

Quantum computation of molecular structure using data from challenging-to-classically-simulate nuclear magnetic resonance experiments

Thomas E. O'Brien,^{1,*} Lev B. Ioffe,^{1,†} Yuan Su,¹ David Fushman,²
Hartmut Neven,¹ Ryan Babbush,^{1,‡} and Vadim Smelyanskiy^{1,§}

¹*Google Quantum AI, Venice, CA 90291, United States*

²*Department of Chemistry and Biochemistry, Center for Biomolecular Structure and Organization, University of Maryland, College Park, MD 20742, United States*

We propose a quantum algorithm for inferring the molecular nuclear spin Hamiltonian from time-resolved measurements of spin-spin correlators, which can be obtained via nuclear magnetic resonance (NMR). We focus on learning the anisotropic dipolar term of the Hamiltonian, which generates dynamics that are challenging-to-classically-simulate in some contexts. We demonstrate the ability to directly estimate the Jacobian and Hessian of the corresponding learning problem on a quantum computer, allowing us to learn the Hamiltonian parameters. We develop algorithms for performing this computation on both noisy near-term and future fault-tolerant quantum computers. We argue that the former is promising as an early beyond-classical quantum application since it only requires evolution of a local spin Hamiltonian. We investigate the example of a protein (ubiquitin) confined in a membrane as a benchmark of our method. We isolate small spin clusters, demonstrate the convergence of our learning algorithm on one such example, and then investigate the learnability of these clusters as we cross the ergodic-MBL phase transition by suppressing the dipolar interaction. We see a clear correspondence between a drop in the multifractal dimension measured across many-body eigenstates of these clusters, and a transition in the structure of the Hessian of the learning cost-function (from degenerate to learnable). Our hope is that such quantum computations might enable the interpretation and development of new NMR techniques for analyzing molecular structure.

I. INTRODUCTION

Quantum computing researchers are struggling to find near-term, ‘beyond-classical’ applications of quantum computers: problems whose solution has scientific or commercial value but that cannot be solved on classical devices alone. Fault-tolerant (FT) quantum computers able to solve valuable beyond-classical problems in chemistry [1–4] and materials science [5–7] are predicted to be years away, leaving us in the noisy intermediate scale quantum (NISQ) era [8]. And, while initial quantum experiments beyond classical [9, 10] or on the beyond-classical boundary [11, 12] have proven to be of great interest, it has been difficult to extend these to solve practical problems in other fields. This is due to the large error rates on quantum computers, and to the overhead that comes from mapping non-native problems onto a quantum computer. Here, non-native means problems beyond studying statics and dynamics of local spin systems, such as fermionic quantum simulation, linear algebra, or optimization. Algorithms to solve these problems incur costs in circuit compilation [13–16] and measurement [17–20], creating a gap in hardware requirements before these can achieve beyond-classical results. Before this gap is crossed, it makes sense to focus on those practical applications that do not require this overhead.

Nuclear magnetic resonance (NMR) spectroscopy has been a lauded cornerstone of analytical and organic chemistry since its discovery 80 years ago [21, 22], being applicable to any molecule or material containing atoms with non-zero nuclear spin. In an NMR experiment, these spinful nuclei are excited by a radio-frequency pulse and allowed to evolve for some period of time under a typically strong (multiple Tesla) magnetic field, yielding a small time-dependent response in said magnetic field. This response, or free induction decay, contains information about its generating nuclear spin Hamiltonian, which itself contains information about the chemical structure of the molecule or material under probe. The nuclear spin Hamiltonian is typically a strongly interacting quantum Hamiltonian due to its strong dipolar coupling. From accurate knowledge of the dipolar couplings between different spins, one can infer the real-space molecular structure. However, when molecules are free to quickly rotate (e.g., when in solution) this interaction averages out, leaving only chemical shifts (local fields) and a weak electron-mediated Heisenberg coupling term [23]. Such Hamiltonians can be easily classically analysed (sometimes even by intuition), one of the reasons why NMR has shown such success to this day. However, NMR spectra are considerably more complex for systems that are not free to tumble in all directions (as they are in a solution field [24]): e.g., solid-state materials [25, 26], molecules in gels [27], stuck on surfaces [28] or in membranes [29, 30]. Spectrum prediction can also pose a challenge for experiments operating in low magnetic field such as ‘zero-field NMR’ experiments [31–34], which promise more affordable and practical NMR technologies not requiring huge magnetic fields. These

