Testing and Learning Structured Quantum Hamiltonians

Srinivasan Arunachalam
IBM Quantum
Almaden Research Center, United States
Srinivasan.Arunachalam@ibm.com

Arkopal Dutt
IBM Quantum
Cambridge, Massachusetts, United States
arkopal@ibm.com

Francisco Escudero Gutiérrez
Qusoft and CWI
Amsterdam, Netherlands
feg@cwi.nl

Abstract

We consider the problems of testing and learning an unknown n-qubit quantum Hamiltonian $H = \sum_x \lambda_x \sigma_x$ expressed in its Pauli basis, from queries to its evolution operator e^{-iHt} under the normalized Frobenius norm. To this end, we prove the following results (with and without quantum memory) for Hamiltonians whose Pauli spectrum involves only k-local terms or has sparsity at most s:

(1) **Local Hamiltonians**: We give a *tolerant* testing protocol to decide if a Hamiltonian is ε_1-close to k-local or ε_2-far from k-local, with $O(1/(\varepsilon_2 - \varepsilon_1)^4)$ queries, thereby solving two open questions posed in a recent work by Bluhm, Caro and Oufkir [8]. For learning a k-local Hamiltonian up to error ε, we give a protocol with query complexity and total time evolution $\exp(O(k^2 + k\log(1/\varepsilon)))$. Our algorithm leverages the non-commutative Bohnenblust-Hille inequality in order to get a complexity independent of n.

(2) **Sparse Hamiltonians**: We give a protocol for testing whether a Hamiltonian is ε_1-close to being s-sparse or ε_2-far from being s-sparse, with $O(s^6/(\varepsilon_2^2 - \varepsilon_1^2)^6)$ queries. For learning up to error ε, we show that $O(s^4/\varepsilon^8)$ queries suffices.

(3) **Learning without quantum memory**: The learning results stated above have no dependence on the system size n, but require n-qubit quantum memory. We give subroutines that allow us to reproduce all the above learning results without quantum memory; increasing the query complexity by a $(\log n)$-factor in the local case and an n-factor in the sparse case.

(4) **Testing without quantum memory**: We give a new subroutine called *Pauli hashing*, which allows one to tolerantly test s-sparse Hamiltonians using $\tilde{O}(s^{14}/(\varepsilon_2^2 - \varepsilon_1^2)^{18})$ query complexity. A key ingredient is showing that s-sparse Pauli channels can be tested in a tolerant fashion as being ε_1-close to being s-sparse or ε_2-far under the diamond norm, using $\tilde{O}(s^2/(\varepsilon_2 - \varepsilon_1)^6)$ queries via Pauli hashing.

In order to prove these results, we prove new structural theorems for local Hamiltonians, sparse Pauli channels and sparse Hamiltonians. We complement our learning algorithms with lower bounds that are polynomially weaker. Furthermore, our algorithms use short time evolutions and do not assume prior knowledge of the terms on which the Pauli spectrum is supported on, i.e., we do not require prior knowledge about the *support* of the Hamiltonian terms.

STOC '25, Prague, Czechia

ACM ISBN 979-8-4007-1510-5/25/06
https://doi.org/10.1145/3717823.3718289

CCS Concepts

• **Theory of computation** → **Quantum complexity theory**.

Keywords

Testing, Learning, Quantum Hamiltonians, Low-degree, Sparse

ACM Reference Format:
Srinivasan Arunachalam, Arkopal Dutt, and Francisco Escudero Gutiérrez. 2025. Testing and Learning Structured Quantum Hamiltonians. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC '25), June 23–27, 2025, Prague, Czechia.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3717823.3718289

1 Introduction

A fundamental and important challenge with building quantum devices is being able to characterize and calibrate its behavior. One approach to do so is *Hamiltonian learning* which seeks to learn the Hamiltonian governing the dynamics of a quantum system given finite classical and quantum resources. Beyond system characterization, it is also carried out during validation of physical systems and designing control strategies for implementing quantum gates [34]. However, learning an n-qubit Hamiltonian is known to be difficult [12], requiring complexity that scales exponentially in the number of qubits.

In practice, however, prior knowledge on the structure of Hamiltonians is available e.g., those of engineered quantum devices [44] where the underlying Hamiltonians primarily involve local interactions with few non-local interactions, and even naturally occurring physical quantum systems such as those with translationally invariant Hamiltonians. To highlight these structural properties, consider an n-qubit Hamiltonian H (which is a self-adjoint operator acting on $(\mathbb{C}^2)^{\otimes n}$) expanded in terms of the n-qubit Pauli operators:

$$H = \sum_{x \in \{0,1\}^{2n}} \lambda_x \sigma_x,$$

where λ_x are real-valued coefficients (also called *interaction strengths*) to the Pauli operators σ_x denoted by the string $x = (a, b) \in \mathbb{F}_2^{2n}$ with $\sigma_{(a,b)} = i^{a\cdot b} \otimes_{i=1}^n X^{a_i} Z^{b_i}$. We call the set of Paulis with non-zero coefficients λ_x as the *Pauli spectrum* of the Hamiltonian denoted by $\mathcal{S} = \{x \in \{0,1\}^{2n} : \lambda_x \neq 0\}$. Of particular relevance are *k-local* Hamiltonians which involve Pauli operators that act non-trivially on all but at most k qubits and *s-sparse* Hamiltonians whose Pauli expansion contains at most s non-zero Pauli operators i.e., $|\mathcal{S}| \leq s$.

There has been a growing suite of Hamiltonian learning results that have shown that when the underlying n-qubit Hamiltonian H satisfies these structural properties, learning can be performed with only poly(n) query complexity, either by making "queries" to the unitary evolution operator $U(t) = \exp(-iHt)$ [20, 21, 29–31, 33, 37, 45, 50, 51], or by assuming one has access to Gibbs state [2, 5, 29, 30, 41, 42]. Notably, [6] considered the problem of learning Hamiltonians that are both local and sparse, without prior knowledge of the support. Several of the learning algorithms mentioned above, however require assumptions on the support of the Hamiltonian beyond locality or sparsity, such as [33] which considers *geometrically-local* Hamiltonians (a subset of local Hamiltonians) and [50] which requires assumptions on the support.

But, before learning, it might be desirable to uncover *what is the structure* of an unknown Hamiltonian in order to choose specialized learning algorithms. Even deciding if a Hamiltonian has a particular structure is a fundamental challenge and constitutes the problem of *testing* if an unknown Hamiltonian satisfies a certain structural property. As far as we know, this line of investigation is nascent with only a few works on Hamiltonian *property* testing [1, 36, 43] with Bluhm, Caro and Oufkir [9] having considered the problem of testing local Hamiltonians. This leads us to the motivating question:

What is the query complexity of learning and testing structured Hamiltonians?

