

Heisenberg-Limited Hamiltonian Learning Without Short-Time Control

Myeongjin Shin^{1,*}, Junseo Lee^{2,†,§} and Changhun Oh^{3,‡}

¹*School of Computational Sciences, Korea Advanced Institute of Science and Technology, Daejeon 34141, Korea*

²*Institute of Computer Technology, Seoul National University, Seoul 08826, Korea*

³*Department of Physics, Korea Advanced Institute of Science and Technology, Daejeon 34141, Korea*



(Received 30 April 2026; revised 23 July 2026; accepted 17 September 2026; published 16 October 2026)

Characterizing quantum systems by learning their underlying Hamiltonians is a central task in quantum information science. While recent algorithmic advances have achieved near-optimal efficiency in this task, they critically rely on accessing arbitrarily short-time dynamics. This reliance poses severe experimental challenges due to finite control bandwidth and transient pulse errors. In this work, we demonstrate that Heisenberg-limited Hamiltonian learning can be achieved without short-time control. We introduce a framework in which every query to the unknown dynamics has duration at least a prescribed minimum time T , and show that this restriction does not preclude Heisenberg-limited scaling. The key ingredient is a method for emulating the continuous quantum control required by iterative learning algorithms using only such lower-bounded evolution times. This reduces the learning task to sparse pure-state tomography. Notably, for logarithmically sparse Hamiltonians, our algorithm achieves the information-theoretically optimal $1/\epsilon$ scaling in total evolution time for *any* arbitrary constant minimum evolution time T . For many-body (polynomially sparse) systems, we uncover a rigorous quantitative tradeoff, showing that the minimum required evolution time can be significantly relaxed from the standard limit at a polynomial cost in total evolution time. Our results affirmatively resolve a prominent open problem in the field and reveal that high-bandwidth, ultra-short pulses are not fundamentally necessary for optimal quantum learning.

DOI: [10.1103/kjvf-jgqj](https://doi.org/10.1103/kjvf-jgqj)

I. INTRODUCTION

Understanding how interactions govern physical systems is a central goal of physics. In quantum systems, these interactions are encoded in the Hamiltonian, which dictates the system's dynamical behavior. Precisely characterizing this Hamiltonian from observed dynamics has become a fundamental problem of broad interest, with applications ranging from quantum metrology to quantum device benchmarking and the verification of analog quantum simulators [1–4]. A central question is therefore whether one can efficiently learn the interactions governing an unknown quantum system directly from its time dynamics, and under what control assumptions this is possible.

The problem of learning an unknown Hamiltonian from time evolution has been extensively studied [1–16]. A standard theoretical framework assumes that the Hamiltonian is m -sparse in the Pauli basis,

$$H = \sum_{x=1}^m \alpha_x P_x, \quad (1)$$

where $\{P_x\}$ are Pauli operators and $\{\alpha_x\}$ are unknown nonzero coefficients. Given oracle access to the time-evolution operator e^{-iHt} at chosen time durations $\{t_i\}_{i=1}^{\mathcal{N}}$, the goal is to output an estimate $\tilde{H}$ that approximates H to the desired accuracy. To evaluate the efficiency of learning algorithms in this setting, two natural complexity measures are commonly considered:

$$t_{\text{tot}} = \sum_{i=1}^{\mathcal{N}} t_i, \quad t_{\text{min}} = \min_{i \in [\mathcal{N}]} t_i, \quad (2)$$

where $\mathcal{N}$ denotes the number of time-evolution queries used by the algorithm. While t_{tot} captures the overall dynamical cost of the protocol, t_{min} quantifies the shortest timescale that the learner must physically access, and is therefore the quantity most directly tied to experimental

*Contact author: hanwoolmj@kaist.ac.kr

†Contact author: junseolee@fas.harvard.edu

‡Contact author: changhun0218@gmail.com

§Present address: Harvard University.

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.

control constraints. Our focus in this work is precisely to understand whether Heisenberg-limited Hamiltonian learning remains possible when $t_{\min}$ is lower bounded by an arbitrary fixed constant $T > 0$.

Recent works have achieved strong guarantees for sparse Hamiltonian learning, including nearly optimal dependence on the target precision in terms of total evolution time [8,9,17]. A key requirement of these state-of-the-art approaches is access to short-time discrete control, namely, the ability to access sufficiently fine-grained time evolution, including arbitrarily short-time dynamics. This capability is crucial in iterative refinement schemes: given a current estimate H_j of the unknown Hamiltonian H , the learner seeks to isolate the remaining unknown interactions by simulating evolution under the residual Hamiltonian $\Delta H_j = H - H_j$. This requirement arises from the need to implement continuous quantum control $e^{-i\Delta H_j t}$ through multiple applications of discrete control, as learning is fundamentally limited to the standard quantum limit in the absence of discrete control [14].

To synthesize this target dynamics, standard approaches combine queries to the physical evolution e^{-iHt} with simulation of the current estimate H_j , and rely on Trotterization [18–20]:

$$e^{-i\Delta H_j t} \simeq (e^{-iHt/N} e^{iH_j t/N})^N. \quad (3)$$

For this product formula to accurately approximate the target evolution, the number of segments N must be sufficiently large. As a result, the elementary query time t/N must become correspondingly small, so access to vanishingly short-time dynamics is built directly into the simulation primitive. In particular, previous algorithms require the minimum evolution time to scale as $t_{\min} = \Theta(1/m)$ [17] or $t_{\min} = \Theta(\sqrt{\epsilon})$ [8,9], so the accessible time window must shrink as the sparsity m grows or the target accuracy ϵ is tightened.

Implementing increasingly short-time evolutions poses a generic challenge in experimental platforms where the target Hamiltonian is established through finite-bandwidth controls [21–23]. The physical control parameters cannot generally be switched instantaneously from an idle configuration to the desired operating point. Instead, the system undergoes a finite transient before reaching an approximately stationary evolution under the target Hamiltonian. When the intended evolution window becomes comparable to or shorter than this response or settling time, the transient can constitute a substantial fraction of the implemented dynamics, and a nominal pulse of duration τ need not realize a clean evolution $e^{-iH\tau}$. This issue is particularly relevant to Trotterization-based Hamiltonian-learning protocols, in which such short unknown-evolution segments are repeatedly interleaved with additional controls.

Recent superconducting experiments provide concrete evidence for this limitation. In flux-tunable qubits and

couplers, control-line distortions can produce delayed response, overshoot, ringing, and longer-time drift before the device settles at the target operating point. Recent in-situ measurements found that, without compensation, a superconducting qubit reached its target frequency only after approximately 50 ns, and characterized distortions over timescales ranging from nanoseconds to tens of microseconds [24]. In a Hamiltonian-identification experiment on Sycamore processors, finite parameter ramps generated nontrivial dynamics, constituted the dominant source of SPAM errors, and were identified as a severe limitation of the architecture [25]. A recent 69-qubit analog-digital superconducting-simulator experiment likewise identified Hamiltonian ramps as a potential vulnerability of sophisticated Hamiltonian-learning techniques [26]. These observations highlight an operational distinction between implementing an isolated stationary evolution of duration τ and resolving the same additive duration difference between two longer evolutions of durations T and $T + \tau$. If T exceeds the settling time, both long queries may contain an established plateau under the target Hamiltonian, and the $T + \tau$ query may be realized by retaining the same calibrated turn-on and turn-off procedures while extending the plateau by τ . Our access model is intended to capture this regime. It does not eliminate switching transients or the need to distinguish T from $T + \tau$, but it avoids requiring a stationary-Hamiltonian evolution window whose duration vanishes as the target accuracy is increased. This motivates the central question considered here:

These physical constraints strongly motivate studying Hamiltonian learning in a fundamentally different regime: one in which *short-time evolution* is not accessible. Concretely, a setting where every available evolution time is lower bounded by a constant T , independent of both the sparsity m and the target accuracy ϵ . This leads to the following central question:

Can one achieve Heisenberg-limited Hamiltonian learning without using short-time evolution?

A. Our contribution and technical overview

In this work, we show that Heisenberg-limited Hamiltonian learning can be achieved without access to arbitrarily short-time dynamics. For logarithmically sparse Hamiltonians, every query to the unknown evolution can have duration at least an arbitrary fixed constant $T > 0$, while the total evolution time still scales as $1/\epsilon$ in the target accuracy and remains efficient in the number of qubits. For polynomially sparse Hamiltonians, we obtain a quantitative tradeoff between the minimum accessible evolution time and the total evolution time. In particular, one can achieve a constant minimum evolution time independent of m at the cost of a quasi-polynomial



Previous approaches: short-time control**This work: *without* short-time control**

FIG. 1. Comparison between previous short-time control approaches and our framework. Previous Heisenberg-limited algorithms rely on stroboscopic control and Trotterization, requiring many ultra-short evolution segments with small $t_{\min}$. Our framework replaces these short-time queries by long-time evolutions interleaved with learned auxiliary operations, ensuring that every query to the unknown Hamiltonian has duration at least T . Thus, the short-time bottleneck is removed while Heisenberg-limited scaling in total evolution time is retained.

overhead in the total evolution time, while still retaining Heisenberg-limited scaling in ϵ (See Fig. 1).

The central difficulty is that existing Heisenberg-limited algorithms rely on residual dynamics. Given a current estimate H_j , one would like to isolate the remaining Hamiltonian $\Delta H_j = H - H_j$ by implementing $e^{-i\Delta H_j t}$. Standard product-formula implementations do this through short steps such as

$$(e^{-iHt/N} e^{iH_j t/N})^N, \quad (4)$$

which directly requires access to $e^{-iHt/N}$. Our main observation is that each such short step can be rewritten as a long-time step plus a correction:

$$e^{-iH\tau} e^{iH_j \tau} = e^{-iH(T+\tau)} C_j e^{iH_j(T+\tau)}, \quad C_j = e^{iHT} e^{-iH_j T}. \quad (5)$$

Thus, all queries to the unknown Hamiltonian have duration at least T .

The correction C_j is not directly accessible because it contains backward evolution under the unknown Hamiltonian. However, its adjoint $C_j^\dagger = e^{iH_j T} e^{-iHT}$ is accessible using only forward long-time evolution under H and simulation of the known Hamiltonian H_j . Our strategy is therefore to learn a Hermitian generator W of

$C_j^\dagger = e^{-iW}$ and use it to implement the inverse correction needed in the long-time product formula. The key technical point is that this auxiliary generator has bounded norm and sparsity, or quasi-sparsity, to be learned and simulated efficiently.

Once residual dynamics can be emulated, the remaining task is coefficient extraction. We reduce this step to sparse pure-state tomography: applying the residual dynamics $e^{-i\Delta H_j t}$ to one half of a maximally entangled state encodes the Pauli coefficients of ΔH_j as first-order amplitudes. By recovering these amplitudes and applying truncation to enforce bounded norm and sparsity, we obtain a refined estimate H_{j+1} .

Iterating this refinement yields the main guarantees of the paper. For logarithmically sparse Hamiltonians, our algorithm allows any fixed minimum evolution time $T > 0$ while retaining Heisenberg-limited scaling in the target accuracy. For polynomially sparse Hamiltonians, it gives an explicit tradeoff between the minimum accessible evolution time and the total evolution time. Thus, the need for ultra-short-time dynamics can be removed entirely in the logarithmically sparse regime and substantially relaxed in the many-body regime. To the best of our knowledge, this is the first work to establish a quantitative tradeoff between the minimum accessible evolution time and the total evolution time in Hamiltonian learning.

Our results resolve an open problem posed by Bakshi *et al.* [17]. More broadly, we suggest a conceptual shift in quantum learning from dynamics: access to ultra-short evolution times is not a fundamental requirement for information-theoretically optimal learning. We now formalize this learning model that avoids short-time evolution and state the corresponding guarantees.

II. PROBLEM SETUP AND MAIN RESULTS

We now formalize the long-time Hamiltonian learning problem and state the main results. The purpose of this section is to specify the Hamiltonian model, the allowed oracle access, and the resource measures used throughout the paper.

A. Hamiltonian model

Let $H \in \mathbb{C}^{2^n \times 2^n}$ be an n -qubit, m -sparse, traceless Hamiltonian of the form

$$H = \sum_{x=1}^m \alpha_x P_x, \quad (6)$$

where $P_x \in \{\mathbb{I}, X, Y, Z\}^{\otimes n}$ are n -qubit Pauli operators and $\alpha_x \in \mathbb{R}$ are unknown nonzero coefficients. The sparsity parameter m reflects the underlying physical structure of the system: the regime $m = \mathcal{O}(\log n)$ characterizes highly structured systems where a few key interactions dominate the dynamics, whereas $m = \text{poly}(n)$ encompasses the broader class of many-body Hamiltonians typically encountered in condensed matter and quantum information science.

B. Access model

We are given oracle access to the time-evolution unitary $U(t) = e^{-iHt}$, for any evolution time $t \geq T$, where $T > 0$ is an arbitrarily large constant. That is, for each query the learner may choose any time $t \geq T$ and obtain access to the corresponding unitary $U(t)$.

C. Complexity measures

We evaluate the efficiency of a learning algorithm using the following quantities:

- (i) the *total evolution time* $t_{\text{tot}} = \sum_{i=1}^{\mathcal{N}} t_i$,
- (ii) the *minimum evolution time* $t_{\min} = \min_{i \in [\mathcal{N}]} t_i$,
- (iii) the *number of queries* $\mathcal{N}$.

Here $t_1, \dots, t_{\mathcal{N}}$ denote the evolution times used by the algorithm. In our setting, the constraint $t_{\min} \geq T$ prohibits access to arbitrarily short-time evolution.

Definitions. Let $H = \sum_{x=1}^m \alpha_x P_x$ be an n -qubit Hamiltonian with eigenvalues $\{\lambda_j\}_{j=1}^{2^n}$. We use the following notions of norm and sparsity:

- (i) Matrix norms:

$$\|H\|_{\infty} := \max_{j \in [2^n]} |\lambda_j|, \quad (7)$$

$$\|H\|_F := \sqrt{\frac{1}{2^n} \sum_{j=1}^{2^n} \lambda_j^2} = \sqrt{\sum_x \alpha_x^2}. \quad (8)$$

Here $\|\cdot\|_{\infty}$ denotes the operator norm, and $\|\cdot\|_F$ denotes the normalized Frobenius norm.

- (ii) Unitary distance:

$$d(U, V) := \min_{\phi} \|U - e^{i\phi} V\|_F. \quad (9)$$

Here $d(U, V)$ denotes the distance between unitaries up to a global phase.

- (iii) (Pauli) coefficient norms: for $p \geq 1$,

$$\|H\|_{\ell_p} := \left(\sum_x |\alpha_x|^p \right)^{1/p}, \quad \|H\|_{\ell_{\infty}} := \max_x |\alpha_x|. \quad (10)$$

- (iv) (Pauli) sparsity:

$$\text{supp}_P(H) := |\{x: \alpha_x \neq 0\}|. \quad (11)$$

In particular, if H is given as in Eq. (6), then $\text{supp}_P(H) = m$.

- (v) (Pauli) quasi-sparsity: for $\varepsilon \geq 0$,

$$\text{supp}_P(H, \varepsilon) := \min \{ \text{supp}_P(\tilde{H}) : \|H - \tilde{H}\|_F \leq \varepsilon \}. \quad (12)$$

Combining the above ingredients, we arrive at the central problem studied in this work.

Problem 1. (Hamiltonian learning with minimum evolution time T). Let H be an n -qubit, m -sparse, traceless Hamiltonian of the form $H = \sum_{x=1}^m \alpha_x P_x$. Suppose we are given oracle access to the time-evolution unitary $U(t) = e^{-iHt}$, for any evolution time $t \geq T$, where $T > 0$ is an arbitrarily large constant. The goal is to output an estimate $\tilde{H}$ such that $\|H - \tilde{H}\|_{\ell_{\infty}} < \varepsilon$.

We say that the problem admits an efficient learning algorithm if the total evolution time is $\text{poly}(n)/\varepsilon$, the number of queries is $\text{poly}(n, 1/\varepsilon)$, and the classical post-processing time is $\text{poly}(n, 1/\varepsilon)$.

This formalization captures a realistic experimental regime in which the learner is restricted to accessing only long-time dynamics, thereby avoiding transient control errors associated with ultra-short evolutions. The key challenge is to extract high-precision information—namely, the individual interaction coefficients $\{\alpha_x\}$ —from such coarse, long-duration evolutions. The following

results show that this is possible in logarithmically sparse systems and remains possible for non short regimes in T , with a quantitative tradeoff, for polynomially sparse systems.

D. Main results

We now state our main guarantees for Hamiltonian learning under the minimum-evolution-time constraint in Problem 1. The first result applies to logarithmically sparse Hamiltonians. In this regime, an exponential dependence on the sparsity m remains efficient in the system size, and the minimum accessible evolution time can be chosen as an arbitrary fixed constant.

Theorem 1. (Learning logarithmically sparse Hamiltonians from arbitrary long-time evolution). Let H be an n -qubit, m -sparse traceless Hamiltonian with $\|H\|_\infty \leq 1$. There exists a learning algorithm that solves Problem 1 with access to e^{-iHt} for all $t \geq T$, using total evolution time

$$t_{\text{tot}} = \tilde{\mathcal{O}}\left(\min\left\{\frac{4^m T^3}{\varepsilon}, \frac{4^m T}{\varepsilon^2}\right\}\right), \quad (13)$$

and minimum evolution time $t_{\min} = T$, with high probability. In particular, when $m = \mathcal{O}(\log n)$, this gives an efficient algorithm for any fixed constant $T > 0$, achieving Heisenberg-limited scaling in the target accuracy while never querying the unknown dynamics for times shorter than T .

The overall complexity is given by the minimum of the Heisenberg-limited and standard quantum limit scalings. In the small-error regime, the Heisenberg-limited scaling $1/\varepsilon$ dominates. In the complementary regime, the $1/\varepsilon^2$ scaling applies; however, since the error ε is bounded below by a constant independent of the desired accuracy, the total evolution time in this regime is effectively a constant in terms of ε . Consequently, combining both regimes, the overall complexity is governed by the Heisenberg-limited scaling.

The next result extends the framework to many-body Hamiltonians with polynomial sparsity. Here the minimum accessible evolution time can still be made substantially larger than the $1/m$ -scale required by standard product-formula approaches, at the cost of an increased polynomial dependence on m .

Theorem 2. (Learning many-body Hamiltonians without short-time evolution). Let H be an n -qubit, m -sparse traceless Hamiltonian with $\|H\|_\infty \leq 1$. There exists a learning algorithm that solves Problem 1 with access to e^{-iHt} for all $t \geq T$, using total evolution time

$$t_{\text{tot}} = \tilde{\mathcal{O}}\left(\min\left\{\frac{m^{K+2}T}{\varepsilon}, \frac{m^K T}{\varepsilon^2}\right\}\right), \quad (14)$$

provided that the minimum evolution time satisfies $t_{\min} = T = \Theta(m^{-1/K})$, where $K \in \mathbb{N}$. The algorithm succeeds with high probability.

Moreover, choosing $K = \Theta(\log m)$ makes the minimum evolution time constant, $t_{\min} = \Theta(1)$, at the expense of quasi-polynomial total evolution time,

$$t_{\text{tot}} = \tilde{\mathcal{O}}\left(\min\left\{\frac{m^{2+\log m}}{\varepsilon}, \frac{m^{\log m}}{\varepsilon^2}\right\}\right). \quad (15)$$

We emphasize that this constant-time guarantee is qualitatively different from that of Theorem 1. More explicitly, under the normalization $\|H\|_\infty, \|H_j\|_\infty \leq 1$, the parameter choice used in the proof is

$$t_{\min} = T = \frac{m^{-1/K}}{16e}.$$

Therefore,

$$t_{\min} \rightarrow \frac{1}{16e} \quad \text{as } K \rightarrow \infty.$$

Thus, under this formalism, the constant minimum evolution time guaranteed by Theorem 2 is at most $1/(16e)$, even with quasi-polynomial total evolution time. In contrast, Theorem 1 allows $t_{\min} = T$ for any arbitrarily prescribed constant $T > 0$.