* teobrien@google.com

† ioffel@google.com

‡ babbush@google.com

§ smelyan@google.com

systems present a region of parameter space where data cannot be analysed by classical computers, yielding a potential area for beyond-classical quantum computation. This ability to analyze classically intractable data sets could enable new types of NMR techniques to characterize previously difficult-to-characterize systems.

Quantum dynamics are generated by a system’s Hamiltonian; from a sufficient set of experimental data, it should be possible to learn the Hamiltonian that generated it. Hamiltonian learning is the inverse problem to predicting experimental outcomes given a system’s Hamiltonian [35], and is well-established in the field of quantum information as an approach to device characterization. Sufficiently-small devices may be characterized via classical post-processing of experimental data via Bayesian [36, 37] or machine-learning methods [38–40]. Exact classical methods are intractable in large systems whenever the forward problem becomes beyond-classical, but this can be avoided when possible by careful experiment design. For example, given sufficient control one can dynamically decouple a small subsystem from its environment, even in the presence of large background magnetic fields, and then learn the global structure piece-by-piece [41–43]. It is also possible to learn Hamiltonians from expectation values of thermal or long-time average states, which are more easy to classically approximate [44, 45]. These methods also allow one to learn a Hamiltonian from highly accurate measurements at ultra-short periods of time, where $e^{iHt} \simeq 1 + iHt$ is a nearly-exact approximation. However, when none of the above methods are viable, the Hamiltonian learning problem becomes classically challenging, giving a potential beyond-classical quantum computing application. This still requires that the experiment yield sufficient data that the Hamiltonian can be inferred.

If a quantum system is chaotic, or ergodic, following some perturbation its state will explore its entire Hilbert space, showing little dependence on the precise Hamiltonian parameters and washing out long-time correlation functions [46]. This suggests that the criteria for a system to be learnable is that it correspond to the breakdown of ergodicity. This breakdown has been well-studied in many-body physics for many years, most famously in the case of many-body localization [47, 48], and corresponds to many interesting phenomena such as area law entanglement scaling [49, 50] and the emergence of fractal many-body wavefunctions [51–55]. To the best of our knowledge, little has been done to tie these fundamental physics concepts to the notion of learnability of a quantum system.

When considering quantum Hamiltonian learning as a beyond-classical experiment, there are actually two classically-intractable quantum computations being performed. The first is the Hamiltonian evolution itself (e.g. NMR experiment), which is an analog experiment being performed by the spectrometer and the sample. This produces a set of data which is beyond-classical whenever the experiment is beyond-classical. (Beyond-classical does

not necessarily require the experiment to be BQP-hard; many of the algorithms that we will consider in this work lie in the DQC-1 complexity theory class [41, 56].) A key part of this work lies in identifying those NMR experiments that are classically challenging to simulate, which is one quality distinguishing our work from previous suggestions to study classically tractable NMR signals on a quantum device [57]. The second classically-intractable quantum computation is to learn the quantum Hamiltonian from the experimental data, which can be executed on a digital quantum computer. We assume that for the foreseeable future the connection between the spectrometer and the computer are classical. This prevents the quantum Hamiltonian learning technique of Ref. [58–60], or algorithms that require access to the quantum state [61, 62] being implemented. (Access to a quantum connection between spectrometer and computer would be of immense interest if it could be achieved, as it could provide additional exponential speedups [62].) Alternatively, one can approach this problem by generating classically hard spectra (from a quantum simulation or NMR experiment) and using this to train a classical machine agent to infer Hamiltonians from spectra in the same phase (in a manner similar to Ref. [63]). However, this requires access to this additional training data. Given a sample and an NMR spectrometer, chemists have many techniques at their disposal to infer molecular structure without requiring quantum computing assistance: magic angle spinning [64, 65], dipolar decoupling and recoupling pulse schemes [66–69], polarization transfer methods [70], and choices of different spin species and use of heteronuclear NMR couplings [71]. However, these methods throw away or approximate information about a system that is potentially valuable for characterization. Here we show that quantum computers can provide an extra tool in the NMR toolbox, and in doing so open up a realm of novel NMR experiments that have not been available before.