1.1 Problem Statement

Before we state our results answering the question above, we clearly mention our learning and testing problems first. If H is the Hamiltonian describing the dynamics of a certain physical system, then the state of that system evolves according to the *time evolution operator* $U(t) = e^{-iHt}$. This means that if $\rho(0)$ is the state at time 0, at time t the state would have evolved to $\rho(t) = U(t)\rho(0)U^\dagger(t)$. Hence, to test and learn a Hamiltonian one can do the following: prepare a desired state, apply $U(t)$ or tensor products of $U(t)$ with identity to the state, and finally measure in a chosen basis. From here onwards, this is what we mean by *querying* the unitary $U(t)$. It is usual to impose the normalization condition $\|H\|_\infty \leq 1$ (i.e., that the eigenvalues of H are bounded in absolute value by 1), as otherwise the complexities scale with the norm of the Hamiltonian. Throughout this paper, we will consider the normalized Frobenius norm as the distance between Hamiltonians, namely

$$d(H, H') = \|H - H'\|_2 = \sqrt{\frac{\mathrm{Tr}[(H-H')^2]}{2^n}},$$

which equals the 2-norm of the Pauli spectrum, $\sqrt{\sum |\lambda_x - \lambda'_x|^2}$. We choose the Frobenius norm here since it's natural to want to approximate the action of a Hamiltonian on most states, rather than *all states* [9, Appendix A].

A *property* of a Hamiltonian, denoted $\mathcal{H}$ is a class of Hamiltonians that satisfy the property (here we will be interested in sparse and local properties). We say that H is ε-far from having a property $\mathcal{H}$ if $d(H, H') > \varepsilon$ for every $H' \in \mathcal{H}$, and otherwise is ε-close. Now, we are ready to state the testing and learning problems.

Let $\mathcal{H}$ be a property and let H be an unknown Hamiltonian with $\|H\|_\infty \leq 1$ and $\mathrm{Tr}[H] = 0$.

Problem 1.1 (Tolerant testing). *Promised H is either ε_1-close or ε_2-far from satisfying property $\mathcal{H}$, decide which is the case by making queries to $U(t)$.*

Problem 1.2 (Hamiltonian learning). *Promised $H \in \mathcal{H}$, output a classical description of $\widetilde{H} \in \mathcal{H}$ such that $\|H - \widetilde{H}\|_2 \leq \varepsilon$ by making queries to $U(t)$.*

1.2 Summary of Results

The main results of this submission are query-efficient algorithms for testing and learning Hamiltonians that are local and/or sparse. We can reproduce these results without using quantum memory by increasing the number of queries. We summarize our results in the table below (for simplicity we state our results for constant accuracy).

Before we discuss our results in more detail, we make a few remarks about our main results.

(*i*) As far as we know, this is the first work: (*a*) with complexities that are *independent* of n (with memory)[1], and (*b*) that does not assume knowledge of the support.

(*ii*) We give the first learning algorithm for Hamiltonians that are only promised to be sparse, and not necessarily local. Similarly, our local Hamiltonian learning algorithm does not further assume geometric locality which has been assumed in several prior works.

(*iii*) As far as we know, this is one of the the first polynomial-time *tolerant testers* for quantum objects. Prior tolerant algorithms [3, 26] require a *multiplicative* gap between $\varepsilon_1, \varepsilon_2$, whereas we have runtime that depends only on $\varepsilon_2 - \varepsilon_1$.

(*iv*) We show that all our learning protocols with quantum memory can be translated to ones which require no quantum memory. In the case of learning structured Hamiltonians, we obtain a protocol with only a factor $\log n$ overhead for local Hamiltonians and a protocol with a factor n overhead for sparse Hamiltonians.

(*v*) We also give a tolerant testing algorithm for s-sparse Hamiltonians that requires no quantum memory based on a new subroutine called Pauli hashing. The query complexity is $O(\mathrm{poly}(s))$ and is notably independent of dimension n.

We remark that most previous work on Hamiltonian learning (that we highlighted earlier) are done under the distance induced by the supremum norm of the Pauli spectrum and with extra constraints apart from locality [12, 20, 21, 29–31, 33, 37, 39, 45, 47, 50, 51]. When transformed into learning algorithms under the finer distance induced by the 2-norm of the Pauli spectrum, these proposals yield complexities that depend polynomially on n^k and only work for a restricted family of k-local Hamiltonians. The works that explicitly consider the problem of learning under the 2-norm have complexities depending on n and assume a stronger access model [6, 13].

[1] There are a few works that achieve n-independent complexities for learning local Hamiltonians in the ∞-norm of the Pauli coefficients, but when transformed into 2-norm learners they yield complexities depending on n^k.

Table 1: Query complexity for learning and testing n-qubit structured Hamiltonians. Dependence on n and the structural property is shown for constant accuracy. Results are indicated with quantum memory (i.e., an n-qubit ancillary system is available) and without quantum memory.

	Testing		Learning	
	with memory	*w/o memory*	*with memory*	*w/o memory*
s-sparse	$\mathrm{poly}(s)$	$\mathrm{poly}(s)$	$\mathrm{poly}(s)$	$n \cdot \mathrm{poly}(s)$
k-local	$O(1)$	$O(1)$ [9]	$\exp(k^2)$	$(\log n) \cdot \exp(k^2)$
k-local & s-sparse	$\mathrm{poly}(s)$	$\mathrm{poly}(s)$	$\min\{\exp(k^2), \mathrm{poly}(s)\}$	$(\log n) \cdot \min\{\exp(k^2), \mathrm{poly}(sk)\}$

In the next few sections we present the proof sketches of our main results. We describe four results: testing as well as learning of sparse and local Hamiltonians in separate sections, first highlighting the technical contribution before discussing the algorithms. See the full version of the paper [4] for more details.

1.3 Result 1: Learning with Quantum Memory

1.3.1 Analytic Operator Inequalities. **Taylor series.** Conceptually, the high-level idea of accessing the Hamiltonian H via the time-dependent unitary evolution operator e^{-iHt} is pretty simple: consider short times t and use Taylor expansion to treat $U(t) = e^{-itH}$ as $\mathrm{Id} - itH$. However, analyzing the propagation of the error of this approximation requires some delicate calculations. For instance, for testing sparsity we need to estimate the sum of the squares of the smallest Pauli coefficients of H, but due to the error in the Taylor expansion we do not have that the smallest Pauli coefficients of H correspond to the smallest coefficients of $U(t)$ even for short time steps. Similar issues occur for learning. Overcoming these to obtain our testing and learning algorithm involves technical calculations in the error analysis of the algorithms.