We emphasize that large minimum evolution time does not by itself relax the required timing resolution. In the long-time control-emulation step, the unknown Hamiltonian is queried at times T and $T + \Delta t_j$, where

$$\Delta t_j = \frac{t_j}{N_j} = \frac{1}{32m\eta_j[1/(2\sqrt{m\eta_j})]} = \Theta\left(\frac{1}{\sqrt{m}}\right).$$

The protocol therefore still requires timing control with resolution Δt_j . However, this required resolution does not become finer as the target accuracy ε decreases, although it may still be demanding for large m . Operationally, the protocol does not require the device to realize $e^{-iH\Delta t_j}$ as a stand-alone short evolution. Instead, it requires the duration of an already established long-time evolution to be controlled accurately enough to distinguish T from $T + \Delta t_j$. Our analysis assumes exact access to these prescribed evolution times and does not account for timing miscalibration or jitter.

These results show that the minimum accessible evolution time and the total evolution time are tradable resources in Hamiltonian learning. For logarithmically sparse Hamiltonians, the short-time requirement can be removed entirely while retaining Heisenberg-limited scaling. For polynomially sparse Hamiltonians, our bounds

give an explicit tradeoff between increasing $t_{\min}$ and increasing t_{tot} . In particular, they allow constant $t_{\min}$ at the expense of quasi-polynomial total evolution time. The main remaining open question is whether polynomially sparse Hamiltonians can be learned with polynomial total evolution time while allowing $t_{\min}$ to be an arbitrary fixed constant independent of m . We next describe the learning architecture that proves Theorems 1 and 2.

III. ALGORITHM OVERVIEW

We now describe how the main algorithm is organized, which is illustrated in Fig. 2. The protocol follows an iterative refinement scheme. Starting from $H_0 = 0$, the learner maintains an estimate H_j with

$$\|H - H_j\|_{\ell_\infty} \leq \eta_j,$$

and aims to produce H_{j+1} with $\eta_{j+1} = \eta_j/2$. Thus each round reduces to learning the residual Hamiltonian

$$\Delta H_j = H - H_j$$

to accuracy comparable to η_j .

Each refinement round has two conceptual steps. First, the learner emulates the residual evolution $e^{-i\Delta H_j t_j}$ using only queries to the unknown dynamics e^{-iHt} with $t \geq T$, together with simulations of known Hamiltonians. Second, the learner converts this residual evolution into estimates of the Pauli coefficients of ΔH_j by applying it to one half of a maximally entangled state and performing sparse pure-state tomography.

Schematically, one round takes the form

$$H_j \xrightarrow{\text{long-time emulation}} \tilde{U}_{\text{res}}^{(j)} \approx e^{-i(H-H_j)t_j} \xrightarrow{\text{sparse tomography}} \widetilde{\Delta H_j} \xrightarrow{\text{truncate}} H_{j+1}. \quad (16)$$

The truncation step keeps the estimate m -sparse and norm bounded, ensuring that the next round satisfies the same structural promises.

The key new ingredient is the first step. Instead of implementing residual dynamics through short product-formula steps, we rewrite each short step as a long-time query plus a correction unitary. This correction is implemented by learning a structured generator W_j from the accessible adjoint dynamics $e^{iH_j T} e^{-iHT}$. The norm and




FIG. 2. Schematic overview of our iterative learning protocol. At iteration j , the current estimate H_j defines a residual Hamiltonian $\Delta H_j = H - H_j$. We first learn an auxiliary Hamiltonian W_j using only long-time evolutions under H , and then use W_j to emulate the residual-control primitive $e^{-i\Delta H_j t}$. This emulated evolution reduces coefficient extraction to sparse pure-state tomography. After estimating ΔH_j , we update the estimate and apply the sparse bounded truncation map $\mathcal{T}_{k,c}$, which projects onto the set of k -sparse Hamiltonians with coefficient magnitude at most c by minimizing the ℓ_∞ -distance, thereby preserving sparsity and bounded coefficients while approximately halving the estimation error at each iteration.

sparsity, or quasi-sparsity, of W_j determine the final complexity.

The algorithm uses this Heisenberg-limited refinement once the current error is sufficiently small. For earlier rounds, where the residual error is large, we use a simpler standard-quantum-limit subroutine. The full protocol and the crossover between the two regimes are given in Sec. VI. We first develop the long-time emulation primitive, which is the main technical contribution.

IV. LONG-TIME EMULATION OF RESIDUAL DYNAMICS

We now construct the first ingredient of the learning algorithm: a long-time implementation of the residual evolution

$$e^{-i\Delta H_j t} = e^{-i(H-H_j)t}, \quad (17)$$

given a current estimate H_j . Throughout this section, H and H_j are m -sparse and satisfy $\|H - H_j\|_{\ell_\infty} \leq \varepsilon$, while the learner may query the unknown dynamics only for times $\tau \geq T$.

This primitive is essential for isolating the remaining unknown coefficients during iterative refinement, but standard Hamiltonian-learning protocols implement it using short-time queries to the unknown dynamics. Our goal in this section is to construct the same primitive using exclusively long-time queries $e^{-iH\tau}$ with $\tau \geq T$.

A. Standard approach: Trotterization

The conventional method for approximating continuous quantum control relies on Trotterization [18–20]:

$$\begin{aligned} & \| (e^{-iHt/N} e^{iH_j t/N})^N - e^{-i(H-H_j)t} \|_F \\ & \leq \frac{t^2 \min\{\|H\|_\infty, \|H_j\|_\infty\}}{N} \|H - H_j\|_F. \end{aligned} \quad (18)$$

This approximation simulates $e^{-i(H-H_j)t}$ by repeatedly applying the discrete control operation

$$e^{-iHt/N} e^{iH_j t/N}. \quad (19)$$

However, this approach intrinsically requires access to short-time evolution $e^{-iHt/N}$, which violates the long-time evolution constraint considered in this work.

This is precisely the short-time bottleneck inherited by previous Heisenberg-limited Hamiltonian-learning algorithms. In an iterative learning scheme, one repeatedly constructs a current estimate H_j and then uses residual dynamics generated by $H - H_j$ to refine the estimate. When this residual dynamics is implemented through the product formula above, every refinement step requires many short applications of the unknown evolution $e^{-iHt/N}$,

interleaved with simulations of the known estimate H_j . Consequently, the minimum accessible evolution time becomes tied to the product-formula step size. As the desired accuracy improves, or as the sparsity-dependent control error becomes more stringent, the required step size must shrink. This is the mechanism behind the ultra-short-time requirements in existing approaches.

Our goal is to break this link. We show that each short-time product-formula step can be replaced by a long-time query to the unknown Hamiltonian, together with a correction unitary that is itself accessible through long-time dynamics.

B. Our approach: Emulating short-time steps via long-time dynamics

To bypass the need for short-time evolution, we rewrite each elementary product-formula step as a long-time evolution interrupted by a fixed correction unitary. Define

$$C_j = e^{iHT} e^{-iH_j T}. \quad (20)$$

Then, for any $\tau \geq 0$,

$$e^{-iH\tau} e^{iH_j \tau} = e^{-iH(T+\tau)} C_j e^{iH_j(T+\tau)}. \quad (21)$$

Indeed, substituting the definition of C_j cancels the additional evolutions of duration T . Setting $\tau = t/N$, we obtain the exact rewriting

$$(e^{-iHt/N} e^{iH_j t/N})^N = (e^{-iH(T+t/N)} C_j e^{iH_j(T+t/N)})^N. \quad (22)$$

Crucially, every query to the unknown evolution on the right-hand side has duration at least T . Thus, if the correction unitary C_j could be implemented, the usual product-formula approximation to $e^{-i(H-H_j)t}$ would require no short-time evolution under H .

At first sight, this rewriting only shifts the difficulty from short-time evolution to the implementation of the correction unitary C_j . The useful viewpoint is to regard C_j not as a black-box unitary to be estimated directly, but as a correction to be implemented through a learned Hermitian generator. That is, we seek a traceless Hermitian matrix W_j satisfying $d(C_j, e^{iW_j}) = 0$. If such a generator W_j is known explicitly, then the missing correction C_j can be implemented by simulating the known Hamiltonian W_j to realize e^{iW_j} . In this sense, the problem of removing short-time queries is reduced to the problem of learning a suitable generator W_j for the correction unitary.

In the actual learning protocol, we only construct a finite-precision estimate $\tilde{W}_j$ of W_j . Suppose that $\tilde{W}_j$ satisfies

$$\|W_j - \tilde{W}_j\|_F \leq \frac{\varepsilon'}{2N}. \quad (23)$$

Then replacing $C_j = e^{iW_j}$ by $e^{i\tilde{W}_j}$ in the long-time product formula introduces at most ε' error over the N product-formula steps. By choosing N sufficiently large such that the right-hand side of Eq. (18) is at most $\varepsilon'/2$, and by learning W_j to the corresponding accuracy, the product-formula and replacement errors can be jointly bounded by ε' :

$$\begin{aligned} d(e^{-i(H-H_j)t}, (e^{-iH(T+t/N)} e^{i\tilde{W}_j} e^{iH_j(T+t/N)})^N) \\ \leq \frac{\varepsilon'}{2} + \frac{\varepsilon'}{2N} \cdot N = \varepsilon'. \end{aligned} \quad (24)$$

This approximate continuous quantum control primitive will be used as a subroutine in the main Hamiltonian learning algorithm (Algorithm 3).

The usefulness of this reduction depends on whether the generator W_j can be learned and simulated efficiently. The generator W_j could be dense in the Pauli basis, have large effective complexity, or be difficult to simulate. The main technical content of our construction is to show that one can work with a generator W_j whose norm and Pauli complexity are controlled well enough for both learning and simulation. We will show below that, in the logarithmically sparse regime, W_j can be supported on a Pauli space of size at most 4^m . In the polynomially sparse regime, W_j may be dense, but it admits a quasi-sparse approximation obtained from the Baker-Campbell-Hausdorff (BCH) expansion. These structural properties are what make the formal rewriting above algorithmically useful.

It remains to explain how this generator W_j can be learned using only the allowed long-time oracle access. This is where the distinction between C_j and $C_j^\dagger$ becomes important.

C. Learning the correction generator W_j

We now explain how the correction generator W_j can be learned using only the allowed long-time evolution oracle access. Recall that W is chosen such that $d(C_j, e^{iW_j}) = 0$. Here, a main difficulty is that C_j itself is not directly accessible because it requires the backward evolution e^{iHT} under the unknown Hamiltonian H . Our key observation to overcome this obstacle is that its adjoint $C_j^\dagger = e^{iH_j T} e^{-iHT}$ is accessible using only forward evolution under H and simulation of the known Hamiltonian H_j . Since $d(C_j, e^{iW_j}) = 0$, the same generator satisfies $d(C_j^\dagger, e^{-iW_j}) = 0$. Hence, for every positive integer q , we can implement, up to a global phase,

$$e^{-iW_j q} = (C_j^\dagger)^q = (e^{iH_j T} e^{-iHT})^q. \quad (25)$$

Thus, long-time evolution access to H induces integer-time evolution access to the auxiliary Hamiltonian W_j . This is precisely the type of access used in Algorithm 1: by querying $e^{-iW_j q}$ for integer q , we learn an approximation $\tilde{W}_j$ in normalized Frobenius norm.

Importantly, this procedure does not violate the minimum-evolution-time constraint. Each application of $C_j^\dagger$ uses one query to e^{-iHT} , whose duration is exactly T . Implementing $e^{-iW_j q}$ simply repeats this allowed long-time query q times, so the minimum query time to the unknown dynamics remains $t_{\min} = T$.

The efficiency of this learning step is governed by the norm and sparsity of W_j . Moreover, the output $\tilde{W}_j$ must be produced in a sparse or quasi-sparse Pauli representation, since this is the Hamiltonian that will later be simulated to implement the correction.

D. Implementing inverse time evolution of $\tilde{W}_j$

For a known Hamiltonian $\tilde{W}_j$, implementing $e^{i\tilde{W}_j}$ is a standard task in Hamiltonian simulation. The circuit complexity of Hamiltonian simulation algorithms typically

ALGORITHM 1. INTEGEREVOL($\mathcal{U}_A, s, c_F, c_\infty, c, \varepsilon$): Hamiltonian learning in normalized Frobenius norm via pure-state tomography using integer time evolution.

Input: Query access to $\mathcal{U}_A(\tau) = e^{-iA\tau}$, $\tau \in \mathbb{N}$ for a traceless Hermitian operator $A = \sum_x \alpha_x P_x$ satisfying $\|A\|_F \leq c_F \varepsilon$ and $\|A\|_\infty \leq c_\infty \varepsilon$; target constant c ; effective support size $s = \text{supp}_P(A, c\varepsilon/8)$.
Output: An estimate $\tilde{A}$ such that $\|A - \tilde{A}\|_F \leq c\varepsilon$.
1 $t \leftarrow \lfloor c/(10c_F c_\infty \varepsilon) \rfloor$
2 $\delta_t \leftarrow c_F c_\infty t^2 \varepsilon^2$
3 Prepare the maximally entangled state $|\Phi\rangle$
4 Let V be a unitary such that $V(P_x \otimes I)|\Phi\rangle = |x\rangle$ for all x
5 Prepare $|\psi_t\rangle \leftarrow V(\mathcal{U}_A(t) \otimes I)|\Phi\rangle$
6 Apply sparse pure-state tomography (Lemma 5) to learn $|\psi_t\rangle$ with accuracy δ_t using $\mathcal{O}(s/\delta_t^2)$ copies, and obtain coefficients $\{\tilde{\alpha}_x\}_x$ satisfying $(\sum_x |\alpha_x - \tilde{\alpha}_x|^2)^{1/2} \leq c\varepsilon$.
7 return $\tilde{A} \leftarrow \sum_x \tilde{\alpha}_x P_x$

TABLE I. Bounds on the correction generator W_j in the two sparsity regimes considered in this work. Here, m denotes the sparsity of H and H_j , ε satisfies $\|H - H_j\|_{\ell_\infty} \leq \varepsilon$, and $K \in \mathbb{N}$.

m	$t_{\min}$	Properties of W_j		
		Upper bound on $\ W_j\ _F$	Upper bound on $\ W_j\ _\infty$	(Quasi-)sparsity
$\mathcal{O}(\log n)$	$\forall T > 0$	$2\pi T \varepsilon \sqrt{m}$	$2\pi T \varepsilon m$	4^m
$\text{poly}(n)$	$\Theta(m^{-1/K})$	$\varepsilon \sqrt{m}$	εm	$\mathcal{O}(m^{K-1})$

scales linearly with the norm and sparsity of $\tilde{W}_j$ and polylogarithmically with the target accuracy [27–32]. Thus, for the learned correction to be useful, it is not enough to approximate W_j accurately in norm; the estimate $\tilde{W}_j$ must also have bounded sparsity or quasi-sparsity.

Our learning subroutine is designed to satisfy this requirement. It outputs $\tilde{W}_j$ within the sparse or quasi-sparse representation guaranteed by the structural bounds on W_j . Consequently, the cost of simulating $e^{i\tilde{W}_j}$ is governed by the same norm and sparsity parameters that enter the learning of W_j .

In addition, $e^{iH_j(T+t/N)}$ can be implemented efficiently, since H_j is known, m -sparse, and satisfies the same bounded-norm assumptions as the target Hamiltonian. Therefore, the entire sequence in Eq. (24) can be realized using only long-time queries to the unknown Hamiltonian H , together with simulations of known Hamiltonians.

Therefore, in order to efficiently learn W_j and then implement $e^{i\tilde{W}_j}$, it suffices to control the norm and sparsity, or quasi-sparsity, of the correction generator W_j . These structural bounds ensure both that W_j can be learned from integer-time access and that the resulting estimate $\tilde{W}_j$ can be efficiently simulated. Concretely, we need bounds on

$$\|W_j\|_F, \quad \|W_j\|_\infty, \quad \text{supp}_P(W_j), \quad (26)$$

or, in the polynomially sparse regime, on the corresponding quasi-sparsity of W_j . These bounds are summarized in Table I for the two operational regimes relevant to our main results. Notably, they are independent of j . We now turn to their derivation.

E. Bounding the norm and (quasi-)sparsity of the auxiliary Hamiltonian W_j

We now establish the structural bounds on W_j that make the above reduction efficient. There are two distinct regimes of interest. When $m = \mathcal{O}(\log n)$, an exponential dependence on m is still polynomial in the system size, so it suffices to show that W lies in the Pauli algebra generated by the supports of H and H_j . This gives an exact sparsity bound of order 4^m . For polynomially sparse Hamiltonians, this exact bound is no longer efficient, and

we instead prove that W_j admits a quasi-sparse approximation by truncating its BCH expansion. The two arguments are different, but serve the same purpose: they ensure that W_j can be learned and that the learned correction $e^{i\tilde{W}_j}$ can be simulated efficiently.

We first consider the logarithmically sparse regime, where the target Hamiltonian satisfies $m = \mathcal{O}(\log n)$ and the minimum evolution time $T > 0$ is an arbitrary constant. Our goal is to bound the norm and sparsity of the auxiliary generator W_j . In this regime, the sparsity parameter grows only logarithmically with the number of qubits, so an exponential dependence on m remains polynomial in n . Thus, it is sufficient to prove an exact sparsity bound for W , even if that bound scales exponentially in m .

The key observation is that W_j can be chosen inside the finite Pauli algebra generated by the Pauli supports of H and H_j . Indeed, the linear span of these supports over $\mathbb{F}_2$ has dimension at most $2m$, and hence contains at most 4^m Pauli labels. This gives the exact sparsity bound $\text{supp}_P(W_j) \leq 4^m$. The norm of W_j , on the other hand, is controlled by Duhamel’s principle, which relates the correction unitary $e^{iHT}e^{-iH_jT}$ to the residual Hamiltonian $H - H_j$.

Lemma 1. (Informal, see Lemma 11 for details). Let H, H_j be n -qubit, m -sparse traceless Hamiltonians satisfying $\|H - H_j\|_{\ell_\infty} \leq \varepsilon$. Then there exists an n -qubit traceless Hamiltonian W_j such that $d(e^{iHT}e^{-iH_jT}, e^{iW_j}) = 0$, and $\|W_j\| \leq \pi T \|H - H_j\|$ and $\text{supp}_P(W_j) \leq 4^m$ for any unitarily invariant norm $\|\cdot\|$.

Proof Sketch. We bound the size of the linear span generated by the Pauli supports of H and H_j to bound the sparsity of W_j . Let $C_j = e^{iHT}e^{-iH_jT}$. By Lemma 7, there exists a traceless Hermitian matrix W_j satisfying $d(U, e^{iW_j}) = 0$, where $d(\cdot, \cdot)$ is the unitary distance defined in Eq. (9), and $\|W_j\| \leq \pi \|I - C_j\|$. Applying Duhamel’s principle (Lemma 6) yields $\|I - C_j\| \leq T \|H - H_j\|$, and hence $\|W_j\| \leq \pi T \|H - H_j\|$.