In this paper we propose a protocol to learn the nuclear spin Hamiltonian of a molecular or material system using a digital quantum computer and time-resolved measurements from an NMR spectroscopy experiment. We expect this protocol to present a beyond-classical quantum computation (in lieu of classical access to additional data such as a training set for a machine-learning algorithm) when the dataset from the NMR experiment is hard to classically simulate. In section II, we design quantum algorithms to estimate the cost function, Jacobian and Hessian of the learning problem, giving implementations for both NISQ and FT cost models. We then attempt to identify those systems and situations where a beyond-classical application can be found, following some general discussions of learnability in section III. In section V, we identify NMR spectroscopy of proteins within cell walls or other membranes as one potential application, as the physical pinning of these systems within a membrane prevents the tumbling that would wash away strong correlations in solution. In section V A, we study an example protein, ubiquitin, as a benchmark with known

molecular structure. We identify sets of clusters of ^1H spins within this molecule with strong intra-cluster coupling and weak coupling to the environment, that should produce a strongly-coupled signature able to be studied by a quantum device. We calculate the multifractal dimension of small clusters to study their ergodic-localized phase transition as the dipolar term is suppressed (e.g. by magic angle spinning or decoupling pulse schemes), and find the ‘quantum-feasible’ region to require a suppression factor between around $\alpha = 5$ and $\alpha = 100$ (assuming a background magnetic field of $23.5T$, corresponding to a proton frequency of 1 GHz). This gives a large window within which quantum computers could be expected to assist in NMR interpretation. We demonstrate the application of our learning algorithm to small clusters within this region, demonstrating its convergence on a small spin cluster in the presence of sampling noise. Finally, we show a direct correspondence between the onset of localization (as measured in the multifractal dimension of ubiquitin spin clusters as their dipolar coupling is suppressed) and the onset of learnability (as measured by the analytical Hessian of our learning problem at the global minimum). To the best of our knowledge, this is the first demonstration of a clear connection between the notion of fractal eigenstates in a quantum system and its learnability by quantum or classical means.

II. QUANTUM-ASSISTED HAMILTONIAN LEARNING

We now consider the problem of learning a Hamiltonian H of some system from a set of time-resolved experimental data $S_x(t)$, where, x indexes different sets of experiments. As the Hamiltonian of a system dictates the time dynamics, this is a natural thing to learn from time series data; we will discuss later how one can infer molecular structure from a Hamiltonian of nuclear spins. Each experiment consists of an initial state preparation ρ_x , time evolution by H plus an external time-dependent driving field $H_x(t)$, and final measurement of some observable O_x . The signal $S_x(t)$ is then given by

$$S_x(t) = \text{Trace}[U_x(t, 0)\rho_x U_x^\dagger(t, 0)O_x], \quad (1)$$

where $U_x(t_2, t_1)$ is the time evolution operator generated by the Hamiltonian $H + H_x(t)$ from $t = t_1$ to t_2

$$U_x(t_2, t_1) = \mathfrak{T} \exp \left\{ i \int_{t_1}^{t_2} [H + H_x(t)] dt \right\}, \quad (2)$$

where \mathfrak{T} is the time-ordering operator. Note that one may consider $\rho_x = \rho$ and $O_x = O$ independent of the experiment x by encoding preparation and measurement terms onto the driving Hamiltonian $H_x(t)$. (As we will discuss later, this is often an accurate description of a real-world NMR experiment.) Alternatively, if the driving Hamiltonian $H_x(t)$ is only used for preparation and measurement, this may be encoded entirely in ρ_x and O_x , setting $H_x(t) = 0$ and $U_x(t_2, t_1) = U(t_2, t_1) = e^{iH(t_2-t_1)}$.