BH inequalities. The Bohnenblust-Hille inequality is well-known since the 1930s in analysis and has become popular more recently in theoretical computer science and quantum computing [10, 32, 46]. Our locality learning algorithm also makes crucial use of Bohnenblust-Hille inequality to attain a complexity independent of n. To see why, consider the problem of learning an n-qubit k-local Hamiltonian with $\|H\|_2 \leq 1$, an ε-packing argument suggests that one would expect $\Omega((3n)^k)$ lower bound just information theoretically. Even though we additionally have the restriction that the Hamiltonians satisfy $\|H\|_\infty \leq 1$, we have the sometimes-tight inequality $\|H\|_2 \leq \|H\|_\infty$, which makes the suspected lower bound of $\Omega((3n)^k)$ plausible. But the non-commutative Bohnenblust-Hille inequality [32, 46] ensures that for Hamiltonians with $\|H\|_\infty \leq 1$ only $O_k(1)$ Pauli coefficients are non-negligible, so it suffices to only find these and estimate them. This surprisingly allows one to obtain n-independent complexity for learning.

1.3.2 Learning Protocols.

Local Hamiltonians. We now present our learning result for local Hamiltonians.

Result 1.3. *There is an algorithm that solves Theorem 1.2 for k-local Hamiltonians by making $\exp(k^2 + k\log(1/\varepsilon))$ queries to the evolution operator with $\exp(k^2 + k\log(1/\varepsilon))$ total evolution time.*

Our learning protocol for local Hamiltonians consists on two steps. In the first, we sample from the Pauli spectrum of the short-time evolution operator, which can be done by measuring the correspondent Choi state in the Bell basis. Then, we store the *big coefficients*, which are those that are sampled with frequency larger than a certain threshold. In the second step we learn these big coefficients one by one. Thanks to the non-commutative Bohnenblust-Hille inequality, the threshold that defines what coefficients are big can be taken to be independent of n, which results in complexities independent of n. A bit more formally, in the first step we find the big coefficients, those that

$$|\lambda_x| = \Omega_{k,\varepsilon}(1).$$

Since $\sum_x |\lambda_x|^2 \leq 1$, there is at most $O_{k,\varepsilon}(1)$ big coefficients, so by using the protocol of Montanaro and Osborne [38] $O_{k,\varepsilon}(1)$ times we can obtain estimates $\tilde{\lambda}_x$ such that

$$\sum_{|\lambda_x| = \Omega_{k,\varepsilon}(1)} |\tilde{\lambda}_x - \lambda_x|^2 \leq \varepsilon^2. \tag{1}$$

Bohnenblust-Hille inequality ensures that the small coefficients are negligible, meaning that

$$\sum_{|\lambda_x| = O_{k,\varepsilon}(1)} |\lambda_x|^2 \leq \varepsilon^2. \tag{2}$$

Putting Eqs. (1) and (2) together, it follows that $\sum_{|\lambda_x| = \Omega_{k,\varepsilon}(1)} \tilde{\lambda}_x \sigma_x$ is $O(\varepsilon)$-close to H.

Sparse Hamiltonians. For learning sparse Hamiltonians, we have the following result, which, as far as we know, is the first result for learning sparse, and non-necessarily local, Hamiltonians.

Result 1.4. *There is an algorithm that solves Theorem 1.2 for s-local Hamiltonians by making $\mathrm{poly}(s/(\varepsilon_2 - \varepsilon_1))$ queries to the evolution operator and with $\mathrm{poly}(s/(\varepsilon_2 - \varepsilon_1))$ total evolution time.*

Our learning protocol for sparse Hamiltonians is similar to the one for local Hamiltonians. The difference is the threshold that defines what is a *big coefficient*. Here, we can take it to be ε^2/s, as

$$\sum_{|\lambda_x| \leq \varepsilon^2/s} |\lambda_x|^2 = \varepsilon^2$$

because H has at most s non-zero Pauli coefficients.

1.4 Result 2: Testing with Quantum Memory

So far, we saw how to use Pauli-analytical tools in order for learning. In this section we move to testing these Hamiltonians. Before talking about the testing results themselves, we present a

few structural lemmas about sparse and local Hamiltonians which will be useful for tolerant testing.

1.4.1 Structural Lemmas. As we can only access a Hamiltonian H through its time evolution operator $U(t)$, in order to test whether H satisfies a property, we first need to show that the property is preserved by $U(t)$. To this end we prove quantitative structural lemmas that ensure that the properties of being local and sparse are preserved by $U(t)$.

Local Hamiltonians. For k-locality, one can show that the distance of H to the space of k-local Hamiltonians is $\sum_{|x|>k} |\lambda_x|^2$. Motivated by this, we introduce the *k-non-local energy of $U(t)$*, given by

$$\mathsf{NonLocalEnergy}(k,t) = \sum_{|x|>k} |\widehat{U}_x(t)|^2.$$

By carefully analyzing Taylor expansion, we show that for short times $\mathsf{NonLocalEnergy}(k,t)$ is small if and only if H is close to local.

Sparse Hamiltonians. For s-sparsity, one can show that the distance of a Hamiltonian H to the set of s-sparse Hamiltonians is given by

$$d(H, s\text{-sparse}) = \sqrt{\sum_{i>s} |\lambda_{x_i}|^2},$$

where $x_1, \ldots, x_{4^n}$ are labels to $\{0,1\}^{2n}$ that satisfy $|\lambda_{x_1}| \geq |\lambda_{x_2}| \geq \cdots \geq |\lambda_{x_{4^n}}|$. Then, one would be tempted to define the *s-non-sparse energy of $U(t)$* as the sum of the squares of its $4^n - s$ smallest Pauli coefficients. However, it is not clear how to efficiently compute this tentative quantity. This is because the way we access the Pauli distribution of $U(t)$ is by sampling, which may incur in an estimation error for every of the $4^n - s$ coefficients, which would imply a dependence on n. To circumvent this we introduce the *top-s energy* of $U(t)$, which is given by

$$\mathsf{TopEnergy}(s,t) = |\widehat{U}_{0^n}|^2 + \sum_{i\leq s} |\widehat{U}_{y_i}(t)|^2, \tag{3}$$

where $y_1, \ldots, y_{4^n-1}$ are labels of $\{0,1\}^{2n} - \{0^n\}$ such that $|\widehat{U}_{y_1}| \geq \cdots \geq |\widehat{U}_{y_{4^n-1}}|$. The top-$s$ energy can be approximated efficiently, as it is the sum of just $s+1$ squared coefficients of $U(t)$. Our structural lemma for sparse Hamiltonians establishes that for short times the top-s energy is close to 1 *if and only if* H is close to sparse. To give intuition of why this is true, consider for a moment the Taylor expansion of $U(t)$ disregarding the higher-order error terms above the second-order term, i.e., treat $U(t)$ as

$$U(t) \approx \mathrm{Id} - itH - \frac{t^2}{2}H^2.$$

Then, as H is traceless, we would have $\widehat{U}_{0^n} \approx 1 - \frac{t^2}{2}\sum_x \lambda_x^2$. Ignoring higher order terms suggests

$$|\widehat{U}_{0^n}|^2 \approx 1 - t^2 \sum_{x\in\{0,1\}^{2n}} \lambda_x^2. \tag{4}$$