For the sparsity bound, let S_H and S_{H_j} denote the Pauli supports of H and H_j , and define $L := \text{span}_{\mathbb{F}_2}(S_H \cup S_{H_j}) \subseteq \mathbb{F}_2^{2n}$. Then $\text{supp}_P(W_j) \leq |L| \leq 4^m$. Using $\|H - H_j\|_{\ell_\infty} \leq \varepsilon$ together with the m -sparsity of H and H_j , the lemma further implies that $\|W_j\|_\infty \leq 2\pi m T \varepsilon$ and $\|W_j\|_F \leq 2\pi \sqrt{m} T \varepsilon$. ■

We established upper bounds on the norm and sparsity of W_j when the target Hamiltonian is $\mathcal{O}(\log n)$ -sparse, namely

$$\begin{aligned} \|W_j\|_\infty &\leq 2\pi m T \varepsilon, \\ \|W_j\|_F &\leq 2\pi \sqrt{m} T \varepsilon, \\ \text{supp}_P(W_j) &\leq 4^m. \end{aligned} \quad (27)$$

Since these quantities scale polynomially with the number of qubits n (in the regime $m = \mathcal{O}(\log n)$), it follows that W_j can be learned efficiently. We later present an explicit algorithm for learning W_j (Algorithm 1).

We next consider the regime where the Hamiltonian is $\text{poly}(n)$ -sparse. Concretely, we assume $m = \text{poly}(n)$ and $T = \Theta(m^{-1/K})$ for a constant $K \in \mathbb{N}$. Our goal is again to bound the norm and sparsity of the correction generator W_j .

In this regime, the exact support bound $\text{supp}_P(W_j) \leq 4^m$ from the logarithmically sparse case is no longer useful, since it is exponential in the system size. Thus, we cannot hope to learn W_j by treating it as an exactly sparse Pauli Hamiltonian with support 4^m . Instead, it is enough to show that W_j is quasi-sparse at the accuracy required by the learning subroutine: although W_j may be dense, it can be approximated in normalized Frobenius norm by a Hamiltonian with only polynomially many Pauli terms.

We establish this quasi-sparsity using the BCH expansion of the logarithm of $C_j = e^{iHT} e^{-iH_j T}$. When T decreases as $m^{-1/K}$, higher-order commutators are increasingly suppressed. Truncating the BCH series therefore yields an approximation to W_j whose Pauli support grows only polynomially with m , while the truncation error remains below the accuracy needed for learning and for implementing $e^{i\tilde{W}_j}$.

Lemma 2. (Informal, see Lemma 12 for details). Let H, H_j be m -sparse traceless Hamiltonians satisfying $\|H - H_j\|_{\ell_\infty} \leq \varepsilon$. Then there exists a traceless Hermitian matrix W_j such that $e^{iHT} e^{-iH_j T} = e^{iW_j}$, and for any $l \in \mathbb{N}$ we have $\|W_j\| \leq \|H - H_j\|$ and

$$\text{supp}_P\left(W_j, \frac{\varepsilon m^{-\ell/K}}{8}\right) \leq \mathcal{O}(m^{\ell+K/2-1}), \quad (28)$$

for any unitarily invariant norm $\|\cdot\|$, where $T = \Theta(m^{-1/K})$ for $K \in \mathbb{N}$.

Proof Sketch. We express W_j using the BCH expansion (Lemma 8). Since $\|TH\|_\infty + \|TH_j\|_\infty \leq \log 2$, the series converges.

Let $W_j^{(k)}$ denote the k th order term. One can show that $\|W - W_j^{(k)}\|_F = \mathcal{O}(T^k \sqrt{m} \varepsilon)$, while its Pauli sparsity scales as $\mathcal{O}(m^k)$. Truncating the expansion at a sufficiently large order yields an approximation of W_j whose effective

sparsity is polynomial in m , establishing the claimed quasi-sparsity bound. Using $\|H - H_j\|_{\ell_\infty} \leq \varepsilon$ together with the m -sparsity of H and H_j , the lemma further implies that $\|W_j\|_\infty \leq m\varepsilon$ and $\|W_j\|_F \leq \sqrt{m} \varepsilon$. ■

We therefore obtain the bounds

$$\begin{aligned} \|W_j\|_\infty &\leq m\varepsilon, \\ \|W_j\|_F &\leq \sqrt{m} \varepsilon, \\ \text{supp}_P\left(W_j, \frac{\varepsilon}{8\sqrt{m}}\right) &\leq \mathcal{O}(m^{K-1}). \end{aligned} \quad (29)$$

Thus, although W_j need not be exactly sparse, it can be replaced at the required precision by a polynomially sparse Hamiltonian. This is sufficient both for learning W_j using Algorithm 1 and for simulating the learned correction $e^{i\tilde{W}_j}$. The price is the tradeoff $T = \Theta(m^{-1/K})$, $K \in \mathbb{N}$, which relates the minimum accessible evolution time to the polynomial overhead in total evolution time. For every constant $K = \mathcal{O}(1)$, the quasi-sparsity of W_j becomes a polynomial with m .

Since these quantities (norm, sparsity) scale polynomially with the number of qubits n in both regimes, it follows that W_j can be learned efficiently. We now analyze the complexity of the algorithm for learning W_j (Algorithm 1).

F. Complexity of learning the auxiliary Hamiltonian W_j

We have shown that the norm and sparsity of W_j admit polynomial bounds. Under these guarantees, standard Hamiltonian learning algorithms require only polynomial total evolution time. Our Hamiltonian learning algorithm (Algorithm 1) for the auxiliary Hamiltonian W_j likewise achieves polynomial complexity. In this section, we explicitly compute the total evolution time required to estimate W_j .

To this end, we first specify the constant c , which controls the target accuracy. For simplicity, we set $c = \mathcal{O}(1/\sqrt{m})$ (see Appendix E for details). The remaining parameters c_F , c_∞ , and s in Algorithm 1 correspond to upper bounds on $\|W_j\|_F$, $\|W_j\|_\infty$, and the (quasi-)sparsity of W_j , respectively, as derived in the previous section and summarized in Table I.

The parameters in Algorithm 1 are given by

$$t = \left\lfloor \frac{c}{10c_F c_\infty \varepsilon} \right\rfloor, \quad \delta_t = c_F c_\infty t^2 \varepsilon^2. \quad (30)$$

Since t must be a positive integer, we require $\varepsilon \leq c/(10c_F c_\infty) = \eta_{\text{sw}}$. Since t must be a positive integer, Algorithm 1 is applicable only when

$$\varepsilon \leq \eta_{\text{sw}} := \frac{c}{10c_F c_\infty}.$$

In the logarithmically sparse regime, substituting $c = 1/(256\sqrt{m})$, $c_F = 2\pi\sqrt{m}T$, and $c_\infty = 2\pi mT$ gives

$$\eta_{\text{sw}} = \frac{1}{10 \cdot 2^{10} m^2 \pi^2 T^2}.$$

Therefore, when Algorithm 1 is applied in round j with $\varepsilon = \eta_j$, our present guarantee for learning the correction generator W_j , and hence for implementing the long-time control-emulation procedure through this construction, requires

$$\|H - H_j\|_{\ell_\infty} \leq \eta_j \leq \eta_{\text{sw}}.$$

This is a sufficient condition arising from the positive-integer evolution-time choice and the error bounds used in Algorithm 1. Our analysis does not establish that this threshold is necessary, and whether the long-time control-emulation procedure remains valid under a weaker accuracy condition is left as an open question. The complementary regime $\eta_j > \eta_{\text{sw}}$ is handled by the standard-quantum-limit subroutine described in Appendix F.

The algorithm requires $M = \mathcal{O}(s/\delta_t^2)$ queries to $e^{-iW_j t}$. Each query to $e^{-iW_j t}$ corresponds to evolution time $\mathcal{O}(tT)$ of H , so the total evolution time is $MtT = \mathcal{O}(stT/\delta_t^2)$. Substituting the values of the parameters c, c_F, c_∞, s , we obtain the following bounds. For logarithmically sparse Hamiltonians,

$$MtT = \tilde{\mathcal{O}}\left(\frac{4^m T^3}{\varepsilon}\right). \quad (31)$$

In particular, this shows that the minimum accessible evolution time can be chosen as an arbitrary constant. For polynomially sparse Hamiltonians, we obtain

$$MtT = \tilde{\mathcal{O}}\left(\frac{m^{K+2} T}{\varepsilon}\right). \quad (32)$$

This demonstrates that the minimum evolution time T can be relaxed from $\Theta(m^{-1})$ to $\Theta(m^{-1/K})$ at the cost of an increased total evolution time. Overall, our approach reveals a quantitative tradeoff between the minimum evolution time and the total evolution time required to approximate continuous quantum control in Hamiltonian learning.

Our framework learns the unitary C_j in order to approximate continuous quantum control. In principle, this task could also be approached via general unitary channel estimation [33]. However, producing an ε -accurate estimate of an arbitrary n -qubit unitary requires $\mathcal{O}(4^n/\varepsilon)$

queries, which is exponential in the system size even when the underlying Hamiltonian is sparse. In contrast, our framework exploits the specific structure that C_j is generated by the time evolution of a logarithmically sparse Hamiltonian, and therefore achieves the same task using only a polynomial number of queries. This highlights a key advantage of our approach and underscores the importance of leveraging dynamical structure in Hamiltonian learning.

The discussion above constructs the residual-evolution primitive required at each refinement step. Once this primitive is available, the remaining task is to convert residual dynamics into coefficient estimates. This is achieved by a separate reduction to sparse pure-state tomography, which we describe next.

V. COEFFICIENT EXTRACTION BY PURE-STATE TOMOGRAPHY

We now explain how the emulated residual dynamics is converted into coefficient estimates. The basic idea is an amplitude encoding of Pauli coefficients into a sparse quantum state through the first-order expansion of the time evolution operator. Recall that there are two uses of pure-state tomography in our protocol. The auxiliary use, described above, learns the correction generator W_j needed to emulate residual dynamics. The primary use, described in this section, extracts the coefficients of the residual Hamiltonian $\Delta H_j = H - H_j$ once this emulated dynamics is available. Both uses rely on the same first-order amplitude-encoding idea: by applying an evolution $e^{-iA t}$ to a maximally entangled state, the Pauli coefficients of A are encoded in the amplitudes of the resulting quantum state.

We now describe this reduction in a unified form. Let $A = \sum_x \alpha_x P_x$ be a Hamiltonian with $P_0 = I$ and $\alpha_0 = 0$, and define the Pauli sparsity $m = \text{supp}_P(A)$. Suppose that $\|A\|_{\ell_\infty} \leq \varepsilon$ and that we can implement the time-evolution operator $e^{-iA t}$. For sufficiently small t , we have the expansion $e^{-iA t} = I - iA t + \mathcal{O}(t^2 \varepsilon^2)$.

Using this evolution, we prepare the state $|\psi_t\rangle = V(e^{-iA t} \otimes I)|\Phi\rangle$, where $|\Phi\rangle$ is a maximally entangled state and V is a unitary satisfying $V(P_x \otimes I)|\Phi\rangle \mapsto |x\rangle$. A convenient choice is the modified Bell-basis decoding circuit $V := \bigotimes_{j=1}^n [D_{S_j A_j} (H_{S_j} \otimes I_{A_j}) \times \text{CNOT}_{S_j \rightarrow A_j}]$, $D := \text{diag}(1, 1, 1, i)$. D acts on the output of each qubit pair, in the basis order $|00\rangle, |01\rangle, |10\rangle, |11\rangle$, and $P_x = i^{a \cdot b} X^a Z^b$, we have $V(P_x \otimes I)|\Phi\rangle_{SA} = |b\rangle_S |a\rangle_A \equiv |x\rangle$.

Substituting the expansion of $e^{-iA t}$ yields $|\psi_t\rangle \simeq |0\rangle - it \sum_x \alpha_x |x\rangle$. Hence, the Pauli coefficients α_x can be recovered via pure-state tomography of $|\psi_t\rangle$. More precisely, the approximation error satisfies

$$\left\| |\psi_t\rangle - |0\rangle + it \sum_x \alpha_x |x\rangle \right\|_{\ell_\infty} \leq mt^2 \|A\|_{\ell_\infty}^2 \leq mt^2 \varepsilon^2. \quad (33)$$

Learning $|\psi_t\rangle$ to accuracy $\mathcal{O}(mt^2\varepsilon^2)$ yields an estimate of $t\alpha_x$ with the same accuracy, and hence an estimate of α_x with error $\mathcal{O}(mt\varepsilon^2)$. Choosing $t = \Theta(1/(m\varepsilon))$ ensures this error is $\mathcal{O}(\varepsilon)$. Therefore, it suffices to learn $|\psi_t\rangle$ to accuracy $\mathcal{O}(1/m)$.

Since A is (quasi-)sparse, the state $|\psi_t\rangle$ is also (quasi-)sparse with sparsity parameter m . Sparse pure-state tomography with accuracy $\mathcal{O}(1/m)$ requires $S = \mathcal{O}(m^3)$ samples, yielding a total evolution time of $\mathcal{O}(St) = \mathcal{O}(m^2/\varepsilon)$, which matches the Heisenberg limit. A detailed analysis, including extensions to other norms such as the normalized Frobenius norm, is provided in Appendix C. We now combine this coefficient-extraction primitive with the long-time emulation primitive to obtain the full iterative learner and its complexity.

VI. FULL LEARNING PROTOCOL AND COMPLEXITY

We now assemble the components from the previous sections into the full Hamiltonian learning protocol and derive the complexity bounds in Theorems 1 and 2. Recall from the algorithmic overview that the protocol follows an iterative refinement strategy. At the beginning of round j , we have an estimate H_j satisfying

$$\|H - H_j\|_{\ell_\infty} \leq \eta_j. \quad (34)$$

The round learns the residual Hamiltonian

$$\Delta H_j := H - H_j \quad (35)$$

to accuracy proportional to η_j , and then updates H_j so that the error is reduced by a constant factor.

In the small-error regime, one refinement round has the concrete form

$$\begin{array}{c} H_j \xrightarrow{\text{INTEGEREVOL}} \tilde{W}_j \xrightarrow{\text{long-time PF}} \tilde{U}_{\text{res}}^{(j)} \\ \xrightarrow{\text{SPARSEHAM}} \widetilde{\Delta H_j} \xrightarrow{T_{m,1}} H_{j+1}. \end{array} \quad (36)$$

Here **INTEGEREVOL** learns the correction generator needed for the long-time emulation primitive developed in Sec. IV. The resulting long-time product formula implements the residual dynamics, which is then converted into coefficient estimates by the pure-state tomography reduction described above. The final truncation step restores the promise that the estimate remains m -sparse and norm bounded.

We now explain the main ingredients of Algorithm 3. The algorithm maintains an estimate H_j inside the feasible set

$$S_{k,c} := \{H_0 : \text{supp}_P(H_0) \leq k, \|H_0\|_\infty \leq c\}. \quad (37)$$

To enforce this condition after each update, we use the sparse bounded truncation map

$$\mathcal{T}_{k,c}(H) := \arg \min_{H_0 \in S_{k,c}} \|H - H_0\|_{\ell_\infty}. \quad (38)$$

Starting from $H_0 = 0$, the algorithm aims to maintain the invariant

$$\|H - H_j\|_{\ell_\infty} \leq \eta_j, \quad \eta_j = 2^{-j}. \quad (39)$$

Thus, each round only needs to learn the residual $\Delta H_j = H - H_j$ well enough to halve the current error. After $J = \lceil \log_2(1/\varepsilon) \rceil$ rounds, the estimate H_j reaches the target accuracy.

The algorithm has two regimes, separated by the crossover point η_{sw} defined in Algorithm 3. We emphasize that $\eta_j \leq \eta_{\text{sw}}$ is not an initial assumption required for applying the full learning protocol. Starting from $H_0 = 0$ and $\eta_0 = 1$, each round with $\eta_j > \eta_{\text{sw}}$ uses the standard-quantum-limit subroutine stated as Algorithm 4 in Appendix F, to construct an update that reduces the maintained error bound by a constant factor. Once $\eta_j \leq \eta_{\text{sw}}$, the algorithm switches to the Heisenberg-limit subroutine based on long-time control emulation. Thus, when the target accuracy satisfies $\varepsilon < \eta_{\text{sw}}$, Algorithm 4 serves as an initial bootstrap stage for reaching the theoretically guaranteed long-time control-emulation regime; when $\varepsilon \geq \eta_{\text{sw}}$, the algorithm may terminate entirely within the standard-quantum-limit regime.

A. Small-error regime ($\eta_j \leq \eta_{\text{sw}}$)

When the current residual error is sufficiently small, we operate in a regime based on long-time control emulation. This regime is responsible for achieving the final Heisenberg-limited scaling in the target accuracy. Intuitively, once we are close enough to the true Hamiltonian, long-time evolution suffices to emulate residual dynamics and extract higher-precision information about the Hamiltonian. Each round consists of the following three steps.

- (1) Quantum control emulation (Algorithm 1). The goal of this step is to coherently simulate the residual evolution $e^{-i\Delta H_j t_j}$ without access to short-time queries. To this end, we first learn an effective Hamiltonian W_j that captures the discrepancy between the true dynamics and the current hypothesis. We then construct an approximate unitary

$\tilde{U}_{\text{res}}^{(j)}$ via a product formula using only long-time evolutions:

$$\tilde{U}_{\text{res}}^{(j)} = (e^{-iH(T+t_j/N_j)} e^{i\tilde{W}_j} e^{iH_j(T+t_j/N_j)})^{N_j}. \quad (40)$$

Here $\tilde{W}_j$ approximates W_j , which is defined implicitly through the exact identity

$$d(e^{-iW_j}, e^{iH_j T} e^{-iHT}) = 0. \quad (41)$$

This construction can be interpreted as a stroboscopic simulation that cancels the known component H_j while isolating and amplifying the residual ΔH_j . Using Algorithm 1 with $c = \mathcal{O}(1/\sqrt{m})$, we obtain

$$\|W_j - \tilde{W}_j\|_F = \mathcal{O}(\eta_j/\sqrt{m}).$$

This error accumulates linearly in the number of repetitions N_j , yielding

$$\|\tilde{U}_{\text{res}}^{(j)} - e^{-i\Delta H_j t_j}\|_F = \mathcal{O}\left(\frac{N_j \eta_j}{\sqrt{m}}\right) = \mathcal{O}\left(\frac{1}{m}\right). \quad (42)$$

Thus, by choosing N_j appropriately, we ensure inverse-polynomial accuracy in the implemented residual evolution.

- (2) Coefficient and structure learning (Algorithm 2). Having obtained an accurate implementation of the residual evolution, we convert this dynamical information into a state encoding the coefficients of ΔH_j . Specifically, we prepare

$$|\psi_j\rangle = V(e^{-i\Delta H_j t_j} \otimes I)|\Phi\rangle \simeq |0\rangle - it_j \sum_{x \neq 0} \alpha_x^{(j)} |x\rangle, \quad (43)$$

where $t_j = \mathcal{O}(1/(m\eta_j))$. This approximation follows from the first-order expansion of the evolution, valid since t_j is sufficiently small relative to $\|\Delta H_j\|_{\ell_\infty}$. Applying sparse tomography, we

estimate these coefficients. Writing $|\psi_j\rangle = \sum_x \beta_x |x\rangle$, the guarantee

$$|\beta_x + it_j \alpha_x^{(j)}| = \mathcal{O}(1/m)$$

implies that each $\alpha_x^{(j)}$ can be estimated within error $\mathcal{O}(1/(mt_j)) = \eta_j/4$. This yields an estimate $\widetilde{\Delta H_j}$ satisfying

$$\|H - (H_j + \widetilde{\Delta H_j})\|_{\ell_\infty} \leq \eta_j/4.$$

Importantly, this step recovers both the support and the coefficients of the residual Hamiltonian.

- (3) Sparse bounded truncation (Lemma 9). We update the hypothesis via

$$H_{j+1} = \mathcal{T}_{m,1}(H_j + \widetilde{\Delta H_j}).$$

The truncation operator $\mathcal{T}_{m,1}$ enforces that the estimate remains m -sparse and has norm at most 1, thereby preventing error accumulation and inefficiency across iterations. By stability of sparse bounded truncation, we obtain

$$\|H - H_{j+1}\|_{\ell_\infty} \leq \eta_j/2.$$

Thus, each refinement round reduces the error by a constant factor while preserving the structural constraints.

