Iowa City, Iowa, September 11–13, 1985* (Springer, New York, NY, 1986), p. 28.
- [151] H. H. Bauschke, and J. M. Borwein, Dykstra’s alternating projection algorithm for two sets, *J. Approx. Theory* **79**, 418 (1994).
- [152] S.-P. Han, A successive projection method, *Math. Program.* **40**, 1 (1988).
- [153] H. H. Bauschke, and A. S. Lewis, Dykstras algorithm with Bregman projections: A convergence proof, *Optimization* **48**, 409 (2000).
- [154] H. H. Bauschke, R. S. Burachik, D. B. Herman, and C. Y. Kaya, On Dykstra’s algorithm: Finite convergence, stalling, and the method of alternating projections, *Optim. Lett.* **14**, 1975 (2020).
- [155] F. Deutsch, and H. Hundal, The rate of convergence of Dykstra’s cyclic projections algorithm: The polyhedral case, *Numer. Funct. Anal. Optim.* **15**, 537 (1994).
- [156] N. Gaffke, and R. Mathar, A cyclic projection algorithm via duality, *Metrika* **36**, 29 (1989).
- [157] H. Hundal, and F. Deutsch, Two generalizations of Dykstra’s cyclic projections algorithm, *Math. Program.* **77**, 335 (1997).
- [158] C. J. Pang, Set intersection problems: Supporting hyperplanes and quadratic programming, *Math. Program.* **149**, 329 (2015).
- [159] F. Deutsch, and F. Deutsch, *Best Approximation in Inner Product Spaces* (Springer, New York, NY, 2001), Vol. 7.
- [160] H. H. Bauschke, and J. M. Borwein, On the convergence of von Neumann’s alternating projection algorithm for two sets, *Set-Valued Analysis* **1**, 185 (1993).
- [161] A. S. Lewis, D. R. Luke, and J. Malick, Local linear convergence for alternating and averaged nonconvex projections, *Found. Comput. Math.* **9**, 485 (2009).
- [162] P. Tseng, and S. Yun, A coordinate gradient descent method for nonsmooth separable minimization, *Math. Program.* **117**, 387 (2009).
- [163] Z.-Q. Luo, and P. Tseng, Error bounds and convergence analysis of feasible descent methods: A general approach, *Ann. Oper. Res.* **46**, 157 (1993).
- [164] A. Shapiro, Sensitivity analysis of generalized equations., *J. Math. Sci.* **115** (2003).
- [165] J. F. Bonnans, and A. Shapiro, *Perturbation Analysis of Optimization Problems* (Springer Science & Business Media, New York, NY, 2013).
- [166] Y. Nesterov *et al.*, *Lectures on Convex Optimization* (Springer, Cham, 2018), Vol. 137.
- [167] T. Surawy-Stepney, J. Kahn, R. Kueng, and M. Guta, Projected least-squares quantum process tomography, *Quantum* **6**, 844 (2022).