Similarly, by ignoring higher order terms for $x \neq 0^n$ we would get $|U_x|^2 \approx t^2\lambda_x^2$. If we further assume that the largest Pauli coefficients of $U(t)$ correspond to the largest Pauli coefficients of H, it would follow from Eq. (4) that

$$\mathsf{TopEnergy}(s,t) \approx 1 - t^2 \sum_{i>s} |\lambda_{x_i}|^2 = 1 - t^2 \cdot d(H, s\text{-sparse}).$$

1.4.2 Testing Protocols.

Testing locality. Recently, Bluhm, Caro and Oufkir proposed a non-tolerant testing algorithm, meaning that it only works for the case $\varepsilon_1 = 0$, whose query complexity is $O(n^{2k+2}/(\varepsilon_2 - \varepsilon_1)^4)$ and with total evolution time $O(n^{k+1}/(\varepsilon_2 - \varepsilon_1)^3)$. They posed as open questions whether the dependence on n could be removed and whether an efficient tolerant-tester was possible [8, Section 1.5]. Our first result gives a positive answer to both questions.

Result 1.5. *There is an algorithm that solves Theorem 1.1 for k-local Hamiltonians by making* $\mathrm{poly}(1/(\varepsilon_2 - \varepsilon_1))$ *queries to the evolution operator and with* $\mathrm{poly}(1/(\varepsilon_2 - \varepsilon_1))$ *total evolution time.*

Thanks to our structural lemma, to test for locality it suffices to estimate $\mathsf{NonLocalEnergy}(k,t)$ for some short time t. To do that we sample from $\{|\widehat{U}_x|^2\}_x$ a number of $\mathrm{poly}(1/(\varepsilon_2 - \varepsilon_1))$ times and compute the empirical frequency of the event $\{|x| > k\}$. This frequency equals $\mathsf{NonLocalEnergy}(k,t)$ on expectation.

Testing sparsity. Despite numerous papers in the classical literature studying problems of testing and learning sparse Boolean functions [23, 27, 40, 48], there have not been many results on learning Hamiltonians that are *sparse* (and not necessarily local) and the only testing result that we are aware of requires $O(sn)$ queries [9, Remark B.2]. Here, we present the first sparsity testing algorithm whose complexity does not depend on n and the first learning algorithm for sparse Hamiltonians which does not require assumptions regarding the support.

Result 1.6. *There is an algorithm that solves Theorem 1.1 for s-local Hamiltonians by making* $\mathrm{poly}(s/(\varepsilon_2 - \varepsilon_1))$ *queries to the evolution operator and with* $\mathrm{poly}(s/(\varepsilon_2 - \varepsilon_1))$ *total evolution time.*

Thanks to our structural lemma, to test for sparsity it suffices to estimate $\mathsf{TopEnergy}(s,t)$ for some short time t. As we do not know which are the largest coefficients of $U(t)$, we cannot compute $\mathsf{TopEnergy}(s,t)$ as the expectation of a random variable. Instead, we obtain an empirical estimate $|\alpha_x|^2$ of $|\widehat{U}_x|^2$, and use the sum of $|\alpha_{0^n}|^2$ and the s largest $|\alpha_x|^2$ as a proxy for $\mathsf{TopEnergy}(s,t)$. To ensure that this is a good proxy, we need to sample from $\{|\widehat{U}_x(t)|^2\}$ a number of $\mathrm{poly}(s/(\varepsilon_2 - \varepsilon_1))$ times.

1.5 Result 3: Learning Without Quantum Memory

Motivated by the limitations of current devices, there has been a series of recent works aiming to understand the power of quantum memory in testing and learning tasks, exhibiting exponential separations in some cases [14, 16, 17]. A natural question is, if the problems that we mentioned above become harder without quantum memory? We surprisingly show that the learning protocols that we mention above, can be implemented efficiently when one has *no quantum memory*. To this end, we provide two subroutines for: (i) estimating the Pauli spectrum of a unitary, (ii) estimating a single Pauli coefficient to make our protocols work in the memory-less setting. Subroutine (ii) incurs in no extra query-cost, and subroutine (i) only incurs in a factor-n overhead in the case of learning s-sparse Hamiltonians and a factor $\log(n)$ in the case of learning k-local Hamiltonians. Hence, we have the following result.

Result 1.7. *There are memory-less algorithms to learn n-qubit Hamiltonians H such that*

- *if H is k-local, require just* $\exp(k^2)\cdot\log(n)$ *queries.*
- *if H is s-sparse, require just* $O(s^4 n)$ *queries.*
- *if H is k-local and s-sparse, require just* $\min\{\exp(k^2), O(s^4 k))\}\cdot\log(n)$ *queries.*

Now, we describe these subroutines and in Section 1.7 we show that they can also be useful in other contexts.

Estimating the Pauli spectrum of a unitary without memory. In order to substitute sampling from the Pauli spectrum of unitaries in our learning algorithms, we propose a memory-less protocol to estimate the Pauli distribution of a unitary. We use the same protocol as other works, but with different post-processing [9, 24, 25, 49]. To describe it, consider 2^n+1 mutually unbiased bases $\mathcal{B}_i = \{|\phi_{i,j}\rangle\}_{j\in[2^n]}$ of stabilizer states. One iteration of this protocol consists on picking $i \in [2^n+1]$ and $j \in [2^n]$ uniformly at random, preparing $|\phi_{i,j}\rangle$, applying U, and finally measuring on $\mathcal{B}_i$. Let $|\phi_{i,l}\rangle$ is the outcome of this measurement. For every $x \in \{0,1\}^{2n}$, we define a random variable M_x that takes value 1 if $|\phi_{i,l}\rangle\langle\phi_{i,l}| = \sigma_x|\phi_{i,j}\rangle\langle\phi_{i,j}|\sigma_x$ and 0 otherwise. In other words, M_x takes value 1 if the change from $|\phi_{i,j}\rangle$ to $|\phi_{i,l}\rangle$ can be attributed to the weight of σ_x in U. Hence, on expectation, M_x roughly equals $|\widehat{U}_x|^2$. By an union bound, repeating this iteration $O(\log(M))$ times allows one to estimate M squares of Pauli coefficients of U.