The total evolution time in the small-error regime is $t_{\text{tot,small}}^{(j)} = \tilde{\mathcal{O}}(4^m T^3/\eta_j)$ for logarithmically sparse Hamiltonians and $\tilde{\mathcal{O}}(m^{K+2}T/\eta_j)$ for polynomially sparse Hamiltonians. In the former case, T can be chosen as an arbitrarily large constant. In the latter case, we require $T = \Theta(m^{-1/K})$ for some constant $K \in \mathbb{N}$ to ensure that the simulation remains well-conditioned. See Appendix E for detailed derivations.

B. Large error regime ($\eta_j \geq \eta_{\text{sw}}$)

The long-time control-emulation step above is used only once the residual error is sufficiently small. For larger η_j ,

ALGORITHM 2. SPARSEHAM($\mathcal{U}_A, m, \varepsilon$): Sparse Hamiltonian learning in ℓ_∞ norm via pure-state tomography.

Input: Query access to $\mathcal{U}_A(\tau) = e^{-iA\tau}$ for a traceless Hamiltonian $A = \sum_x \alpha_x P_x$ satisfying $\|A\|_{\ell_\infty} \leq \varepsilon$ and $m = \text{supp}_P(A)$.

Output: An estimate $\tilde{A}$ such that $\|A - \tilde{A}\|_{\ell_\infty} \leq \varepsilon/8$.

1 $t \leftarrow 1/(32m\varepsilon)$

2 Prepare the maximally entangled state $|\Phi\rangle$

3 Let V be a unitary such that $V(P_x \otimes I)|\Phi\rangle = |x\rangle$ for all x

4 Prepare $|\psi_t\rangle \leftarrow V(\mathcal{U}_A(t) \otimes I)|\Phi\rangle$

5 Apply sparse pure-state tomography (Lemma 4) to learn $|\psi_t\rangle$ with accuracy $mt^2\varepsilon^2$ using $\mathcal{O}(m^3)$ copies, and obtain coefficients $\{\tilde{\alpha}_x\}_x$ satisfying $|\alpha_x - \tilde{\alpha}_x| \leq \varepsilon/8$ for all x

6 **return** $\tilde{A} \leftarrow \sum_x \tilde{\alpha}_x P_x$

 ALGORITHM 3. MAIN: Hamiltonian learning in Heisenberg limit without short time control.

Input: Oracle access to $e^{-iH\tau}$ for all $\tau \geq T$, where H is traceless, m -sparse, and $\|H\|_\infty \leq 1$; target accuracy $\varepsilon \in (0, 1)$.
Output: An m -sparse estimate $\tilde{H}$ satisfying $\|H - \tilde{H}\|_{\ell_\infty} \leq \varepsilon$.

```

1 Choose  $J = \lceil \log_2(1/\varepsilon) \rceil$ 
2 if  $m = \mathcal{O}(\log n)$  then
3    $s \leftarrow 4^m$ ,  $c_F \leftarrow 2\pi\sqrt{mT}$ ,  $c_\infty \leftarrow 2\pi mT$ ,  $c \leftarrow 1/(256\sqrt{m})$ 
4 else if  $m = \text{poly}(n)$  and  $T = \Theta(m^{-1/K})$  for some  $K \in \mathbb{N}$  then
5    $s \leftarrow \mathcal{O}(m^{K-1})$ ,  $c_F \leftarrow 2\sqrt{m}$ ,  $c_\infty \leftarrow 2m$ ,  $c \leftarrow 1/(256\sqrt{m})$ 
6  $\eta_{\text{sw}} \leftarrow c/(10c_F c_\infty)$ 
7 Initialize  $H_0 \leftarrow 0$ 
8 for  $j = 0$  to  $J - 1$  do
9    $\eta_j \leftarrow 2^{-j}$ ,  $t_j \leftarrow 1/(32m\eta_j)$ ,  $N_j \leftarrow \lceil 1/(2\sqrt{m}\eta_j) \rceil$ 
10   $C_j \leftarrow e^{iHT} e^{-iH_j T}$ 
11  if  $\eta_j \leq \eta_{\text{sw}}$  then
12    Let  $W_j$  be Hermitian such that  $C_j = e^{iW_j}$ 
13    Implement integer-time access to  $\mathcal{U}_{W_j}(t) = e^{-iW_j t}$  via  $(C_j^\dagger)^t$ 
14     $\tilde{W}_j \leftarrow \text{INTERGEREVOL}(\mathcal{U}_{W_j}, s, c_F, c_\infty, c, \eta_j)$ 
15     $\tilde{C}_j \leftarrow e^{i\tilde{W}_j}$ 
16     $\tau_j \leftarrow T + t_j/N_j$ 
17     $\tilde{U}_{\text{res}}^{(j)} \leftarrow (e^{-iH\tau_j} \tilde{C}_j e^{iH_j \tau_j})^{N_j}$ 
18     $\Delta H_j \leftarrow \text{SPARSEHAM}(\tilde{U}_{\text{res}}^{(j)}, m, \eta_j)$ 
19  else
20     $\Delta H_j \leftarrow \text{LEARNSTANDARDQUANTUMLIMIT}(\mathcal{U}_H, H_j, m, \eta_j)$ 
21   $H_{j+1} \leftarrow T_{m,1}(H_j + \Delta H_j)$ 
22 return  $\tilde{H} \leftarrow H_J$ 
    
```

the integer-time learning parameter in Algorithm 1 would be too small to yield a useful Heisenberg-limited primitive. In these early rounds, we instead use the simpler sub-routine Algorithm 4.

In this regime, we avoid quantum control emulation and instead directly learn $\mathcal{U}_{A,A_0}(t)$ using Algorithm 4. The unitary is reconstructed via sparse pure-state tomography, from which we extract $\tilde{\Delta H}_j$ using its first-order expansion.

Compared to the small-error regime, this approach incurs a worse dependence on η_j but avoids the overhead of simulating continuous control. After truncation, we again obtain $H_{j+1} \in \mathcal{S}_{m,1}$ with $\|H - H_{j+1}\|_{\ell_\infty} \leq \eta_j/2$.

The total evolution time in the larger error regime is $t_{\text{tot},\text{large}}^{(j)} = \tilde{\mathcal{O}}(4^m T/\eta_j^2)$ for logarithmically sparse Hamiltonians and $\tilde{\mathcal{O}}(m^K T/\eta_j^2)$ for polynomially sparse Hamiltonians. In the former case, T can be chosen as an arbitrarily large constant, whereas in the latter case we require $T = \Theta(m^{-1/K})$ for some constant $K \in \mathbb{N}$. See Appendix E for detailed derivations.

C. Total complexity and regime crossover

At the crossover point $\eta_j = \eta_{\text{sw}}$ (defined in Algorithm 3), the small-error regime and large error regime complexities coincide, and each is optimal in its respective domain. Hence $t_{\text{tot}}^{(j)} = \min\{t_{\text{tot},\text{small}}^{(j)}, t_{\text{tot},\text{large}}^{(j)}\}$ and $t_{\text{tot}} = \sum_{j=1}^J t_{\text{tot}}^{(j)}$.

For logarithmically sparse Hamiltonians, we obtain

$$t_{\text{tot}} = \tilde{\mathcal{O}}\left(\min\left\{\frac{4^m T^3}{\varepsilon}, \frac{4^m T}{\varepsilon^2}\right\}\right), \quad (44)$$

with minimum evolution time $t_{\min} = T$.

For polynomially sparse Hamiltonians, we obtain

$$t_{\text{tot}} = \tilde{\mathcal{O}}\left(\min\left\{\frac{m^{K+2} T}{\varepsilon}, \frac{m^K T}{\varepsilon^2}\right\}\right), \quad (45)$$

with $t_{\min} = \Theta(m^{-1/K}) = T$ for $K \in \mathbb{N}$. This establishes Theorems 1 and 2.

The resulting bounds show that the minimum accessible evolution time is not merely a technical parameter of the algorithm, but a resource that can be traded against total evolution time. We now discuss the implications of this tradeoff and several directions opened by the long-time learning model.

D. Implementation cost of the quantum controls

Our main complexity bounds concern access to the unknown evolution and do not explicitly count the gates required to implement the remaining quantum controls. We briefly summarize these costs here. We assume that the known Hamiltonians are supplied through their explicit Pauli decompositions and suppress polylogarithmic factors in the target simulation accuracy.

First, consider simulation of the current estimate H_j . Every occurrence of an unknown evolution $e^{-iH\tau}$, both when learning W_j and when emulating the residual dynamics, is accompanied by an evolution under the known Hamiltonian H_j for a comparable duration. Hence, summing over the complete protocol, the total simulated evolution time under H_j is $\mathcal{O}(t_{\text{tot}})$. If H_j is geometrically local on a constant-dimensional lattice, this gives gate complexity $\tilde{\mathcal{O}}(nt_{\text{tot}})$ using lattice-Hamiltonian simulation techniques with ancillary qubits [34], or $\mathcal{O}((nt_{\text{tot}})^{1+\frac{1}{K}})$ using a $(2k)$ th-order product formula [35]. Since the truncation step preserves the Pauli sparsity and coefficient bounds of H_j , its simulation has essentially the same implementation cost as simulating a known Hamiltonian from the target class.

The long-time product formula additionally requires one implementation of $e^{i\tilde{W}_j}$ per product-formula segment. Each such segment also contains one query to the unknown Hamiltonian, so the total number of correction operations is at most the total query count $\mathcal{N}$. Under the standard sparse-Pauli simulation model, in which a bounded Hamiltonian containing r explicitly specified Pauli terms can be simulated for unit time using $\tilde{\mathcal{O}}(nr)$ two-qubit gates, the total gate complexity of the correction operations is bounded by $\tilde{\mathcal{O}}(n4^m\mathcal{N})$ for logarithmically sparse Hamiltonians, and by $\tilde{\mathcal{O}}(nm^{K-1}\mathcal{N})$ for polynomially sparse Hamiltonians with fixed K . These are general Pauli-simulation bounds. Unlike H_j , the learned correction $\tilde{W}_j$ need not preserve the geometric locality of the original Hamiltonian, and its simulation may therefore constitute the dominant digital-control overhead.

It remains to account for the learning and tomography circuits. Both IntegerEvol and SparseHam prepare a Choi state by generating n maximally entangled pairs, applying the relevant evolution to one half, and applying the fixed transformation V from the Bell basis to the computational basis. Preparing the Bell pairs and implementing V require $\mathcal{O}(n)$ Clifford gates and $\mathcal{O}(n)$ two-qubit gates per circuit execution. Support identification is then performed by computational-basis measurement and requires a single measurement setting. The subsequent recovery of the relative amplitudes may be implemented using pairwise interference measurements between the reference amplitude and each recovered support element. A direct implementation therefore uses $\mathcal{O}(4^m)$ or $\mathcal{O}(m^{K-1})$ distinct measurement settings when learning W_j , and $\mathcal{O}(m)$ distinct settings when performing tomography of the residual Hamiltonian. On a generic all-to-all architecture, each such two-level basis change can be synthesized using $\mathcal{O}(n^2)$ two-qubit gates, or $\mathcal{O}(n)$ gates given suitable ancillary qubits and multi-controlled gate constructions. Every circuit execution ends with $\mathcal{O}(n)$ single-qubit measurements, and no collective measurement across different experimental repetitions is required.

Thus, besides the analog evolution time t_{tot} , the principal implementation costs are the simulation of H_j for total time $\mathcal{O}(t_{\text{tot}})$, at most $\mathcal{N}$ simulations of the sparse correction Hamiltonians $\tilde{W}_j$, and $\mathcal{O}(n)$ -size state-preparation and measurement circuits for each tomography experiment. The main practical advantage of our protocol is therefore the removal of vanishingly short evolutions under the unknown Hamiltonian, at the expense of additional digital simulation of the learned correction Hamiltonians.

Overall, relaxing the minimum evolution time can increase other implementation costs. In particular, the tradeoff between t_{min} and t_{tot} may also lead to larger gate complexity through more applications and more costly simulations of the learned correction Hamiltonians. Achieving constant t_{min} while retaining near-optimal total evolution time and circuit complexity remains an open question.

VII. DISCUSSION

To the best of our knowledge, existing Heisenberg-limited protocols for learning an m -sparse Hamiltonian without continuous quantum control rely on access to very short evolution times, with the minimum evolution time scaling as $t_{\text{min}} = \Theta(\sqrt{\epsilon})$ [9] or $t_{\text{min}} = \Theta(1/m)$ [36]. In sharp contrast, our framework demonstrates that Heisenberg-limited Hamiltonian learning is possible without access to very short-time evolutions, thereby eliminating the need for ultra-short-time control in the logarithmically sparse regime and substantially relaxing it in the polynomially sparse regime. This represents a qualitative shift in the paradigm of Hamiltonian learning: rather than relying on fine-grained temporal access, one can instead leverage long-time dynamics together with suitable algorithmic reductions to achieve information-theoretically optimal performance. In particular, for logarithmically sparse Hamiltonians, our results show that Heisenberg-limited scaling can still be attained even when the minimum accessible evolution time is fixed to be a constant independent of the system size and target accuracy.

An important limitation of the present approach is that an efficient algorithm with an arbitrary fixed minimum evolution time is obtained only for logarithmically sparse Hamiltonians. This excludes extensive local models, such as nearest-neighbor Ising-type Hamiltonians, whose number of interaction terms typically grows linearly with the system size. The factor 4^m in our current bound arises from considering the full Pauli algebra generated by the supports of H and H_j , without using the geometry of the interaction terms. We highlight the fundamental open problem of characterizing the optimal tradeoff between the total evolution time t_{tot} and the minimum evolution time t_{min} . Can one design algorithms that achieve $t_{\text{tot}} = \text{poly}(m, T)/\epsilon$ while maintaining $t_{\text{min}} = T$? More broadly,

what is the optimal dependence of t_{tot} on the sparsity m and the minimum evolution time T ? Resolving these questions would provide a complete understanding of the role of temporal resources in Hamiltonian learning. One can consider a local Hamiltonian on a bounded-degree interaction graph. On this Hamiltonian, nested commutators can be nonzero only when the participating terms are connected through overlapping supports. This suggests that the correction generator may be described or approximated using connected local clusters, rather than by enumerating the full Pauli algebra. A natural open question is whether such a locality-sensitive analysis can remove the exponential dependence on m and yield a polynomial-time long-time-only learning algorithm with a constant minimum evolution time. Establishing such a result even for geometrically local Ising-type models, including non-commuting variants, would be an important first step toward extending our framework to physically relevant many-body systems.

From an experimental perspective, this temporal-resource tradeoff also exchanges one implementation challenge for another. Our protocol avoids stationary evolution intervals that become arbitrarily short with the target accuracy, but the emulation of residual dynamics requires product-formula sequences containing repeated long-time evolutions. The potential advantage is that each long-time query may contain a well-established stationary plateau under the target Hamiltonian, rather than requiring the device to reach, maintain, and leave the target operating point within a vanishingly short interval. The cost is a longer coherent sequence and greater circuit depth, which can increase the effects of decoherence and accumulated control errors. This tradeoff may therefore be favorable when finite ramping and settling times are the dominant obstacle to short-time access, while evolution under the settled Hamiltonian remains sufficiently stable and coherent throughout the longer sequence. Conversely, it may offer little practical advantage when coherence time or accumulated control errors are the dominant limitations, or when short stationary evolutions can already be implemented reliably. Thus, our results identify a hardware-dependent operating regime rather than establishing that long-time access is universally easier than short-time access.

This conceptual shift brings crucial clarity to the distinction between minimum evolution time and time resolution. While the minimum evolution time ($t_{\text{min}} := \min_i t_i$) is often referred to as the minimum time resolution in the literature [5,9,36], from a hardware perspective, it is more natural to formalize time resolution as a fundamental discretization step δ (e.g., the clock speed of a control generator). Under this definition, accessible evolutions are restricted to $U(t_j) = e^{-iHt_j}$ with $t_j \in \delta\mathbb{N}$, even if all such times are required to satisfy $t_j \geq t_{\text{min}}$. Thus, t_{min} and δ capture distinct experimental limitations: the former specifies the shortest duration for which the system must be

coherently evolved, whereas the latter specifies how finely the evolution time must be programmable. Our results address the former by showing that Heisenberg-limited learning can be achieved without accessing ultra-short evolutions. Determining the minimum time resolution δ required for efficient Hamiltonian learning remains an important open problem.

Beyond temporal constraints, a crucial practical direction is to study Hamiltonian learning under the limited availability of ancillary systems. In many physical platforms, introducing and maintaining clean ancillas is highly restricted or costly. Inspired by recent progress in ancilla-free Hamiltonian learning [8], it would be highly desirable to develop ancilla-free algorithms that operate solely using long-time evolutions. Furthermore, understanding the interplay between long-time dynamics and additional control resources, such as static single-qubit fields [37], may provide new avenues for improving practical performance while maintaining theoretical guarantees.

From an experimental perspective, several additional directions merit investigation. One natural question concerns even more restrictive access models, where only a limited subset of the system can be probed, without requiring global control over the dynamics. In recent work on quantum probe tomography [38], it has been shown that, for certain structured Hamiltonians, Heisenberg-limited scaling can be achieved even under such locality constraints, using only local probe operations. It is therefore natural to ask whether similar guarantees can be established in the long-time evolution regime considered in this work. In particular, can one achieve Heisenberg-limited Hamiltonian learning using only local probes combined with long-time dynamics, thereby completely eliminating the need for global control?

Another important experimental consideration concerns the precision with which long-time evolution can be implemented. While our framework removes the need for ultra-short-time control, it still implicitly assumes the ability to realize time evolution for a prescribed duration T with high accuracy. In realistic experimental settings, however, such timing may only be achievable up to multiplicative uncertainties, e.g., implementing an evolution time in the range $[(1 - \delta)T, (1 + \delta)T]$ for some constant $\delta > 0$. It is therefore an important open question whether Heisenberg-limited scaling can still be achieved under such timing imprecision. Addressing this question would significantly enhance the robustness and practical relevance of long-time Hamiltonian learning protocols.

ACKNOWLEDGMENTS

M. S. thanks Yu Tong for insightful discussions. J. L. thanks Anurag Anshu, Sitan Chen, Francisco Escudero Gutiérrez, Weiyuan Gong, Antonio Anna Mele, and Chirag Wadhwa for helpful feedback on an earlier draft and for insightful discussions. J. L. also thanks Constantin

Cedillo Vayson de Pradenne, Jordan Cotler, and Hsin-Yuan Huang for sharing their concurrent work. J. L. acknowledges support from the National Research Foundation of Korea (NRF) grants funded by the MSIT (Grant No. RS-2025-00515537) and the Institute for Information & Communications Technology Promotion (IITP) grant funded by the Korean government (MSIP) (Grants No. RS-2025-02304540 and No. RS-2019-II190003). C. O. was supported by the NRF Grants (No. RS-2024-00431768 and No. RS-2025-00515456) funded by the Korean government (MSIT) and IITP grants funded by the Korea government (MSIT) (No. RS-2024-00437284, No. IITP-2025-RS-2025-02283189 and No. IITP-2025-RS-2025-02263264) and by Global Partnership Program of Leading Universities in Quantum Science and Technology (RS-2025-08542968) through the National Research Foundation of Korea (NRF) funded by the Korean government (Ministry of Science and ICT(MSIT)).