Estimating one Pauli coefficient We also propose a protocol to estimate one Pauli coefficient of a Hamiltonian that requires no quantum memory. Additionally, it only involves querying $U(t)$ with inputs of Pauli eigenstates and performing Pauli measurements.

1.6 Result 4: Testing Without Quantum Memory

We provide a testing algorithm for s-sparse Hamiltonians which does not use any quantum memory and only requires $\mathrm{poly}(s/\varepsilon)$ queries and total evolution time, notably avoiding any dependence on n. To obtain this tester, we introduce a new subroutine called *Pauli hashing* which allows us to construct random partitions of the Pauli operators and requires no quantum memory.

We first discuss Pauli hashing in the context of the following testing problem for Pauli channels:

Problem 1.8. *Given query access to a Pauli channel* $\mathcal{E}:\rho\mapsto\sum_x p(x)\sigma_x\rho\sigma_x$, *decide if* $\mathcal{E}$ *is* ε_1*-close to being s-sparse or* ε_2*-far away from all such Pauli channels, under the diamond norm.*

We call a Pauli channel $\mathcal{E}:\rho\mapsto\sum_x p(x)\sigma_x\rho\sigma_x$ as s-sparse if $\mathcal{E}$ has at most s many non-zero error rates $p(x)$. Of relevance is the Pauli fidelity $\lambda:\mathbb{F}_2^{2n}\to[-1,1]$ defined as $\lambda(x)=\frac{1}{2^n}\mathrm{Tr}(\sigma_x\mathcal{E}(\sigma_x))$ and which has the symplectic Fourier transform

$$\lambda(x)=\sum_{\alpha\in\mathbb{F}_2^{2n}} p(\alpha)(-1)^{[x,\alpha]},$$

where $[x,\alpha]$ is the symplectic inner product that equals 0 iff σ_x and σ_α commute, and one otherwise. Note that the symplectic Fourier coefficients of λ are the error rates $\{p(x)\}$. Testing sparsity of $\mathcal{E}$ is then equivalent to testing sparsity of the symplectic Fourier spectrum of λ.

We prove a simple structural lemma regarding sparse Pauli channels to aid us. Let us define the *energy* of the top s error rates of $\mathcal{E}$ as $\mathsf{Energy}(\mathcal{E};s):=\max_{T\subseteq\mathbb{F}_2^{2n}:\,|T|=s}\left\{\sum_{x\in T}p(x)\right\}$. We show if $\mathcal{E}$ is ε_1-close to being s-sparse under the total variation distance (which is same as the diamond norm up to a factor of 2) then $\mathsf{Energy}(\mathcal{E};s)\geq 1-2\varepsilon_1$ and if ε_2-far away then $\mathsf{Energy}(\mathcal{E};s)\leq 1-\varepsilon_2$.

1.6.1 Pauli Hashing.

The high-level idea behind our testing algorithm (as well as that of testing sparse Hamiltonians) is that of *hashing* or isolating the n-qubit Pauli operators (i.e., symplectic Fourier coefficients of λ) into different *buckets*. These buckets are defined by the commutation relations of the n-qubit Pauli operators with an appropriately defined subgroup G of Paulis. This is analogous to Fourier hashing developed by Gopalan, O'Donnell, Servedio, Shpilka, and Wimmer [27] for testing sparsity of the Fourier spectrum of Boolean functions. There, they utilized inner product relations of the standard Fourier characters with a particular subspace. Our Pauli hashing can thus be equivalently described as *symplectic* Fourier hashing. We now describe the Pauli hashing process.

Bucketting. We start off by choosing a random subspace G of $\mathbb{F}_2^{2n}$ of dimension t (to be specified later). Note that the elements of G correspond to a subgroup of Pauli operators, which we will also denote by G in a slight abuse of notation. Consider a set of t Pauli generators $g_1,\ldots,g_t$ such that $G=\mathrm{span}\{g_1,\ldots,g_t\}$. To describe the commutation relations of an arbitrary n-qubit Pauli operator P with the elements in G, it is enough to note down a vector $b\in\mathbb{F}_2^t$ containing the t commutation relations of P with each of the generators g_i for $i\in[t]$. Particularly, we can then partition all the n-qubit Pauli operators into 2^t buckets of the following form for each $b\in\mathbb{F}_2^t$

$$C(b):=\{\alpha\in\mathbb{F}_2^{2n}:[\alpha,g_j]=b_j\;\forall j\in[t]\},$$

where we have denoted the Paulis by their corresponding $2n$-bit strings. These buckets $\{C(b)\}_{b\in\mathbb{F}_2^t}$ are the cosets of the commutant or centralizer of G, denoted by $\mathsf{C}(G)$, constituted by the Paulis that commute with all the elements in G. Our first observation is that for a random subgroup G, defining buckets as above corresponds to a pairwise random hashing process.

Energy of buckets. We then define the energy $E(b)$ of a bucket $C(b)$ as the sum of all the error rates in $C(b)$, $E(b)=\sum_{x\in C(b)}p(x)$. We show $E(b)$ can be efficiently estimated with queries to the Pauli channel $\mathcal{E}$ involving simple Pauli eigenstate preparation and Pauli measurements, requiring query complexity $O(2^t/\varepsilon^2)$ for accuracy ε. If $t\geq 2\log s$ (i.e., $|G|=O(s^2)$), and $\mathcal{E}$ was exactly s-sparse, then each of the error rates $p(x)$ of $\mathcal{E}$ would likely land in a distinct bucket i.e., coset of $\mathsf{C}(G)$. One approach would then be to check if top s buckets (ordered by energy) out of the total $O(s^2)$ buckets, have energies summing up to one. However, it turns out one can do better requiring a lower value $t\geq\log(2s)$ (or $O(s)$ many buckets) by leveraging the structural lemma regarding sparse Pauli channels noted earlier, combined with careful error analysis.

Correctness. We earlier noted that $\mathsf{Energy}(\mathcal{E};s)$ or the sum of the top s error rates of $\mathcal{E}$ was a good proxy for checking sparsity. Our main technical contribution is to show that the total energy of the top s buckets i.e., $\max_{b\in S,|S|=s}E(b)$ is a good estimate

of $\mathsf{Energy}(\mathcal{E}; s)$ for $t \geq \log(2s)$. This requires a careful analysis of two sources of error: (i) *hashing error* which is the error between the total energy of the top s buckets and top s error rates i.e., $\left(\max_{b\in S,|S|=s} E(b)\right) - \mathsf{Energy}(\mathcal{E}; s)$, and the (ii) *estimation error* committed when estimating $E(b)$. To bound the hashing error, we bound the collision error or $E(b) - \max_{x\in C(b)} p(x)$ which accounts for the energy mismatch due to two or more error rates colliding in the same bucket, and show this can be made sufficiently small for $t \geq \log(2s)$. The estimation error can be controlled as a consequence of the guarantees of our energy estimation algorithm. Putting everything together, we obtain the following result.

Result 1.9. *There is an algorithm requiring no quantum memory that tests if a Pauli channel $\mathcal{E}$ is ε_1-close to or ε_2-far from being s-sparse in diamond norm, using $\widetilde{O}(s^2/(\varepsilon_2 - \varepsilon_1)^6)$ queries to $\mathcal{E}$.*

We remark that Pauli hashing only requires inputs of Pauli eigenstates to $\mathcal{E}$ and Pauli measurements, making it suitable for the near-term. We expect Pauli hashing to be useful for testing other quantum objects. We emphasize that the key ingredients were (i) converting the testing problem to testing sparsity of the symplectic Fourier spectrum of a function f that is easy to compute via queries, and (ii) a structural lemma commenting on sparsity of the quantum object in terms of f.

1.6.2 Testing Sparse Hamiltonians via Pauli Hashing. We reduce the problem of testing sparse Hamiltonian to testing the sparsity of an associated Pauli channel. The Hamiltonian evolution channel $\mathcal{H}_t : \rho \to U(t)\rho U^\dagger(t)$ can be transformed into a Pauli channel by applying Pauli twirling [7, 11]. We denote the *Pauli-twirled* Hamiltonian evolution channel as $\mathcal{H}_t^{\mathcal{T}}(\rho) = \mathbb{E}_x[\sigma_x \mathcal{H}_t(\sigma_x \rho \sigma_x)\sigma_x]$. Our first observation is that the error rates of $\mathcal{H}_t^{\mathcal{T}}$ are the squares of the Pauli coefficients $\{|\widehat{U}(t)|^2\}$ of the unitary $U(t) = \exp(-iHt)$. Combined with the structural lemma regarding s-sparse Hamiltonians (Eq. 3), we can thus test sparsity of H by testing sparsity of the Pauli channel $\mathcal{H}_t^{\mathcal{T}}$, whose error rates we denote by $\{p_t^{\mathcal{T}}(x)\}$.

The goal of the testing algorithm is approximating $\mathsf{TopEnergy}(s, t)$ of H by estimating $p_t^{\mathcal{T}}(0^{2n}) + \mathsf{Energy}(\mathcal{E}; s)$, where $\mathsf{Energy}(\mathcal{E}; s)$ is the energy of the top s error rates not including $p_t^{\mathcal{T}}(0^{2n})$. This is done due to the form of $\mathsf{TopEnergy}(s, t)$ which requires the estimation of $|\widehat{U}(0^{2n})|^2$. We then show that $p_t^{\mathcal{T}}(0^{2n}) + \mathsf{Energy}(\mathcal{E}; s)$ can be obtained efficiently via Pauli hashing with dimension of the random group G being $d \geq \log(2s)$ (denoting as d here so as not to be confuse with time t) and estimation of total energy across bucket $C(0^d)$ and the top s buckets excluding $C(0^d)$. We further give an efficient energy estimation algorithm which does not require query access to $\mathcal{H}_t^{\mathcal{T}}$ but simply queries the Hamiltonian evolution channel $\mathcal{H}_t$ with inputs of Pauli eigenstates, interleaved applications of random Pauli gates and Pauli measurements. The analysis of the hashing and estimation error then follows similar to what we described for testing Pauli channels via Pauli hashing. This leads to the following result.

Result 1.10. *There is an algorithm requiring no quantum memory that solves Theorem 1.2 for s-sparse Hamiltonians by making $\mathrm{poly}(s/\varepsilon)$ queries to the evolution operator with $\mathrm{poly}(s/\varepsilon)$ total evolution time.*

1.7 Other Results

Tolerant junta unitary testing. Additionally in this work we obtain a tolerant testing algorithm to decide if an unknown unitary is a k-junta that makes $O(4^k)$ queries making progress on a question of Chen, Nadimpalli and Yuen [19, Section 1.3], and then we use subroutine (i) to turn it into a memory-less tester that only makes $O(4^k n)$ queries.

Lower bounds. One drawback of our learning and testing algorithms is the exponent of the sparsity parameter s, locality parameter k and the tolerance $(\varepsilon_2 - \varepsilon_1)$. Reducing to classical Boolean functions, we show lower bounds without memory that cerfity that the dependence on these parameters cannot be completely avoided.

1.8 Discussion and Open Questions

Our work opens up several interesting directions which we state here and leave for future work.

(1) **Dependence on parameters $\varepsilon_1, \varepsilon_2$.** Our main objective in this work was to obtain query complexities for testing and learning with good dependence on the structural parameters. It is natural to ask if we could improve the dependence on the error parameters and perhaps achieve Heisenberg limited scaling as has been shown to be possible in some particular cases for Hamiltonian learning [6, 33].
(2) **Robustness to SPAM noise.** It would be desirable to make the protocols introduced in this work to be robust to SPAM noise. A potential approach is to adapt strategies in [25].
(3) **Adaptivity.** For learning structured Hamiltonians, adaptive strategies [22, 28] can improve query complexity by shedding constant factors over baseline learning algorithms, thereby improving performance in practice. Another direction is to then explore adaptive protocols for testing structured Hamiltonians and the performance gains they may bring.
(4) **Testing and learning with limited quantum memory.** For estimating properties of quantum states, Chen, Cotler, Huang, and Li [14] showcased the utility of the resource of quantum memory or a k-qubit ancillary system ($k < n$). Large separations in query complexity when learning with memory (even for $k \ll n$) and without any have been reported for learning Pauli channels [15, 18] and shadow tomography [16, 35]. We could thus imagine having access to only limited quantum memory during learning or testing structured Hamiltonians as well. However, it should be noted that given the separation between the query complexities (see Table 1) with access to n-qubit quantum memory and without any, only marginal gains in complexity are expected with access to limited memory.
(5) **Testing and learning Hamiltonians from Gibbs states.** Another natural learning model is that of having access to copies of the Gibbs state of a quantum Hamiltonian at a certain inverse temperature. There has been a suite of work investigating learning local Hamiltonians from Gibbs states [2, 5] but answering the question of testing structured Hamiltonians given access to copies of the Gibbs state remains wide open.