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

APPENDIX A: RELATED WORK

The interactions of a quantum system are encoded in its Hamiltonian, and characterizing this Hamiltonian from real-time dynamics is a fundamental problem in quantum information and many-body physics. In particular, learning an unknown Hamiltonian from access to its time-evolution operator has been extensively studied; see, e.g., [1–15]. Beyond full reconstruction, more recent works have also considered weaker primitives such as Hamiltonian testing [39–45] and certification [45–48], which aim to verify structural or spectral properties of a Hamiltonian without fully learning it.

In this work, we focus on learning the full description of an unknown Hamiltonian using only real-time dynamics. It is by now well understood that Hamiltonian learning can be performed quite generally using derivative-based estimation techniques, which achieve total evolution time scaling $t_{\text{tot}} = \mathcal{O}(\log(n)/\epsilon^3)$ [15].

More efficient guarantees are known under additional structural assumptions. In particular, when the interaction structure is known and the Hamiltonian is low-interaction (e.g., bounded-degree local), one may assume that the Pauli operators P_x are fixed and only the coefficients α_x are unknown. In this setting, Huang *et al.* [5] show that the Hamiltonian can be learned to accuracy ϵ using total evolution time $t_{\text{tot}} = \Theta(\log(n)/\epsilon)$ with minimum evolution time $t_{\text{min}} = \Theta(\sqrt{\epsilon})$, achieving Heisenberg-limited scaling in the error parameter and improving upon the standard $1/\epsilon^2$ dependence.

However, the assumption that the interaction structure is known is often unrealistic in practice. Existing approaches relying on this assumption typically require explicit knowledge of the interaction terms, either to construct commutator expansions [16] or to perform Hamiltonian reshaping procedures that decouple the dynamics [5].

Learning the interaction structure itself has therefore attracted significant attention [8,9,17]. In particular, for $\mathcal{O}(1)$ -local Hamiltonians, Bakshi *et al.* [17] construct Goldreich-Levin-type queries using classical shadows, enabling efficient identification of the support. Their approach achieves Heisenberg-limited total evolution time $t_{\text{tot}} = \mathcal{O}(m/\epsilon)$ via a hierarchical learning scheme, while maintaining a ϵ -independent minimum evolution time $t_{\text{min}} = \Theta(1/m)$ for general local Hamiltonians and $t_{\text{min}} = \Theta(1)$ for local Hamiltonians when every site is acted on by $\mathcal{O}(1)$ terms, through approximate quantum control implemented using constant-time Trotterization. Since the fundamental lower bound is $t_{\text{tot}} = \Omega(1/\epsilon)$ [14], this result is optimal in its dependence on ϵ for local Hamiltonians.

Hamiltonian learning without locality assumptions has also been explored [8,9]. The ansatz-free approach of [8] achieves total evolution time $t_{\text{tot}} = \mathcal{O}(m^2/\epsilon)$ via a hierarchical learning scheme combined with Hamiltonian reshaping. This method isolates individual Pauli components using Pauli twirling, allowing efficient estimation of coefficients. However, implementing the reshaped dynamics requires repeated applications of ultra-short time evolutions, resulting in a minimum evolution time of $t_{\text{min}} = \Theta(\sqrt{\epsilon})$. The work of Abbas *et al.* [9] improves the total evolution time to $\mathcal{O}(m/\epsilon)$, while retaining the same requirement on t_{min} .

Lower bounds for Hamiltonian learning have also been extensively studied [7,9,37,49]. It is known that any algorithm that learns an m -sparse Hamiltonian to accuracy ϵ in the Pauli-coefficient ℓ_1 norm requires a total evolution time of at least $\Omega(m)$ [7]. This was later strengthened to a lower bound of $\Omega(m/\epsilon)$ on the total evolution time (Theorem 4.5 of [9]). Furthermore, learning a m -sparse Hamiltonian H in the ℓ_∞ norm requires at least $\Omega(m/n)$ experiments with H (Theorem 4.1 of [9]). Naturally, it follows that for $t_{\text{min}} = T$, the total evolution t_{tot} is bounded by $\Omega(mT/n)$. It has also been shown that learning a Hamiltonian using at most $\mathcal{L}$ discrete control operations to accuracy ϵ in the ℓ_∞ norm requires total evolution time $\Omega(1/(\mathcal{L}\epsilon^2))$ [8]. Consequently, achieving Heisenberg-limited scaling necessitates at least $\Omega(1/\epsilon)$ discrete control operations.

Practical aspects of Hamiltonian learning in the presence of static single-qubit fields have also been investigated [37]. In this setting, achieving Heisenberg-limited scaling in the total evolution time requires either $\mathcal{O}(1/\epsilon)$ discrete control operations or a static field strength $v = \Omega(1)$. Consequently, in the absence of static single-qubit control, one must interact with the Hamiltonian at

least $\mathcal{O}(1/\varepsilon)$ times. This implies that only with discrete control, for $t_{\min} = T$, the total evolution t_{tot} is bounded by $\Omega(T/\varepsilon)$.

Finally, recent work on quantum probe tomography [38] provides an alternative framework for Hamiltonian learning under restricted access models. In particular, it has been shown that certain structured Hamiltonians (e.g., translation- and rotation-invariant nearest-neighbor systems) can be reconstructed from local probes of Gibbs states with query complexity polynomial in $1/\varepsilon$ and classical post-processing time polylogarithmic in $1/\varepsilon$. These results demonstrate that Hamiltonian learning remains feasible even under severely constrained measurement access.

Concurrent work. Concurrently with and independently of this work, Cedillo Vayson de Pradenne, Cotler, and Huang study Hamiltonian learning from repeated access to a single fixed unitary $U = e^{-iHt}$, where t is unknown and may be arbitrarily large [50]. Their protocol uses random product-state inputs and randomized Pauli measurements and, under structural and genericity assumptions on local Hamiltonian families, recovers the normalized Hamiltonian direction up to sign with inverse-polynomial accuracy using polynomial resources. In contrast, our protocol allows multiple chosen evolution times $t_i \geq T$ and uses coherent access, including maximally entangled probes, to recover the full coefficients of any norm-bounded Pauli-sparse Hamiltonian to an arbitrary ℓ_∞ error ϵ . Moreover, our guarantees explicitly quantify the tradeoff between the minimum and total evolution times while retaining Heisenberg-limited scaling in ϵ .

APPENDIX B: PRELIMINARIES

We represent any traceless n -qubit Hamiltonian H using its Pauli decomposition

$$H = \sum_x \alpha_x P_x, \quad (\text{B1})$$

where $\alpha_x \in \mathbb{R}$ and P_x ranges over the non-identity Pauli operators.

We use several norms to quantify the size of a Hamiltonian. The (unnormalized) Frobenius norm is defined as

$$\|H\|_{\text{Frob}} = \sqrt{\text{Tr}(H^\dagger H)}. \quad (\text{B2})$$

The normalized Frobenius norm is

$$\begin{aligned} \|H\|_F &= \sqrt{\frac{\text{Tr}(H^\dagger H)}{2^n}} \\ &= 2^{-n/2} \|H\|_{\text{Frob}} = \left(\sum_x |\alpha_x|^2 \right)^{1/2}, \end{aligned} \quad (\text{B3})$$

which coincides with the ℓ_2 norm of the Pauli coefficient vector.

The operator norm is defined as

$$\|H\|_\infty = \|H\|_{\text{op}} = \sup_{\|v\|=1} \|Hv\|. \quad (\text{B4})$$

For Hermitian H , this equals the spectral norm $\max_i |\lambda_i|$. Since all Hamiltonians considered in this work are Hermitian, we use the notation $\|\cdot\|_\infty$ throughout.

We will also use the unitary distance

$$d(U, V) = \min_\phi \|U - e^{i\phi} V\|_F. \quad (\text{B5})$$

Since d is a metric, it satisfies the triangle inequality.

We also define the Pauli-coefficient ℓ_p norm of H as

$$\|H\|_{\ell_p} = |\mu|_p = \left(\sum_x |\alpha_x|^p \right)^{1/p}. \quad (\text{B6})$$

In particular, $\|H\|_{\ell_\infty} = \max_x |\alpha_x|$.

We will frequently use the following lemma relating the Pauli-coefficient norms.

Lemma 3. (Pauli-coefficient growth under powers). Let $A = \sum_x \alpha_x P_x$ be an m -sparse Hamiltonian in the Pauli basis. Then for every integer $k \geq 1$,

$$\|A^k\|_{\ell_\infty} \leq m^{k-1} \|A\|_{\ell_\infty}^k. \quad (\text{B7})$$

Proof. Let $S := \text{supp}(\alpha) = \{x: \alpha_x \neq 0\}$, $|S| \leq m$. Then

$$A^k = \sum_{x_1, \dots, x_k \in S} \alpha_{x_1} \cdots \alpha_{x_k} P_{x_1} \cdots P_{x_k}. \quad (\text{B8})$$

Fix a Pauli operator P_y . The coefficient of P_y in A^k is obtained by summing over all tuples $(x_1, \dots, x_k) \in S^k$ satisfying $P_{x_1} \cdots P_{x_k} = \omega P_y$ for some phase $\omega \in \{\pm 1, \pm i\}$.

Once $x_1, \dots, x_{k-1}$ are fixed, there is at most one choice of x_k such that the above equality holds. Hence the number of contributing tuples is at most m^{k-1} . Each term in the sum has magnitude at most $\|A\|_{\ell_\infty}^k$, which implies

$$|\text{coeff}_{P_y}(A^k)| \leq m^{k-1} \|A\|_{\ell_\infty}^k. \quad (\text{B9})$$

Taking the maximum over y yields the desired bound. ■

1. Bell sampling

We recall the Bell basis and the Bell-sampling primitive used throughout this work. For two qubits, the Bell states are

$$|\sigma_{00}\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle), \quad (\text{B10})$$

$$|\sigma_{01}\rangle = \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle), \quad (\text{B11})$$

$$|\sigma_{10}\rangle = \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle), \quad (\text{B12})$$

$$|\sigma_{11}\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle). \quad (\text{B13})$$

The 2^n -qubit Bell basis is given by all tensor products $|\sigma_s\rangle = |\sigma_{s_1}\rangle \otimes \cdots \otimes |\sigma_{s_n}\rangle$, where $s_i \in \{00, 01, 10, 11\}$. Equivalently, we may index outcomes by $\mathbf{s} \in \{0, 1\}^{2n}$ via the standard identification $\{00, 01, 10, 11\} \cong \{0, 1\}^2$ per qubit-pair.

Given an n -qubit state ρ , *Bell sampling* [51,52] means measuring $\rho^{\otimes 2}$ in the n -fold Bell basis $\{|\sigma_s\rangle\}_{\mathbf{s} \in \{0,1\}^{2n}}$. The resulting outcome $\mathbf{s}$ has distribution $q_\rho(\mathbf{s})$. For a pure state $\rho = |\psi\rangle\langle\psi|$, this distribution is given by

$$q_\rho(\mathbf{s}) = 2^{-n} |\langle\psi^*|P_s|\psi\rangle|^2. \quad (\text{B14})$$

Here $|\psi^*\rangle$ is the complex conjugate of $|\psi\rangle$ in the computational basis, and P_s is the n -qubit Pauli operator associated with $\mathbf{s}$.

For a unitary with Pauli expansion $U = \sum_x U_x \sigma_x$, measuring the Choi state $(U \otimes I)|\Phi\rangle$ in the Bell basis returns the label x with probability $|U_x|^2$. Equivalently, there exists a fixed unitary V such that

$$V(U \otimes I)|\Phi\rangle = \sum_x U_x |x\rangle, \quad (\text{B15})$$

where V transforms $(\sigma_x \otimes I)|\Phi\rangle$ into $|x\rangle$.

2. Pure-state tomography

Our Hamiltonian learning algorithm reduces the problem to sparse pure-state tomography. We therefore analyze the sample complexity of reconstructing a sparse pure quantum state under two different norms. In particular, we provide guarantees for recovery in the ℓ_∞ and ℓ_2 norms of the coefficient vector.

Lemma 4. (Sparse pure-state tomography in ℓ_∞ norm). Let $|\psi\rangle = \sum_x \alpha_x |x\rangle$ be a pure quantum state, and let $\varepsilon, \delta \in (0, 1)$. Assume there exists a set $S_\star$ with $|S_\star| \leq s$ such that $|\alpha_0 - 1| \leq \varepsilon/3$ and $|\alpha_x| \leq \varepsilon/3$ for all $x \notin S_\star$. Then there exists a randomized procedure that, using

$$\mathcal{O}\left(\frac{s}{\varepsilon^2} \log \frac{s}{\delta}\right) \quad (\text{B16})$$

copies of $|\psi\rangle$, outputs a coefficient vector $\tilde{\alpha}$ satisfying $|\tilde{\alpha} - \alpha|_\infty \leq \varepsilon$ with probability at least $1 - \delta$.

Proof. First measure $\mathcal{O}(\varepsilon^{-2} \log(s/\delta))$ copies of $|\psi\rangle$ in the computational basis and keep all indices whose empirical probability exceeds a suitable constant multiple of ε^2 .

By standard heavy-hitters guarantees, with probability at least $1 - \delta/3$, the procedure returns a set T such that $|\alpha_x| > 3\varepsilon/4$ implies $x \in T$ and $|\alpha_x| < 3\varepsilon/8$ implies $x \notin T$. Since $|\alpha_x| \leq \varepsilon/3 < 3\varepsilon/8$ for $x \notin S_\star$, we have $T \subseteq S_\star$ and $|T| \leq s$.

Moreover, every $x \notin T$ satisfies $|\alpha_x| \leq 3\varepsilon/4$, and hence $|\alpha_{T^c}|_\infty \leq 3\varepsilon/4$. Also, $|\alpha_0| \geq 1 - \varepsilon/3 > 3/4$, so $0 \in T$.

Next perform standard pure-state tomography restricted to the support T to accuracy $\varepsilon/20$. Using

$$\mathcal{O}\left(\frac{|T|}{\varepsilon^2} \log \frac{|T|}{\delta}\right) = \mathcal{O}\left(\frac{s}{\varepsilon^2} \log \frac{s}{\delta}\right) \quad (\text{B17})$$

copies, we obtain coefficients $\hat{\alpha}$ supported on T such that, with probability at least $1 - \delta/3$,

$$\|\hat{\alpha} - e^{i\phi} \alpha_T\|_2 \leq \frac{\varepsilon}{20} \quad (\text{B18})$$

for some $\phi \in \mathbb{R}$.

Define

$$\tilde{\alpha}_x := \begin{cases} \overline{\hat{\alpha}_0} \hat{\alpha}_x, & x \in T, \\ 0, & x \notin T. \end{cases} \quad (\text{B19})$$

Since

$$|\hat{\alpha}_0 - e^{i\phi}| \leq |\hat{\alpha}_0 - e^{i\phi} \alpha_0| + |\alpha_0 - 1| \quad (\text{B20})$$

$$\leq \frac{\varepsilon}{20} + \frac{\varepsilon}{3}, \quad (\text{B21})$$

the phase correction ensures $\|\tilde{\alpha}_T - \alpha_T\|_\infty \leq \varepsilon/2$. Therefore $\|\tilde{\alpha} - \alpha\|_\infty \leq \max\{\|\tilde{\alpha}_T - \alpha_T\|_\infty, \|\alpha_{T^c}\|_\infty\} \leq \varepsilon$.

The claimed sample complexity follows by summing the two stages, and the failure probability is at most δ by a union bound. ■

Lemma 5. (Sparse pure-state tomography in ℓ_2 norm). Let $|\psi\rangle = \sum_x \alpha_x |x\rangle$ be a pure quantum state, and let $\varepsilon, \delta \in (0, 1)$. Assume there exists a set $S_\star$ with $|S_\star| \leq s$ such that $|\alpha_0 - 1| \leq \varepsilon/4$ and $(\sum_{x \notin S_\star} |\alpha_x|^2)^{1/2} \leq \varepsilon/4$. Then there exists a randomized procedure that, using

$$\mathcal{O}\left(\frac{s}{\varepsilon^2} \log \frac{s}{\delta}\right) \quad (\text{B22})$$

copies of $|\psi\rangle$, outputs a coefficient vector $\tilde{\alpha}$ satisfying $|\tilde{\alpha} - \alpha|_2 \leq \varepsilon$ with probability at least $1 - \delta$.

Proof. First measure $\mathcal{O}(s\varepsilon^{-2} \log(s/\delta))$ copies of $|\psi\rangle$ in the computational basis and keep all indices whose empirical probability exceeds a suitable constant multiple of ε^2/s .

By standard heavy-hitters guarantees, with probability at least $1 - \delta/3$, the procedure produces a set T such that $|\alpha_x| > \varepsilon/(2\sqrt{s})$ implies $x \in T$ and $|\alpha_x| < \varepsilon/(4\sqrt{s})$ implies $x \notin T$, and hence $|T| = O(s)$.

Moreover,

$$\|\alpha_{T^c}\|_2 \leq \|\alpha_{S^c}\|_2 + \|\alpha_{S^c \cap T^c}\|_2 \quad (\text{B23})$$

$$\leq \frac{\varepsilon}{4} + \sqrt{s} \cdot \frac{\varepsilon}{2\sqrt{s}} \quad (\text{B24})$$

$$= \frac{3\varepsilon}{4}. \quad (\text{B25})$$

Also, $|\alpha_0| \geq 1 - \varepsilon/4 > 3/4$, so $0 \in T$.

Next perform standard pure-state tomography restricted to the support T to accuracy $\varepsilon/20$. Since $|T| = O(s)$, this uses

$$\mathcal{O}\left(\frac{|T|}{\varepsilon^2} \log \frac{|T|}{\delta}\right) = \mathcal{O}\left(\frac{s}{\varepsilon^2} \log \frac{s}{\delta}\right) \quad (\text{B26})$$

copies and returns coefficients $\hat{\alpha}$ supported on T such that, with probability at least $1 - \delta/3$,

$$\|\hat{\alpha} - e^{i\phi} \alpha_T\|_2 \leq \frac{\varepsilon}{20} \quad (\text{B27})$$

for some $\phi \in \mathbb{R}$.

Define

$$\tilde{\alpha}_x := \begin{cases} \overline{\hat{\alpha}_0} \hat{\alpha}_x, & x \in T, \\ 0, & x \notin T. \end{cases} \quad (\text{B28})$$

Since

$$|\hat{\alpha}_0 - e^{i\phi}| \leq |\hat{\alpha}_0 - e^{i\phi} \alpha_0| + |\alpha_0 - 1| \quad (\text{B29})$$

$$\leq \frac{\varepsilon}{20} + \frac{\varepsilon}{4}, \quad (\text{B30})$$

the same phase-correction argument yields $\|\tilde{\alpha}_T - \alpha_T\|_2 \leq \varepsilon/4$. Therefore

$$\|\tilde{\alpha} - \alpha\|_2 \leq \|\tilde{\alpha}_T - \alpha_T\|_2 + \|\alpha_{T^c}\|_2 \quad (\text{B31})$$

$$\leq \frac{\varepsilon}{4} + \frac{3\varepsilon}{4} \quad (\text{B32})$$

$$= \varepsilon. \quad (\text{B33})$$

The claimed sample complexity follows by summing the two stages, and the failure probability is at most δ by a union bound. ■

3. Technical tools

In this section, we collect several useful mathematical lemmas that will be used throughout the proofs. We begin with Duhamel's principle, which provides a convenient way to express and bound $e^{-iXt} e^{iYt}$.

Lemma 6. (Duhamel's principle [53]). Let X, Y be Hermitian matrices, and let $V(t) := e^{-iXt} e^{iYt}$ and $\Delta := X - Y$. Then, for every $t \geq 0$,

$$V(t) = I - i \int_0^t e^{-iXs} \Delta e^{iYs} ds. \quad (\text{B34})$$

Consequently, for any unitarily invariant norm $\|\cdot\|$, we have $\|V(t) - I\| \leq t\|\Delta\|$.