Acknowledgements. We thank the anonymous referees for their helpful feedback. S.A. and A.D. thank the Institute for Pure and Applied Mathematics (IPAM) for its hospitality throughout the long program "Mathematical and Computational Challenges in Quantum Computing" in Fall 2023 during which part of this work was initiated. This work was done in part while S.A. was visiting the Simons Institute for the Theory of Computing, supported by DOE QSA grant #FP00010905. This research was supported by the Europea union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no. 945045, and by the NWO Gravitation project NETWORKS under grant no. 024.002.003. We thank Amira Abbas, Francesco Anna Mele, Andreas Bluhm, Jop Briët, Matthias Caro, Nunzia Cerrato, Aadil Oufkir, and Daniel Liang for useful comments and discussions. A.D. thanks Patrick Rall for multiple conversations on stabilizer subgroups and Pauli twirling. A.D. thanks Isaac Chuang for discussions on the problem of testing Hamiltonians. F.E.G. is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC-2047/1 – 390685813. We thank the referees for their detailed comments. F.E.G. is partially supported by the MICINN project PID2023-146758NB-I00 funded by MICIU/AEI/10.13039/501100011033, by TEC-2024/COM-84-QUITEMAD-CM, funded by Comunidad de Madrid and by Grant CEX2023-001347-S funded by MICIU/AEI/10.13039/501100011033

References

[1] Dorit Aharonov, Jordan Cotler, and Xiao-Liang Qi. 2022. Quantum algorithmic measurement. *Nature Communications* 13, 1 (2022), 1–9. doi:10.1038/s41467-021-27922-0

[2] Anurag Anshu, Srinivasan Arunachalam, Tomotaka Kuwahara, and Mehdi Soleimanifar. 2021. Sample-efficient learning of interacting quantum systems. *Nature Physics* 17, 8 (2021), 931–935. doi:10.1038/s41567-021-01232-0

[3] Srinivasan Arunachalam and Arkopal Dutt. 2024. Polynomial-time tolerant testing stabilizer states. *arXiv:2408.06289* (2024).

[4] Srinivasan Arunachalam, Arkopal Dutt, and Francisco Escudero Gutiérrez. 2024. Testing and learning structured quantum Hamiltonians. *arXiv preprint arXiv:2411.00082* (2024). [Full version].

[5] Ainesh Bakshi, Allen Liu, Ankur Moitra, and Ewin Tang. 2023. Learning quantum Hamiltonians at any temperature in polynomial time. arXiv:2310.02243 [quant-ph]

[6] Ainesh Bakshi, Allen Liu, Ankur Moitra, and Ewin Tang. 2024. Structure learning of Hamiltonians from real-time evolution. In FOCS'24.

[7] Ewout van den Berg and Pawel Wocjan. 2023. Techniques for learning sparse Pauli-Lindblad noise models. *arXiv preprint arXiv:2311.15408* (2023).

[8] Andreas Bluhm, Matthias C Caro, and Aadil Oufkir. 2024. Hamiltonian Property Testing (version 1). *arXiv preprint 2403.02968v1* (2024). arXiv:2403.02968v1

[9] Andreas Bluhm, Matthias C Caro, and Aadil Oufkir. 2024. Hamiltonian Property Testing (version 2). *arXiv preprint 2403.02968v2* (2024). arXiv:2403.02968v2

[10] Henri Frédéric Bohnenblust and Einar Hille. 1931. On the absolute convergence of Dirichlet series. *Annals of Mathematics* (1931), 600–622.

[11] Zhenyu Cai and Simon C Benjamin. 2019. Constructing smaller pauli twirling sets for arbitrary error channels. *Scientific reports* 9, 1 (2019), 11281.

[12] Matthias C. Caro. 2023. Learning Quantum Processes and Hamiltonians via the Pauli Transfer Matrix. arXiv:2212.04471 [quant-ph]

[13] Juan Castaneda and Nathan Wiebe. 2023. Hamiltonian Learning via Shadow Tomography of Pseudo-Choi States. arXiv:2308.13020 [quant-ph]

[14] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. 2022. Exponential separations between learning with and without quantum memory. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 574–585. doi:10.1109/FOCS52979.2021.00063

[15] Sitan Chen and Weiyuan Gong. 2023. Efficient Pauli channel estimation with logarithmic quantum memory. (9 2023). arXiv:2309.14326 [quant-ph]

[16] Sitan Chen, Weiyuan Gong, and Qi Ye. 2024. Optimal tradeoffs for estimating Pauli observables. *arXiv preprint arXiv:2404.19105* (2024).

[17] Sitan Chen, Jerry Li, and Allen Liu. 2024. An optimal tradeoff between entanglement and copy complexity for state tomography. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 1331–1342.

[18] Senrui Chen, Sisi Zhou, Alireza Seif, and Liang Jiang. 2022. Quantum advantages for Pauli channel estimation. *Phys. Rev. A* 105 (Mar 2022), 032435. Issue 3. doi:10.1103/PhysRevA.105.032435

[19] Thomas Chen, Shivam Nadimpalli, and Henry Yuen. 2023. Testing and Learning Quantum Juntas Nearly Optimally. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1163–1185.

[20] Marcus P. da Silva, Olivier Landon-Cardinal, and David Poulin. 2011. Practical characterization of quantum devices without tomography. *Physical Review Letters* 107, 21 (2011), 210404. doi:10.1103/PhysRevLett.107.210404

[21] Alicja Dutkiewicz, Thomas E. O'Brien, and Thomas Schuster. 2023. The advantage of quantum control in many-body Hamiltonian learning. arXiv:2304.07172 [quant-ph]

[22] Arkopal Dutt, Edwin Pednault, Chai Wah Wu, Sarah Sheldon, John Smolin, Lev Bishop, and Isaac L. Chuang. 2023. Active learning of quantum system Hamiltonians yields query advantage. *Phys. Rev. Res.* 5 (Jul 2023), 033060. Issue 3. doi:10.1103/PhysRevResearch.5.033060

[23] Alexandros Eskenazis, Paata Ivanisvili, and Lauritz Streck. 2022. Low-degree learning and the metric entropy of polynomials. *arXiv preprint arXiv:2203.09659* (2022).

[24] Omar Fawzi, Aadil Oufkir, and Daniel Stilck França. 2023. Lower Bounds on Learning Pauli Channels. arXiv:2301.09192 [quant-ph]

[25] Steven T. Flammia and Joel J. Wallman. 2020. Efficient Estimation of Pauli Channels. *ACM Transactions on Quantum Computing* 1, 1, Article 3 (dec 2020), 32 pages. doi:10.1145/3408039

[26] András Gilyén and Tongyang Li. 2020. Distributional Property Testing in a Quantum World. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 151)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 25:1–25:19. doi:10.4230/LIPIcs.ITCS.2020.25

[27] Parikshit Gopalan, Ryan O'Donnell, Rocco A. Servedio, Amir Shpilka, and Karl Wimmer. 2011. Testing Fourier Dimensionality and Sparsity. *SIAM J. Comput.* 40, 4 (2011), 1075–1100. arXiv:https://doi.org/10.1137/100785429

[28] Christopher E Granade, Christopher Ferrie, Nathan Wiebe, and David G Cory. 2012. Robust online Hamiltonian learning. *New Journal of Physics* 14, 10 (oct 2012), 103013. doi:10.1088/1367-2630/14/10/103013

[29] Andi Gu, Lukasz Cincio, and Patrick J. Coles. 2024. Practical Hamiltonian learning with unitary dynamics and Gibbs states. *Nature Communications* 15, 1 (2024). doi:10.1038/s41467-023-44008-1

[30] Jeongwan Haah, Robin Kothari, and Ewin Tang. 2022. Optimal learning of quantum Hamiltonians from high-temperature Gibbs states. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 135–146. doi:10.1109/FOCS54457.2022.00020

[31] M Holzäpfel, T Baumgratz, M Cramer, and Martin B Plenio. 2015. Scalable reconstruction of unitary processes and Hamiltonians. *Physical Review A* 91, 4 (2015), 042129.

[32] Hsin-Yuan Huang, Sitan Chen, and John Preskill. 2023. Learning to predict arbitrary quantum processes. *PRX Quantum* 4, 4 (2023), 040337. doi:10.1103/PRXQuantum.4.040337

[33] Hsin-Yuan Huang, Yu Tong, Di Fang, and Yuan Su. 2023. Learning many-body Hamiltonians with Heisenberg-limited scaling. *Physical Review Letters* 130, 20 (2023), 200403. doi:10.1103/PhysRevLett.130.200403

[34] Luca Innocenti, Leonardo Banchi, Alessandro Ferraro, Sougato Bose, and Mauro Paternostro. 2020. Supervised learning of time-independent Hamiltonians for gate design. *New Journal of Physics* 22, 6 (2020), 065001.

[35] Robbie King, David Gosset, Robin Kothari, and Ryan Babbush. 2024. Triply efficient shadow tomography. *arXiv preprint arXiv:2404.19211* (2024).

[36] Margarite L LaBorde and Mark M Wilde. 2022. Quantum algorithms for testing Hamiltonian symmetry. *Physical Review Letters* 129, 16 (2022), 160503.

[37] Haoya Li, Yu Tong, Hongkang Ni, Tuvia Gefen, and Lexing Ying. 2023. Heisenberg-limited Hamiltonian learning for interacting bosons. arXiv:2307.04690 [quant-ph]

[38] Ashley Montanaro and Tobias J Osborne. 2008. Quantum boolean functions. *arXiv preprint arXiv:0810.2435* (2008).

[39] Tim Möbus, Andreas Bluhm, Matthias C. Caro, Albert H. Werner, and Cambyse Rouzé. 2023. Dissipation-enabled bosonic Hamiltonian learning via new information-propagation bounds. arXiv:2307.15026 [quant-ph]

[40] Sahand Negahban and Devavrat Shah. 2012. Learning sparse boolean polynomials. In *2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2032–2036.

[41] Emilio Onorati, Cambyse Rouzé, Daniel Stilck França, and James D. Watson. 2023. Efficient learning of ground & thermal states within phases of matter. arXiv:2301.12946 [quant-ph]

[42] Cambyse Rouzé and Daniel Stilck França. 2023. Learning quantum many-body systems from a few copies. arXiv:2107.03333 [quant-ph]

[43] Adrian She and Henry Yuen. 2023. Unitary Property Testing Lower Bounds by Polynomials. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 251)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 96:1–96:17. doi:10.4230/LIPIcs.ITCS.2023.96

[44] Sarah Sheldon, Easwar Magesan, Jerry M. Chow, and Jay M. Gambetta. 2016. Procedure for systematically tuning up cross-talk in the cross-resonance gate. *Phys. Rev. A* 93 (Jun 2016), 060302(R). Issue 6. doi:10.1103/PhysRevA.93.060302

[45] Daniel Stilck França, Liubov A. Markovich, V. V. Dobrovitski, Albert H. Werner, and Johannes Borregaard. 2024. Efficient and robust estimation of many-qubit Hamiltonians. *Nature Communications* 15 (2024), 311. doi:10.1038/s41467-023-44012-5

[46] Alexander Volberg and Haonan Zhang. 2023. Noncommutative Bohnenblust–Hille inequalities. *Math. Ann.* (2023), 1–20. doi:10.1007/s00208-023-02680-0

[47] Frederik Wilde, Augustine Kshetrimayum, Ingo Roth, Dominik Hangleiter, Ryan Sweke, and Jens Eisert. 2022. Scalably learning quantum many-body Hamiltonians from dynamical data. arXiv:2209.14328 [quant-ph]

[48] Grigory Yaroslavtsev and Samson Zhou. 2020. Fast Fourier Sparsity Testing. In *Symposium on Simplicity in Algorithms*. SIAM, 57–68.

[49] Nengkun Yu. 2021. Sample Efficient Identity Testing and Independence Testing of Quantum States. In *12th Innovations in Theoretical Computer Science Conference (ITCS 2021) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 185)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 11:1–11:20. doi:10.4230/LIPIcs.ITCS.2021.11

[50] Wenjun Yu, Jinzhao Sun, Zeyao Han, and Xiao Yuan. 2023. Robust and efficient Hamiltonian learning. *Quantum* 7 (2023), 1045. doi:10.22331/q-2023-06-29-1045

[51] Assaf Zubida, Elad Yitzhaki, Netanel H. Lindner, and Eyal Bairey. 2021. Optimal short-time measurements for Hamiltonian learning. arXiv:2108.08824 [quant-ph]

Received 2024-11-04; accepted 2025-02-01