We next show that any unitary can be written, up to a global phase, as the exponential of a traceless Hermitian matrix with a controlled norm.

Lemma 7. (Traceless Hermitian representation of a unitary). Let U be a unitary matrix. Then there exists a traceless Hermitian matrix H such that $d(U, e^{-iH}) = 0$, where $d(U, V) := \min_{\phi \in \mathbb{R}} \|U - e^{i\phi} V\|$ and $\|\cdot\|$ is any unitarily invariant norm. Moreover, $\|H\| \leq \pi \|I - U\|$.

Proof. By the spectral theorem, there exists an orthonormal basis $\{|\psi_j\rangle\}_{j=1}^d$ and phases $\theta_j \in (-\pi, \pi]$ such that

$$U = \sum_{j=1}^d e^{-i\theta_j} |\psi_j\rangle \langle \psi_j|. \quad (\text{B35})$$

Define

$$\bar{\theta} := \frac{1}{d} \sum_{j=1}^d \theta_j, \quad (\text{B36})$$

$$H := \sum_{j=1}^d (\theta_j - \bar{\theta}) |\psi_j\rangle \langle \psi_j|. \quad (\text{B37})$$

Then H is Hermitian and traceless, since

$$\text{Tr}(H) = \sum_{j=1}^d (\theta_j - \bar{\theta}) = 0. \quad (\text{B38})$$

Moreover,

$$e^{-iH} = \sum_{j=1}^d e^{-i(\theta_j - \bar{\theta})} |\psi_j\rangle \langle \psi_j| = e^{i\bar{\theta}} U, \quad (\text{B39})$$

and hence $d(U, e^{-iH}) = 0$.

We now prove the norm bound. Since $\|\cdot\|$ is unitarily invariant, it depends only on the singular values. In the eigenbasis of U , both H and $I - U$ are diagonal:

$$H = \text{diag}(\theta_1 - \bar{\theta}, \dots, \theta_d - \bar{\theta}), \quad (\text{B40})$$

$$I - U = \text{diag}(1 - e^{-i\theta_1}, \dots, 1 - e^{-i\theta_d}). \quad (\text{B41})$$

For each j , we have

$$|1 - e^{-i\theta_j}| = 2 \left| \sin \frac{\theta_j}{2} \right| \geq \frac{2|\theta_j|}{\pi}, \quad (\text{B42})$$

and hence

$$|\theta_j| \leq \frac{\pi |1 - e^{-i\theta_j}|}{2}. \quad (\text{B43})$$

Let $M := \max_j |\theta_j|$. Then $|\bar{\theta}| \leq M$, and therefore

$$|\theta_j - \bar{\theta}| \leq |\theta_j| + |\bar{\theta}| \leq 2M \leq \pi \max_k |1 - e^{-i\theta_k}|. \quad (\text{B44})$$

Since Φ is monotone with respect to coordinate-wise inequalities, it follows that $\|H\| \leq \pi \|I - U\|$. ■

We next invoke the BCH formula to express the product of exponentials as a single exponential, enabling a term-by-term analysis of the resulting generator.

Lemma 8. (Baker-Campbell-Hausdorff formula, Dynkin form [54]). Let X, Y be bounded operators such that $\|X\|_\infty + \|Y\|_\infty < \log 2$. Then the Baker-Campbell-Hausdorff series converges absolutely, and

$$\log(e^X e^Y) = \sum_{r \geq 1} \text{BCH}_r(X, Y). \quad (\text{B45})$$

Here $\text{BCH}_r(X, Y)$ is the homogeneous Lie polynomial of degree r given by

$$\text{BCH}_r(X, Y) = \sum_{n=1}^r \frac{(-1)^{n-1}}{n} \sum_{\substack{r_i, s_i \geq 0 \\ r_i + s_i > 0 \\ \sum_{i=1}^n (r_i + s_i) = r}} \frac{[X^{r_1} Y^{s_1} \dots X^{r_n} Y^{s_n}]}{r \prod_{i=1}^n (r_i! s_i!)}. \quad (\text{B46})$$

Here $[X^{r_1} Y^{s_1} \dots X^{r_n} Y^{s_n}]$ denotes the left-normed commutator associated with the word consisting of r_1 copies of X , followed by s_1 copies of Y , and so on. Equivalently, whenever $s_n \geq 1$, we have

$$[X^{r_1} Y^{s_1} \dots X^{r_n} Y^{s_n}] = (\text{ad}_X)^{r_1} (\text{ad}_Y)^{s_1} \dots (\text{ad}_X)^{r_n} (\text{ad}_Y)^{s_n-1} (Y), \quad (\text{B47})$$

where $\text{ad}_X(Z) := [X, Z]$, and ad_X^k denotes its k -fold iteration.

We next derive a bound on each BCH term in the normalized Frobenius norm, which will be used to control the contribution of higher-order terms.

Corollary 1. (Degree-wise normalized Frobenius bound for BCH coefficients). Assume the hypotheses of

Lemma 8, and let $M := \max\{\|X\|_\infty, \|Y\|_\infty\}$ and $\varepsilon := \|X + Y\|_F$. Then

$$\|\text{BCH}_r(X, Y)\|_F \leq C_r M^{r-1} \varepsilon, \quad (\text{B48})$$

where $C_r := 2^{2r-1} r^r / r! \leq (4e)^r$.

Proof. Define the left-normed commutator $L_r(Z_1, \dots, Z_r) := [Z_1, [Z_2, \dots, [Z_{r-1}, Z_r] \dots]]$. We will use the inequality $\|[U, V]\|_F \leq 2\|U\|_F \|V\|_\infty$.

By induction on r , it follows that for every $1 \leq j \leq r$,

$$\|L_r(Z_1, \dots, Z_r)\|_F \leq 2^{r-1} \|Z_j\|_F \prod_{i \neq j} \|Z_i\|_\infty. \quad (\text{B49})$$

Now fix $r \geq 2$ (the case $r = 1$ is trivial), $n \in \{1, \dots, r\}$, and a tuple $(r_1, s_1, \dots, r_n, s_n)$ appearing in Lemma 8. Let $W = [X^{r_1} Y^{s_1} \dots X^{r_n} Y^{s_n}] = L_r(Z_1, \dots, Z_r)$, where $Z_j \in \{X, Y\}$.

Define the reference tuple $(\tilde{Z}_1, \dots, \tilde{Z}_r)$ by

$$\tilde{Z}_j = \begin{cases} X, & Z_j = X, \\ -X, & Z_j = Y. \end{cases} \quad (\text{B50})$$

Since each $\tilde{Z}_j$ is a scalar multiple of X , we have $L_r(\tilde{Z}_1, \dots, \tilde{Z}_r) = 0$ for $r \geq 2$. By multilinearity and telescoping,

$$L_r(Z_1, \dots, Z_r) = \sum_{j=1}^r L_r(\tilde{Z}_1, \dots, \tilde{Z}_{j-1}, Z_j - \tilde{Z}_j, Z_{j+1}, \dots, Z_r). \quad (\text{B51})$$

If $Z_j = X$, then $Z_j - \tilde{Z}_j = 0$, while if $Z_j = Y$, then $Z_j - \tilde{Z}_j = X + Y$. Hence only positions with $Z_j = Y$ contribute. Let q be the number of occurrences of Y , so $q \leq r$. Using Eq. (B49), we obtain

$$\|W\|_F \leq q 2^{r-1} M^{r-1} \|X + Y\|_F \leq r 2^{r-1} M^{r-1} \varepsilon. \quad (\text{B52})$$

Applying Eq. (B52) to the Dynkin formula in Lemma 8, the factor r cancels the denominator r , yielding

$$\begin{aligned} \|\text{BCH}_r(X, Y)\|_F &\leq (2^{r-1} M^{r-1} \varepsilon) \sum_{n=1}^r \left(\frac{1}{n} \sum_{\substack{r_i, s_i \geq 0 \\ r_i + s_i > 0 \\ \sum_{i=1}^n (r_i + s_i) = r}} \frac{1}{\prod_{i=1}^n r_i! s_i!} \right). \end{aligned} \quad (\text{B53})$$

To bound the inner sum, set $d_i := r_i + s_i \geq 1$. Then

$$\sum_{r_i + s_i = d_i} \frac{1}{r_i! s_i!} = \frac{2^{d_i}}{d_i!}, \quad (\text{B54})$$

and hence

$$\sum_{\substack{r_i, s_i \geq 0 \\ r_i + s_i > 0 \\ \sum_{i=1}^n (r_i + s_i) = r}} \frac{1}{\prod_{i=1}^n r_i! s_i!} = 2^r \sum_{\substack{d_1, \dots, d_n \geq 1 \\ d_1 + \dots + d_n = r}} \frac{1}{d_1! \dots d_n!}. \quad (\text{B55})$$

Define

$$T_{r,n} := \sum_{\substack{d_1, \dots, d_n \geq 1 \\ d_1 + \dots + d_n = r}} \frac{1}{d_1! \dots d_n!}. \quad (\text{B56})$$

Then $(e^t - 1)^n = \sum_{r \geq n} T_{r,n} t^r$, so

$$T_{r,n} = [t^r](e^t - 1)^n = \frac{n! S(r, n)}{r!}, \quad (\text{B57})$$

where $S(r, n)$ is a Stirling number of the second kind. Since $n! S(r, n) \leq n^r$, we obtain $T_{r,n} \leq n^r / r!$.

Substituting into Eq. (B53), we obtain

$$\|\text{BCH}_r(X, Y)\|_F \leq (2^{2r-1} M^{r-1} \varepsilon) \cdot \sum_{n=1}^r \frac{1}{n} T_{r,n} \quad (\text{B58})$$

$$\leq (2^{2r-1} M^{r-1} \varepsilon) \cdot \frac{1}{r!} \sum_{n=1}^r n^{r-1}. \quad (\text{B59})$$

Using $\sum_{n=1}^r n^{r-1} \leq r^r$, we conclude

$$\|\text{BCH}_r(X, Y)\|_F \leq \frac{2^{2r-1} r^r}{r!} M^{r-1} \varepsilon, \quad (\text{B60})$$

for $r \geq 2$.

Finally, Stirling's bound $r! \geq \sqrt{2\pi r} (r/e)^r$ implies

$$C_r = \frac{2^{2r-1} r^r}{r!} \leq \frac{2^{2r-1} e^r}{\sqrt{2\pi r}} = \frac{(4e)^r}{2\sqrt{2\pi r}} \leq (4e)^r, \quad (\text{B61})$$

completing the proof. $\blacksquare$

We next introduce a projection step onto the set of sparse, norm-bounded Hamiltonians, which will be used to maintain feasibility during the iteration.

Definition 1. (Sparse bounded truncation). Let $\mathcal{S}_{k,c} := \{H_0 : \text{supp}_P(H_0) \leq k, \|H_0\|_\infty \leq c\}$ be the set of k -sparse Hamiltonians with operator norm at most c . For a Hamiltonian H , define

$$\mathcal{T}_{k,c}(H) := \arg \min_{H_0 \in \mathcal{S}_{k,c}} \|H - H_0\|_\infty. \quad (\text{B62})$$

Lemma 9. (Stability of sparse bounded truncation). Let A be an m -sparse Hamiltonian with $\|A\|_\infty \leq c$, and let B be a Hamiltonian satisfying $\|A - B\|_\infty \leq \varepsilon$. Then $\|A - \mathcal{T}_{m,c}(B)\|_\infty \leq 2\varepsilon$.

Proof. Since $A \in \mathcal{S}_{m,c}$, optimality gives $\|B - \mathcal{T}_{m,c}(B)\|_\infty \leq \|B - A\|_\infty \leq \varepsilon$. The claim follows by the triangle inequality. $\blacksquare$

APPENDIX C: REDUCTION TO PURE-STATE TOMOGRAPHY

1. Correctness of Algorithm 2

In this section, we reduce Hamiltonian learning to pure-state tomography. For convenience, define a unitary V satisfying $V(P_x \otimes I)|\Phi\rangle = |x\rangle$ for all x . Such a unitary exists since the states $\{(P_x \otimes I)|\Phi\rangle : x \in \{0, 1\}^{2n}\}$ form an orthonormal basis.

Theorem 3. (Sparse Hamiltonian learning via pure-state tomography). Let $A = \sum_{x \in \{0,1\}^{2n}} \alpha_x P_x$ be an n -qubit traceless Hermitian operator, where $\{P_x\}$ is the Pauli basis and $|\text{supp}(\alpha)| \leq m$. Assume that $\|A\|_\infty \leq \varepsilon$. Then, using queries to the real-time evolution operator e^{-iAt} with $t = \mathcal{O}(1/(m\varepsilon))$, we can learn an m -sparse Hamiltonian A_1 such that $\|A - A_1\|_\infty \leq \varepsilon/8$. The total evolution time is $\mathcal{O}(m^2/\varepsilon)$.

To prove Theorem 3, we first introduce the following lemma.

Lemma 10. (First-order reduction to pure-state tomography). Let $A = \sum_x \alpha_x P_x$ be an m -sparse traceless Hamiltonian with $\|A\|_\infty \leq \varepsilon$. Define

$$|\psi_t\rangle = V(e^{-iAt} \otimes I)|\Phi\rangle = \sum_x \beta_x |x\rangle, \quad (\text{C1})$$

where $t \leq 1/(m\varepsilon)$. Then $|1 - \beta_0| \leq mt^2 \varepsilon^2$, and $|\beta_x + it\alpha_x| \leq mt^2 \varepsilon^2$ for all $x \neq 0$. In particular, learning $|\psi_t\rangle$ allows one to recover the coefficients α_x .

Proof. Let

$$R := e^{-iAt} - (I - iAt) = \sum_{k \geq 2} \frac{(-itA)^k}{k!}. \quad (\text{C2})$$

Then

$$\|R\|_\infty \leq \sum_{k \geq 2} \frac{t^k}{k!} \|A^k\|_\infty \quad (\text{C3})$$

$$\leq \sum_{k \geq 2} \frac{t^k}{k!} m^{k-1} \varepsilon^k \quad (\text{C4})$$

$$= mt^2 \varepsilon^2 \sum_{k \geq 2} \frac{(tm\varepsilon)^{k-2}}{k!} \quad (\text{C5})$$

$$\leq mt^2 \varepsilon^2 \sum_{k \geq 2} \frac{1}{k!} \quad (\text{C6})$$

$$< mt^2 \varepsilon^2, \quad (\text{C7})$$

where we used Lemma 3.

Since $\alpha_0 = \text{Tr}(A) = 0$, we have

$$V((I - iAt) \otimes I)|\Phi\rangle = |0\rangle - it \sum_{x \neq 0} \alpha_x |x\rangle. \quad (\text{C8})$$

On the other hand,

$$V(e^{-iAt} \otimes I)|\Phi\rangle = \sum_x \beta_x |x\rangle. \quad (\text{C9})$$

Since the Pauli coefficients of R are bounded by $\|R\|_\infty$, we have $|\beta_0 - 1| \leq mt^2 \epsilon^2$ and $|\beta_x + it\alpha_x| \leq mt^2 \epsilon^2$ for all $x \neq 0$. ■

Now we are ready to prove Theorem 3.

Proof of Theorem 3. Since A is m -sparse, there exists a set $S_\star$ with $|S_\star| = m$ such that $\alpha_x = 0$ for all $x \notin S_\star$ and $0 \notin S_\star$.

By Lemma 10, we have $|\beta_0 - 1| \leq mt^2 \epsilon^2$ and $|\beta_x| \leq mt^2 \epsilon^2$ for all $x \notin S_\star$. Applying Lemma 4, we obtain an estimate $\tilde{\beta} = \{\tilde{\beta}_x\}$ satisfying $|\tilde{\beta}_x - \beta_x| \leq 3mt^2 \epsilon^2$ using $\tilde{\mathcal{O}}(1/(mt^4 \epsilon^4))$ samples of $|\psi_t\rangle$.

Since $\alpha_x \in \mathbb{R}$, for $x \neq 0$ define $\tilde{\alpha}_x := -\text{Im}(\tilde{\beta}_x)/t$. Then

$$|\tilde{\alpha}_x - \alpha_x| \leq \frac{|\tilde{\beta}_x - \beta_x| + |\beta_x + it\alpha_x|}{t} \quad (\text{C10})$$

$$\leq \frac{3mt^2 \epsilon^2 + mt^2 \epsilon^2}{t} \quad (\text{C11})$$

$$= 4mt\epsilon^2. \quad (\text{C12})$$

Define $\tilde{A} := \sum_x \tilde{\alpha}_x P_x$. Then $\tilde{A}$ is Hermitian. Setting $t = 1/(32m\epsilon)$ yields

$$\|A - \tilde{A}\|_\infty \leq 4mt\epsilon^2 \leq \frac{\epsilon}{8}. \quad (\text{C13})$$

The number of samples is $\tilde{\mathcal{O}}(1/(mt^4 \epsilon^4)) = \tilde{\mathcal{O}}(m^3)$, and hence the total evolution time is $\tilde{\mathcal{O}}(m^3 t) = \tilde{\mathcal{O}}(m^2/\epsilon)$. ■

2. Correctness of Algorithm 1

The correctness of Algorithm 1 is guaranteed provided that $\epsilon \leq c/(10c_F c_\infty)$.

Theorem 4. (Sparse Hamiltonian learning in Frobenius norm via pure-state tomography). Let $A = \sum_{x \in \{0,1\}^{2n}} \alpha_x P_x$ be an n -qubit traceless Hermitian operator, where $\{P_x\}$ is the Pauli basis and $|\text{supp}_P(A, c\epsilon/8)| \leq s$. Assume that $\|A\|_\infty \leq c_\infty \epsilon$ and $\|A\|_F \leq c_F \epsilon$. Then, using queries to the real-time evolution operator e^{-iAt} with $t = \mathcal{O}(c/(c_F c_\infty \epsilon))$, we can learn an s -sparse Hamiltonian A_1 such that $\|A - A_1\|_F \leq c\epsilon$.

The total evolution time is

$$\tilde{\mathcal{O}}\left(\frac{sc_F c_\infty}{c^3 \epsilon}\right). \quad (\text{C14})$$

Proof. Let $R := e^{-iAt} - (I - iAt)$. Then $\|R\|_F \leq t^2 \|A\|_F \|A\|_\infty \leq c_F c_\infty t^2 \epsilon^2$, provided that $\|A\|_\infty t \leq 1$, which holds for $t \leq 1/(c_\infty \epsilon)$.

Define $|\psi_t\rangle = V(e^{-iAt} \otimes I)|\Phi\rangle = \sum_x \beta_x |x\rangle$. We also have

$$V((I - iAt) \otimes I)|\Phi\rangle = |0\rangle - it \sum_{x \neq 0} \alpha_x |x\rangle. \quad (\text{C15})$$

Since the Pauli coefficients of R are bounded by $\|R\|_F$, we have $|\beta_0 - 1| \leq c_F c_\infty t^2 \epsilon^2$ and

$$\left(\sum_{x \neq 0} |\beta_x + it\alpha_x|^2\right)^{1/2} \leq c_F c_\infty t^2 \epsilon^2, \quad (\text{C16})$$

for $x \neq 0$.

Since $s = \text{supp}_P(A, c\epsilon/8)$, there exists a set $S_\star$ with $|S_\star| = s$ and $0 \notin S_\star$ such that

$$\left(\sum_{x \notin S_\star} |\alpha_x|^2\right)^{1/2} \leq \frac{1}{8} c\epsilon. \quad (\text{C17})$$

It follows that $|\beta_0 - 1| \leq c_F c_\infty t^2 \epsilon^2$ and

$$\left(\sum_{x \notin S_\star} |\beta_x|^2\right)^{1/2} \leq c_F c_\infty t^2 \epsilon^2 + \frac{1}{8} c\epsilon. \quad (\text{C18})$$

Applying Lemma 5, we obtain $\tilde{\beta}$ satisfying $\|\tilde{\beta} - \beta\|_2 \leq 4c_F c_\infty t^2 \epsilon^2 + c\epsilon/2$.

Define $\tilde{\alpha}_x := -\text{Im}(\tilde{\beta}_x)/t$. Then

$$\|\tilde{\alpha} - \alpha\|_2 \leq \frac{\|\tilde{\beta} - \beta\|_2 + \|\beta + it\alpha\|_2}{t} \quad (\text{C19})$$

$$\leq \frac{4c_F c_\infty t^2 \epsilon^2 + c\epsilon/2 + c_F c_\infty t^2 \epsilon^2}{t} \quad (\text{C20})$$

$$= 5c_F c_\infty t\epsilon^2 + c\epsilon/2. \quad (\text{C21})$$

Define $\tilde{A} := \sum_x \tilde{\alpha}_x P_x$, which is Hermitian. Setting $t = \lfloor c/(10c_F c_\infty \epsilon) \rfloor$ yields $\|A - \tilde{A}\|_F \leq c\epsilon$, provided that $c/(10c_F c_\infty \epsilon) \geq 1$.

The total evolution time is

$$\tilde{\mathcal{O}}\left(\frac{st}{(c_F c_\infty t^2 \epsilon^2)^2}\right) = \tilde{\mathcal{O}}\left(\frac{sc_F c_\infty}{c^3 \epsilon}\right). \quad (\text{C22})$$

■

APPENDIX D: APPROXIMATING CONTINUOUS QUANTUM CONTROL USING LONG-TIME EVOLUTION

In this section, we show how to approximate continuous quantum control using only long-time evolution.

Define $C_j := e^{iHT} e^{-iH_j T}$. By Lemma 7, there exists a Hermitian matrix W_j such that $d(C_j, e^{iW_j}) = 0$. We assume that H and H_j are m -sparse traceless Hamiltonians satisfying $\|H - H_j\|_{\ell_\infty} \leq \varepsilon$. Our goal is to learn a Hermitian matrix $\tilde{W}$ such that $\|W - \tilde{W}\|_F \leq c\varepsilon$, for a desired constant c , using queries to e^{-iHT} and with total evolution time achieving Heisenberg-limited scaling.

This procedure will serve as a subroutine for Hamiltonian learning.

1. Approximating continuous quantum control with arbitrary long time evolution

To efficiently learn the Hamiltonian W_j , we first bound its norm and sparsity. Using Lemma 6 and Lemma 7, we have

$$\|W_j\|_F \leq \pi \|I - C_j\|_F \leq \pi T \|H - H_j\|_F \leq 2\pi T \sqrt{m} \varepsilon, \quad (\text{D1})$$

and

$$\|W_j\|_\infty \leq \pi \|I - C_j\|_\infty \leq \pi T \|H - H_j\|_\infty \leq 2\pi T m \varepsilon, \quad (\text{D2})$$

since $\|H - H_j\|_{\ell_\infty} \leq \varepsilon$ and both H, H_j are m -sparse.

We now bound the sparsity of W_j .

Lemma 11. (Pauli sparsity of a bounded logarithm). Let $H = \sum_{a \in S_H} \alpha_a P_a$ and $H_j = \sum_{b \in S_{H_j}} \beta_b P_b$ be m -sparse Hamiltonians. Then any Hermitian W_j satisfying $d(e^{iHT} e^{-iH_j T}, e^{iW_j}) = 0$ and $\|W_j\|_\infty < \pi$ obeys $|\text{supp}_P(W_j)| \leq 4^m$.

Proof. Let $L := \text{span}_{\mathbb{F}_2}(S_H \cup S_{H_j}) \subseteq \mathbb{F}_2^n$ and $r := \dim L$. Then $r \leq 2m$ and $|L| = 2^r$.

Let $\mathcal{A}_L := \text{span}_{\mathbb{C}}\{P_v : v \in L\}$. Since Pauli multiplication is closed modulo phase, $\mathcal{A}_L$ is a finite-dimensional unital $*$ -subalgebra. As $H, H_j \in \mathcal{A}_L$, it follows from the power-series expansion that $e^{iHT}, e^{-iH_j T} \in \mathcal{A}_L$, and hence

$$C_j := e^{iHT} e^{-iH_j T} \in \mathcal{A}_L.$$

By the assumption $d(C_j, e^{iW_j}) = 0$, there exists a phase ϕ such that

$$e^{iW_j} = e^{i\phi} C_j.$$

Set $C_{j,\phi} := e^{i\phi} C_j$. Then $C_{j,\phi} \in \mathcal{A}_L$ is unitary. Write its spectral decomposition as $C_{j,\phi} = \sum_k \lambda_k \Pi_k$, where the λ_k 's

are distinct and $|\lambda_k| = 1$. Since $C_{j,\phi}$ has a finite spectrum, each spectral projector Π_k can be expressed as a polynomial in $C_{j,\phi}$, e.g., via Lagrange interpolation, and hence $\Pi_k \in \mathcal{A}_L$.

It remains to show that the given W_j lies in $\mathcal{A}_L$. Since $\|W_j\|_\infty < \pi$, the spectrum of W_j lies in $[-\pi, \pi]$, where the map $\theta \mapsto e^{i\theta}$ is injective. Therefore, for each eigenvalue λ_k of $C_{j,\phi} = e^{iW_j}$, there is a unique $\theta_k \in [-\pi, \pi]$ such that $\lambda_k = e^{i\theta_k}$. Hence W_j is scalar on each eigenspace $\text{ran}(\Pi_k)$, and we can write

$$W_j = \sum_k \theta_k \Pi_k.$$

Since each $\Pi_k \in \mathcal{A}_L$, we obtain $W_j \in \mathcal{A}_L$.

Finally, since $W_j \in \mathcal{A}_L = \text{span}_{\mathbb{C}}\{P_v : v \in L\}$, its Pauli support is contained in L . Therefore $|\text{supp}_P(W_j)| \leq |L| = 2^r \leq 4^m$. ■

We now apply Theorem 4 to learn W_j in normalized Frobenius norm. Let $s := |\text{supp}_P(W_j)| \leq 4^m$, $c_F := 2\pi T \sqrt{m}$, and $c_\infty := 2\pi T m$. Choosing $c = \mathcal{O}(1/\sqrt{m})$, we obtain a total evolution time of

$$\tilde{\mathcal{O}}\left(\frac{sc_F c_\infty}{c^3 \varepsilon}\right) = \tilde{\mathcal{O}}\left(\frac{4^m T^2}{\varepsilon}\right) \quad (\text{D3})$$

with respect to evolution under W_j .

We assume $\varepsilon \leq 1/(m^2 T^2)$ and set $t = \mathcal{O}(1/(m^2 T^2 \varepsilon))$ as an integer. This enables queries to $(C_j^\dagger)^q = e^{-iW_j q}$ for $q \in \mathbb{N}$. Since each query to $C_j^\dagger = e^{iH_j T} e^{-iHT}$ requires evolution time T under H , the total evolution time with respect to H is $\tilde{\mathcal{O}}(4^m T^3 / \varepsilon)$.

2. Approximating continuous quantum control for many-body Hamiltonians

When H is a many-body Hamiltonian, its sparsity scales as $m = \text{poly}(n)$. From the previous section, we have $|\text{supp}_P(W_j)| \leq 4^m$, which becomes exponential in n . Thus, directly learning W_j is infeasible.

To address this, we restrict the evolution time to scale with the sparsity, setting $T = \mathcal{O}(m^{-1/K})$ for some constant $K \in \mathbb{N}$. In this regime, we show that W_j admits a quasi-sparse approximation via the BCH expansion.

Lemma 12. (Quasi Pauli sparsity of the BCH generator). Let $e^{iH_j T} e^{-iHT} = e^{-iW_j}$, where H, H_j are m -sparse traceless Hermitian Hamiltonians satisfying $\|H\|_\infty, \|H_j\|_\infty \leq C$ and $\|H - H_j\|_{\ell_\infty} \leq \varepsilon$, for some constants $C \geq 1$, $\varepsilon > 0$, and $K \in \mathbb{N}$. Then for every $\ell \geq 1$, there exists $T = \Theta(m^{-1/K}/C)$ such that

$$\text{supp}_P\left(W_j, \frac{m^{-\ell/K}\varepsilon}{8}\right) = \mathcal{O}(m^{\ell+K/2-1}), \quad (\text{D4})$$

and satisfies $\|W_j\|_F \leq \sqrt{m}\varepsilon$ and $\|W_j\|_\infty \leq m\varepsilon$.

Proof. Assuming $T \leq 1/(8eC)$ satisfies the convergence condition of the Baker-Campbell-Hausdorff equation (Lemma 8), we have

$$W_j = i \sum_{r \geq 1} \text{BCH}_r(iH_j T, -iHT). \quad (\text{D5})$$

Since $\| -iHT + iH_j T \|_F = T \|H - H_j\|_F \leq T \sqrt{m}\varepsilon$ and $\max\{\| -iHT \|_\infty, \|iH_j T\|_\infty\} \leq TC$, Corollary 1 yields

$$\|\text{BCH}_r(iH_j T, -iHT)\|_F \leq (4TeC)^r \sqrt{m}\varepsilon, \quad (\text{D6})$$

$$\|\text{BCH}_r(iH_j T, -iHT)\|_\infty \leq (4TeC)^r m\varepsilon. \quad (\text{D7})$$

Moreover, $\text{supp}_P(\text{BCH}_r(iH_j T, -iHT)) \leq (2m)^r$.

Define $W_j^{(k)} := i \sum_{r=1}^k \text{BCH}_r(iH_j T, -iHT)$. Then

$$\|W_j - W_j^{(k)}\|_F \leq \sum_{r \geq k+1} a^r \sqrt{m}\varepsilon = \frac{a^{k+1}}{1-a} \sqrt{m}\varepsilon. \quad (\text{D8})$$

Furthermore,

$$\text{supp}_P(W_j^{(k)}) \leq \sum_{r=1}^k (2m)^r \leq k(2m)^k, \quad (\text{D9})$$

and hence $\text{supp}_P(W_j, \frac{a^{k+1}}{1-a} \sqrt{m}\varepsilon) \leq k(2m)^k$.

For $T = m^{-1/K}/(16eC)$, we have $a = m^{-1/K}/4 \leq 1/4$. When $k = K/2 + \ell - 1$ is an integer with $k \geq 1$,

$$\frac{a^{k+1}}{1-a} \sqrt{m}\varepsilon \leq \frac{4}{3} 4^{-(k+1)} m^{-\ell/K} \leq \frac{1}{12} m^{-\ell/K} < \frac{1}{8} m^{-\ell/K}.$$

$$\text{supp}_P\left(W_j, \frac{m^{-\ell/K}\varepsilon}{8}\right) = \mathcal{O}(m^{\ell+K/2-1}). \quad (\text{D10})$$

Thus the 1/8 target in Eq. (D10) is retained for these integer truncation orders. This calculation does not specify a rounding rule for noninteger k .

For this choice of T , we also have

$$\|W_j\|_F \leq \sum_{r \geq 1} (4TeC)^r \sqrt{m}\varepsilon \quad (\text{D11})$$

$$\leq \sum_{r \geq 1} (1/2)^r \sqrt{m}\varepsilon \quad (\text{D12})$$

$$= \sqrt{m}\varepsilon, \quad (\text{D13})$$

and similarly $\|W_j\|_\infty \leq m\varepsilon$. ■

We now aim to learn W_j within error $c\varepsilon$ in normalized Frobenius norm. Suppose $c = \mathcal{O}(1/\sqrt{m})$. Setting $\ell = K/2$, we obtain $s = |\text{supp}_P(W_j, c\varepsilon/8)| = \mathcal{O}(m^{K-1})$, and parameters $c_F = \sqrt{m}$ and $c_\infty = m$.

Applying Theorem 4, we obtain a total evolution time of

$$\tilde{\mathcal{O}}\left(\frac{sc_\infty c_F T}{c^3 \varepsilon}\right) = \tilde{\mathcal{O}}\left(\frac{m^{K+2} T}{\varepsilon}\right). \quad (\text{D14})$$

Notably, Theorem 4 remains applicable under this quasi-sparsity condition.

APPENDIX E: MAIN ALGORITHM

In this section, we elaborate on the details of the main algorithm. We first define the global parameters s, c_F, c_∞ , where s denotes the (quasi-)sparsity of W_j , and c_F, c_∞ are constants such that $\|W_j\|_F \leq c_F \eta_j$ and $\|W_j\|_\infty \leq c_\infty \eta_j$. The desired error constant c will be specified later. The values of these parameters are summarized in Algorithm 3.

The algorithm follows a hierarchical refinement scheme that halves the error at each iteration:

$$\|H - H_j\|_{\ell_\infty} \leq \eta_j \Rightarrow \|H - H_{j+1}\|_{\ell_\infty} \leq \frac{\eta_j}{2} = \eta_{j+1}. \quad (\text{E1})$$

At iteration j , we define the local parameters

$$\eta_j = 2^{-j}, \quad t_j = \frac{1}{32m\eta_j}, \quad N_j = \frac{1}{2\sqrt{m}\eta_j}. \quad (\text{E2})$$

After $J = \lceil \log_2(1/\varepsilon) \rceil$ iterations, the estimate H_J satisfies $\|H - H_J\|_{\ell_\infty} \leq \varepsilon$.

We now describe the j th iteration in the regime $\eta_j \leq \eta_{\text{sw}}$.

(1) Quantum control emulation (Algorithm 1). We define the emulated control unitary

$$\tilde{U}_{\text{res}}^{(j)} := (e^{-iH(T+t_j/N_j)} e^{i\tilde{W}_j} e^{iH_j(T+t_j/N_j)})^{N_j}, \quad (\text{E3})$$

where W_j satisfies $e^{-iW_j} = e^{iH_j T} e^{-iHT}$. In particular, for every $t \in \mathbb{N}$,

$$e^{-iW_j t} = (e^{iH_j T} e^{-iHT})^t, \quad (\text{E4})$$

which can be implemented using only long-time evolutions. Applying Algorithm 1 with parameters s, c_F, c_∞ and error constant $c = 1/(256\sqrt{m})$, we obtain $\tilde{W}_j$ such that

$$\|W_j - \tilde{W}_j\|_F \leq c\eta_j = \frac{\eta_j}{256\sqrt{m}}. \quad (\text{E5})$$

Consequently,

$$\begin{aligned}
 & \| (e^{-iHt/N_j} e^{iH_j t/N_j})^{N_j} - \tilde{U}_{\text{res}}^{(j)} \|_F \\
 & \leq \| (e^{-iHt_j} e^{iW_j} e^{iH_j t_j})^{N_j} - (e^{-iHt_j} e^{i\tilde{W}_j} e^{iH_j t_j})^{N_j} \|_F \\
 & \leq N_j \| e^{iW_j} - e^{i\tilde{W}_j} \| \leq N_j \| W_j - \tilde{W}_j \| \\
 & \leq \frac{1}{512m}. \tag{E6}
 \end{aligned}$$

By the Trotter formula,

$$\begin{aligned}
 & \| (e^{-iHt_j/N_j} e^{iH_j t_j/N_j})^{N_j} - e^{-i\Delta H_j t_j} \|_F \\
 & \leq \frac{t_j^2}{N_j} \sqrt{m} \eta_j \leq \frac{1}{512m}. \tag{E7}
 \end{aligned}$$

Therefore,

$$\| \tilde{U}_{\text{res}}^{(j)} - e^{-i\Delta H_j t_j} \|_F \leq \frac{1}{256m}. \tag{E8}$$

This $\mathcal{O}(1/m)$ control error is precisely the regime required for the next step. The cost of learning $\tilde{W}_j$ is

$$\tilde{\mathcal{O}}\left(\frac{4^m T^3}{\eta_j}\right) \quad \text{or} \quad \tilde{\mathcal{O}}\left(\frac{m^{K+2} T}{\eta_j}\right), \tag{E9}$$

for $m = \mathcal{O}(\log n)$ and $m = \text{poly}(n)$, respectively.

(2) Coefficient and structure learning (Algorithm 2).

Using $\tilde{U}_{\text{res}}^{(j)}$, we approximate

$$|\psi_j\rangle = V(e^{-i\Delta H_j t_j} \otimes I)|\Phi\rangle \simeq |0\rangle - it_j \sum_{x \neq 0} \alpha_x^{(j)} |x\rangle. \tag{E10}$$

Applying Algorithm 2, we can learn ΔH_j up to error $\eta_j/8$ assuming access to $e^{-i\Delta H_j t_j}$. Since we instead use $\tilde{U}_{\text{res}}^{(j)}$, this introduces an additional error. Because

$$\| \tilde{U}_{\text{res}}^{(j)} - e^{-i\Delta H_j t_j} \|_F \leq \frac{1}{256m}, \tag{E11}$$

this contributes at most $1/(256mt_j) = \eta_j/8$ to the coefficient estimation error. Hence,

$$\| \Delta H_j - \widetilde{\Delta H}_j \|_{\ell_\infty} \leq \frac{\eta_j}{4}. \tag{E12}$$

This step uses $\mathcal{O}(m^3)$ queries to $\tilde{U}_{\text{res}}^{(j)}$. Since each query requires N_j calls to $e^{-iH(T+t_j/N_j)}$, the total

evolution time is

$$\mathcal{O}(m^3 N_j T) = \mathcal{O}\left(\frac{m^{2.5} T}{\eta_j}\right). \tag{E13}$$

(3) Sparse bounded truncation and update (Lemma 9). We restore the sparsity and norm constraints by

$$H_{j+1} := \mathcal{T}_{m,1}(H_j + \widetilde{\Delta H}_j). \tag{E14}$$

By Lemma 9, this ensures $H_{j+1} \in \mathcal{S}_{m,1}$ and

$$\begin{aligned}
 \| H - H_{j+1} \|_{\ell_\infty} & \leq 2 \| H - (H_j + \widetilde{\Delta H}_j) \|_{\ell_\infty} \\
 & \leq \frac{\eta_j}{2} = \eta_{j+1}. \tag{E15}
 \end{aligned}$$

This step is purely classical and requires no additional quantum resources.

The overall complexity is dominated by the cost of learning $\tilde{W}_j$. Therefore,

$$t_{\text{tot}} = \sum_{j=1}^J \tilde{\mathcal{O}}\left(\frac{4^m T^3}{\eta_j}\right) = \tilde{\mathcal{O}}\left(\frac{4^m T^3}{\varepsilon}\right), \tag{E16}$$

for $m = \mathcal{O}(\log n)$, and

$$t_{\text{tot}} = \sum_{j=1}^J \tilde{\mathcal{O}}\left(\frac{m^{K+2} T}{\eta_j}\right) = \tilde{\mathcal{O}}\left(\frac{m^{K+2} T}{\varepsilon}\right), \tag{E17}$$

for $m = \text{poly}(n)$ with $T = \Theta(m^{-1/K})$, where $K \in \mathbb{N}$.

APPENDIX F: HAMILTONIAN LEARNING WITH LONG TIME DYNAMICS IN THE LARGE ERROR REGIME

Algorithm 4 reconstructs the classical representation (Pauli coefficients) of the unitary $\mathcal{U}_{A,A_0}(t) = e^{-iAt} e^{iA_0 t}$. Assuming $\|A - A_0\|_{\ell_\infty} \leq \varepsilon$, for $t \leq 1/\sqrt{m}$ we have

$$\begin{aligned}
 & \| \mathcal{U}_{A,A_0}(t) - I + i(A - A_0)t \|_{\ell_\infty} \\
 & \leq \| \mathcal{U}_{A,A_0}(t) - I + i(A - A_0)t \|_F \leq \sqrt{m} \varepsilon t^2 = \delta_t. \tag{F2}
 \end{aligned}$$

Suppose we obtain $\tilde{\mathcal{U}}_{A,A_0}(t)$ satisfying $\min_\phi \| \tilde{\mathcal{U}}_{A,A_0}(t) - e^{i\phi} \mathcal{U}_{A,A_0}(t) \|_F \leq \delta_t$, and write $\mathcal{U}_{A,A_0}(t) = \sum_x \alpha_x P_x$ and $\tilde{\mathcal{U}}_{A,A_0}(t) = \sum_x \tilde{\alpha}_x P_x$. Then $|e^{i\phi} \tilde{\alpha}_x - \alpha_x| \leq \delta_t$ and $|1 - \alpha_0| \leq \delta_t$, which implies

$$|e^{i\phi} - \tilde{\alpha}_0| \leq |1 - \alpha_0| + |e^{i\phi} \alpha_0 - \tilde{\alpha}_0| \leq 2\delta_t. \tag{F3}$$

Hence $|\alpha_x - \tilde{\alpha}_0^* \tilde{\alpha}_x| \leq 3\delta_t$ for all x , and thus $\| \tilde{\alpha}_0^* \tilde{\mathcal{U}}_{A,A_0}(t) - I + i(A - A_0)t \|_{\ell_\infty} \leq 4\delta_t$.

ALGORITHM 4. LEARNSTANDARDQUANTUMLIMIT($\mathcal{U}_A, A_0, m, \varepsilon$).

Input: Query access to $\mathcal{U}_A(\tau) = e^{-iA\tau}$ for an m -sparse traceless Hermitian operator A , and a known m -sparse Hamiltonian A_0 satisfying $\|A - A_0\|_{\ell_\infty} \leq \varepsilon$ with $\text{supp}_P(A) = \text{supp}_P(A_0) = m$.
Output: An estimate $\tilde{A}$ such that $\|A - A_0 - \tilde{A}\|_{\ell_\infty} \leq \varepsilon/4$.
1 Set $t \leftarrow 1/(16\sqrt{m})$ and $\delta_t \leftarrow \sqrt{m}\varepsilon t^2$
2 Prepare the maximally entangled state $|\Phi\rangle$ and a unitary V such that $V(P_x \otimes I)|\Phi\rangle = |x\rangle$ for all x
3 Define $\mathcal{U}_{A,A_0}(T) \leftarrow \mathcal{U}_A(T)e^{iA_0T}$ and $\mathcal{U}_{A,A_0}(T+t) \leftarrow \mathcal{U}_A(T+t)e^{iA_0(T+t)}$
4 Prepare $|\psi_T\rangle \leftarrow V(\mathcal{U}_{A,A_0}(T) \otimes I)|\Phi\rangle$ and $|\psi_{T+t}\rangle \leftarrow V(\mathcal{U}_{A,A_0}(T+t) \otimes I)|\Phi\rangle$
5 **if** $m = \mathcal{O}(\log n)$ **then**
6 $s \leftarrow 4^m$
7 **else if** $m = \text{poly}(n)$ and $T = \Theta(m^{-1/K})$ for some $K \in \mathbb{N}$ **then**
8 $s \leftarrow \mathcal{O}(m^{K-1})$
9 Apply sparse pure-state tomography to learn $|\psi_T\rangle$ and $|\psi_{T+t}\rangle$ to accuracy δ_t using $\mathcal{O}(s/\delta_t^2)$ copies, obtaining coefficients $\{\tilde{\alpha}_x^{(T)}\}_x$ and $\{\tilde{\alpha}_x^{(T+t)}\}_x$ satisfying

$$\left(\min_\phi \sum_x |\alpha_x^{(T)} - e^{i\phi} \tilde{\alpha}_x^{(T)}|^2\right)^{1/2} \leq \delta_t/2, \quad \left(\min_\phi \sum_x |\alpha_x^{(T+t)} - e^{i\phi} \tilde{\alpha}_x^{(T+t)}|^2\right)^{1/2} \leq \delta_t/2. \quad (\text{F1})$$

10 Form $\tilde{\mathcal{U}}_{A,A_0}(T) \leftarrow \sum_x \tilde{\alpha}_x^{(T)} P_x$ and $\tilde{\mathcal{U}}_{A,A_0}(T+t) \leftarrow \sum_x \tilde{\alpha}_x^{(T+t)} P_x$
11 Compute $\tilde{\mathcal{U}}_{A,A_0}(t) \leftarrow e^{iA_0T} \tilde{\mathcal{U}}_{A,A_0}(T)^\dagger \tilde{\mathcal{U}}_{A,A_0}(T+t) e^{-iA_0T}$
12 Let $\tilde{\alpha}_0$ be the identity Pauli coefficient of $\tilde{\mathcal{U}}_{A,A_0}(t)$
13 **return** $\tilde{A} \leftarrow -\text{Im}(\tilde{\alpha}_0^* \tilde{\mathcal{U}}_{A,A_0}(t))/t$

Finally,

$$\|\tilde{A} - (A - A_0)\|_{\ell_\infty} \quad (\text{F4})$$

$$\leq \frac{1}{t} \|\tilde{\alpha}_0^* \tilde{\mathcal{U}}_{A,A_0}(t) - I + i(A - A_0)t\|_{\ell_\infty} \quad (\text{F5})$$

$$\leq \frac{4\delta_t}{t} \leq \frac{\varepsilon}{4}, \quad (\text{F6})$$

where $\tilde{A} = -\text{Im}(\tilde{\alpha}_0^* \tilde{\mathcal{U}}_{A,A_0}(t))/t$.

We now show how to reconstruct $\mathcal{U}_{A,A_0}(t)$ using only access to $\mathcal{U}_A(\tau) = e^{-iA\tau}$ for $\tau \geq T$. Observe that

$$\mathcal{U}_{A,A_0}(t) = e^{iA_0T} \mathcal{U}_{A,A_0}(T)^\dagger \mathcal{U}_{A,A_0}(T+t) e^{-iA_0T}. \quad (\text{F7})$$

Thus, learning the Pauli coefficients of $\mathcal{U}_{A,A_0}(T)$ and $\mathcal{U}_{A,A_0}(T+t)$ suffices.

From Appendix D, we define $s_T := \text{supp}_P(\mathcal{U}_{A,A_0}(T))$ and $s_{T+t} := \text{supp}_P(\mathcal{U}_{A,A_0}(T+t))$.

When $m = \mathcal{O}(\log n)$, we have $s_T, s_{T+t} \leq 4^m$.

When $m = \text{poly}(n)$ and $T = \Theta(m^{-1/K})$, we instead define the quasi-sparsities $s_T := \text{supp}_P(\mathcal{U}_{A,A_0}(T), \delta_t)$ and $s_{T+t} := \text{supp}_P(\mathcal{U}_{A,A_0}(T+t), \delta_t)$, for which $s_T, s_{T+t} = \mathcal{O}(m^{K-1})$.

Applying sparse pure-state tomography (Lemma 5), we obtain estimates $\tilde{\mathcal{U}}_{A,A_0}(T)$ and $\tilde{\mathcal{U}}_{A,A_0}(T+t)$ satisfying $d(\mathcal{U}_{A,A_0}(T), \tilde{\mathcal{U}}_{A,A_0}(T)) \leq \delta_t/2$ and $d(\mathcal{U}_{A,A_0}(T+t), \tilde{\mathcal{U}}_{A,A_0}(T+t)) \leq \delta_t/2$ using $\tilde{\mathcal{O}}(\max\{s_T, s_{T+t}\}/\delta_t^2)$ queries.

This yields $d(\tilde{\mathcal{U}}_{A,A_0}(t), \mathcal{U}_{A,A_0}(t)) \leq \delta_t$, where $\tilde{\mathcal{U}}_{A,A_0}(t) = e^{iA_0T} \tilde{\mathcal{U}}_{A,A_0}(T)^\dagger \tilde{\mathcal{U}}_{A,A_0}(T+t) e^{-iA_0T}$.

Each query to $\mathcal{U}_{A,A_0}(T)$ or $\mathcal{U}_{A,A_0}(T+t)$ requires time T . Hence, the total evolution time is $\tilde{\mathcal{O}}(4^m T/\varepsilon^2)$ or $\tilde{\mathcal{O}}(m^K T/\varepsilon^2)$, while the classical post-processing time is $\text{poly}(s)/\varepsilon^2$, where $s := \max\{s_T, s_{T+t}\}$.

-
- [1] J. Wang, S. Paesani, R. Santagati, S. Knauer, A. A. Gentile, N. Wiebe, F. Petruccione, J. L. O'Brien, J. G. Rarity, and A. Laing, *Experimental quantum Hamiltonian learning*, *Nat. Phys.* **13**, 551 (2017).
 - [2] S. Pang and T. A. Brun, *Quantum metrology for a general Hamiltonian parameter*, *Phys. Rev. A* **90**, 022117 (2014).
 - [3] R. Shaffer, E. Megidish, J. Broz, W.-T. Chen, and H. Häffner, *Practical verification protocols for analog quantum simulators*, *npj Quantum Inf.* **7**, 46 (2021).
 - [4] A. Gu, L. Cincio, and P. J. Coles, *Practical Hamiltonian learning with unitary dynamics and Gibbs states*, *Nat. Commun.* **15**, 312 (2024).
 - [5] H.-Y. Huang, Y. Tong, D. Fang, and Y. Su, *Learning many-body Hamiltonians with Heisenberg-limited scaling*, *Phys. Rev. Lett.* **130**, 200403 (2023).
 - [6] W. Yu, J. Sun, Z. Han, and X. Yuan, *Robust and efficient Hamiltonian learning*, *Quantum* **7**, 1045 (2023).
 - [7] M. Ma, S. T. Flammia, J. Preskill, and Y. Tong, *Learning k -body Hamiltonians via compressed sensing*, *IEEE Trans. Info. Theor.* **72**, 7553 (2026). [10.1109/TIT.2026.3720996](https://arxiv.org/abs/10.1109/TIT.2026.3720996)
 - [8] H.-Y. Hu, M. Ma, W. Gong, Q. Ye, Y. Tong, S. T. Flammia, and S. F. Yelin, *Hamiltonian learning with Heisenberg-limited scaling*, *PRX Quantum* **6**, 040315 (2025).

- [9] A. Abbas, N. Cerrato, F.E. Gutiérrez, D. Grinko, F.A. Mele, and P. Sinha, *Nearly optimal algorithms to learn sparse quantum Hamiltonians in physically motivated distances*, (2026), [arXiv:2509.09813v2](#).
- [10] H. Li, Y. Tong, T. Gefen, H. Ni, and L. Ying, *Heisenberg-limited Hamiltonian learning for interacting bosons*, *npj Quantum Inf.* **10**, 83 (2024).
- [11] A. Zhao, *Learning the structure of any Hamiltonian from minimal assumptions*, in *Proceedings of the 57th Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2025), pp. 1201–1211, [10.1145/3717823.3718115](#).
- [12] M. P. da Silva, O. Landon-Cardinal, and D. Poulin, *Practical characterization of quantum devices without tomography*, *Phys. Rev. Lett.* **107**, 210404 (2011).
- [13] N. Wiebe, C. Granade, C. Ferrie, and D. G. Cory, *Hamiltonian learning and certification using quantum resources*, *Phys. Rev. Lett.* **112**, 190501 (2014).
- [14] A. Dutkiewicz, T. E. O’Brien, and T. Schuster, *The advantage of quantum control in many-body Hamiltonian learning*, *Quantum* **8**, 1537 (2024).
- [15] A. Zubida, E. Yitzhaki, N.H. Lindner, and E. Bairey, *Optimal short-time measurements for Hamiltonian learning* (2021), [arXiv:2108.08824](#).
- [16] J. Haah, R. Kothari, and E. Tang, *Optimal learning of quantum Hamiltonians from high-temperature Gibbs states*, in *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE Computer Society, Los Alamitos, CA, USA, 2022), pp. 135–146, [10.1109/FOCS54457.2022.00020](#).
- [17] A. Bakshi, A. Liu, A. Moitra, and E. Tang, *Structure learning of Hamiltonians from real-time evolution*, in *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE Computer Society, Los Alamitos, CA, 2024), pp. 1037–1050, [10.1109/FOCS61266.2024.00069](#).
- [18] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, *Theory of trotter error with commutator scaling*, *Phys. Rev. X* **11**, 011020 (2021).
- [19] S. Blanes and F. Casas, *A Concise Introduction to Geometric Numerical Integration* (CRC Press, Boca Raton, FL, USA, 2026).
- [20] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, *Efficient quantum algorithms for simulating sparse Hamiltonians*, *Commun. Math. Phys.* **270**, 359 (2007).
- [21] S. Gustavsson, O. Zwiernik, J. Bylander, F. Yan, F. Yoshihara, Y. Nakamura, T. P. Orlando, and W. D. Oliver, *Improving quantum gate fidelities by using a qubit to measure microwave pulse distortions*, *Phys. Rev. Lett.* **110**, 040502 (2013).
- [22] M. Jerger, A. Kulikov, Zénon Vasselin, and A. Fedorov, *In situ characterization of qubit control lines: A qubit as a vector network analyzer*, *Phys. Rev. Lett.* **123**, 150501 (2019).
- [23] M. A. Rol, L. Ciorciaro, F. K. Malinowski, B. M. Tarasinski, R. E. Sagastizabal, C. C. Bultink, Y. Salathe, N. Haandbaek, J. Sedivy, and L. DiCarlo, *Time-domain characterization and correction of on-chip distortion of control pulses in a quantum processor*, *Appl. Phys. Lett.* **116**, 054001 (2020).
- [24] C. Hellings, N. Lacroix, A. Remm, R. Boell, J. Herrmann, S. Lazăr, S. Krinner, F. Swiadek, C. K. Andersen, C. Eichler *et al.*, *Calibrating magnetic flux control in superconducting circuits by compensating distortions on timescales from nanoseconds up to tens of microseconds*, *Phys. Rev. Res.* **7**, 043142 (2025).
- [25] D. Hangleiter, I. Roth, J. Fuksa, J. Eisert, and P. Roushan, *Robustly learning the Hamiltonian dynamics of a superconducting quantum processor*, *Nat. Commun.* **15**, 9595 (2024).
- [26] T. I. Andersen, N. Astrakhantsev, A. H. Karamlou, J. Berndtsson, J. Motruk, A. Szasz, J. A. Gross, A. Schuckert, T. Westerhout, Y. Zhang *et al.*, *Thermalization and criticality on an analogue–digital quantum simulator*, *Nature (London)* **638**, 79 (2025).
- [27] G. H. Low and I. L. Chuang, *Optimal Hamiltonian simulation by quantum signal processing*, *Phys. Rev. Lett.* **118**, 010501 (2017).
- [28] G. H. Low and I. L. Chuang, *Hamiltonian simulation by qubitization*, *Quantum* **3**, 163 (2019).
- [29] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, *Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics*, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2019), pp. 193–204, [10.1145/3313276.3316366](#).
- [30] A. M. Childs and N. Wiebe, *Hamiltonian simulation using linear combinations of unitary operations*, *Quantum Inf. Comput.* **12**, 901 (2012).
- [31] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, *Simulating Hamiltonian dynamics with a truncated Taylor series*, *Phys. Rev. Lett.* **114**, 090502 (2015).
- [32] X.-M. Zhang and X. Yuan, *Circuit complexity of quantum access models for encoding classical data*, *npj Quantum Inf.* **10**, 42 (2024).
- [33] J. Haah, R. Kothari, R. O’Donnell, and E. Tang, *Query-optimal estimation of unitary channels in diamond distance*, in *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE Computer Society, Los Alamitos, CA, 2023), pp. 363–390, [10.1109/FOCS57990.2023.00028](#).
- [34] J. Haah, M. B. Hastings, R. Kothari, and G. H. Low, *Quantum algorithm for simulating real time evolution of lattice Hamiltonians*, *SIAM J. Comput.* **52**, FOCS18-250 (2021), [10.1137/18M1231511](#).
- [35] A. M. Childs and Y. Su, *Nearly optimal lattice simulation by product formulas*, *Phys. Rev. Lett.* **123**, 050503 (2019).
- [36] A. Bakshi, A. Liu, A. Moitra, and E. Tang, *Learning quantum Hamiltonians at any temperature in polynomial time*, in *Proceedings of the 56th Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2024), pp. 1470–1477, [10.1145/3618260.3649619](#).
- [37] S. Brahmachari, S. Zhu, I. Marvian, and Y. Tong, *Learning Hamiltonians in the Heisenberg limit with static single-qubit fields* (2026), [arXiv:2601.10380](#).
- [38] S. Chen, J. Cotler, and H.-Y. Huang, *Quantum probe tomography* (2025) [arXiv:2510.08499](#).

- [39] D. Aharonov, J. Cotler, and X.-L. Qi, *Quantum algorithmic measurement*, *Nat. Commun.* **13**, 887 (2022).
- [40] M. L. LaBorde and M. M. Wilde, *Quantum algorithms for testing Hamiltonian symmetry*, *Phys. Rev. Lett.* **129**, 160503 (2022).
- [41] A. She and H. Yuen, *Unitary property testing lower bounds by polynomials*, in *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, *Leibniz International Proceedings in Informatics*, Vol. 251, edited by Y. Tauman Kalai (Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2023), pp. 96:1–96:17.
- [42] A. Bluhm, M. C. Caro, and A. Oufkir, *Hamiltonian property testing*, *Quantum* **10**, 1979 (2026).
- [43] S. Arunachalam, A. Dutt, and F. Escudero Gutiérrez, *Testing and learning structured quantum Hamiltonians*, in *Proceedings of the 57th Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2025), pp. 1263–1270, [10.1145/3717823.3718289](#).
- [44] J. Kallaughner and D. Liang, *Hamiltonian Locality Testing via Trotterized Postselection*, in 20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025), *Leibniz International Proceedings in Informatics* **350**, edited by B. Fefferman (2025), 10:1, [10.4230/LIPIcs.TQC.2025.10](#)
- [45] M. Gao, Z. Ji, Q. Wang, W. Yu, and Q. Zhao, *Quantum Hamiltonian certification*, in *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2026), pp. 1424–1467, [10.1137/1.9781611978971.53](#).
- [46] A. Bluhm, M. C. Caro, F. E. Gutiérrez, A. Oufkir, and C. Rouzé, *Certifying and learning quantum Ising Hamiltonians* (2025) [arXiv:2509.10239](#).
- [47] J. Lee and M. Shin, *Optimal certification of constant-local Hamiltonians* (2025), [arXiv:2512.09778](#).
- [48] A. Bluhm, M. C. Caro, F. E. Gutiérrez, J. Lee, A. Oufkir, C. Rouzé, and M. Shin, *Certifying and learning local quantum Hamiltonians* (2026), [arXiv:2603.29809](#).
- [49] Z. Chen, J. Li, and J. Sloate *Lower bounds for learning Hamiltonians from time evolution* (2026), [arXiv:2509.20665v4](#).
- [50] C. Cedillo Vayson de Pradenne, J. Cotler, and H.-Y. Huang, *Learning Hamiltonians at long times* (2026), [arXiv:2606.05690](#).
- [51] A. Montanaro, *Learning stabilizer states by Bell sampling* (2017), [arXiv:1707.04012](#).
- [52] D. Hangleiter and M. J. Gullans, *Bell sampling from quantum circuits*, *Phys. Rev. Lett.* **133**, 020601 (2024).
- [53] R. Mathias, *Evaluating the Frechet derivative of the matrix exponential*, *Numer. Math.* **63**, 213 (1992).
- [54] F. Casas and A. Murua, *An efficient algorithm for computing the Baker-Campbell-Hausdorff series and some of its applications*, *J. Math. Phys.* **50**, 033513 (2009).