To define our learning problem, we write our system Hamiltonian in the form

$$H = \sum_n h_n V_n, \quad (3)$$

where h_n are a set of parameters and V_n a set of Hermitian operators. We then consider the case where some or all of the h_n are unknown, and we wish to estimate these by a set $\{\bar{h}_n\}$; this defines our learning problem. Given some prior $\{h_n^{(0)}\}$ with standard deviation w_n , and assuming that each datapoint $S_x(t)$ is drawn from a normally-distributed experimental population with standard deviation $\sigma_{x,t}^2$, the maximum-likelihood estimation of the true parameters can be found by minimizing the cost function

$$C[\bar{H}] = \sum_n \frac{(\bar{h}_n - h_n^{(0)})^2}{2w_n^2} + \sum_{x,t} \frac{(\bar{S}_x(t) - S_x(t))^2}{2\sigma_{x,t}^2}, \quad (4)$$

where $\bar{S}_x(t) = \text{Trace}[\bar{U}_x(t, 0)\rho_x \bar{U}_x^\dagger(t, 0)O_x]$ is the estimated signal with our estimates of the parameters \bar{h}_n . (Throughout this work, we use bars to denote quantities derived from estimated parameters rather than hidden ones.) Though $\bar{S}_x(t)$ cannot be estimated on a classical device, implementing it on a quantum computer simply requires a circuit to simulate the time evolution $\bar{U}_x(t, 0)$. However, performing such an optimization gradient-free on a higher-dimensional surface is a costly endeavour. The first key result of this work is to give a practical form for the gradient and Hessian of Eq. 4

$$\begin{aligned} \frac{dC[H]}{d\bar{h}_n} &= \frac{1}{w_n^2} (\bar{h}_n - h_n^{(0)}) \\ &+ \sum_{x,t} \frac{i}{\sigma_{x,t}^2} [\bar{S}_x(t) - S_x(t)] \bar{J}_x^n(t) \end{aligned} \quad (5)$$

$$\begin{aligned} \frac{d^2C[H]}{d\bar{h}_n d\bar{h}_m} &= \frac{\delta_{nm}}{w_n^2} - \sum_{x,t} \frac{1}{\sigma_{x,t}^2} \\ &\times \left[\bar{J}_x^n(t) \bar{J}_x^m(t) + 2[\bar{S}_x(t) - S_x(t)] \bar{K}_x^{n,m}(t) \right], \end{aligned} \quad (6)$$

where we define

$$\bar{J}_x^n(t) = \int_0^t ds \bar{j}_x^n(t, s), \quad (7)$$

$$\bar{j}_x^n(t, s) = \text{Trace} \left[O_x [\bar{V}_{n,x}(t, s), \bar{\rho}_x(t)] \right], \quad (8)$$

$$\bar{K}_x^{n,m}(t) = \int_0^t ds \int_0^s dr k_x^{n,m}(t, s, r), \quad (9)$$

$$\begin{aligned} &\bar{k}_x^{n,m}(t, s, r) \\ &= \text{Trace} \left[O_x [\bar{V}_{n,x}(t, s), [\bar{V}_{m,x}(t, r), \bar{\rho}_x(t)]] \right], \end{aligned} \quad (10)$$

and $\bar{V}_{n,x}(t, s) = \bar{U}_x(t, s)V_n\bar{U}_x^\dagger(t, s)$ is the (estimated) operator V_n evolved forwards in time from s to t , and $\bar{\rho}_x(t) = \bar{U}_x(t, 0)\rho_x\bar{U}_x^\dagger(t, 0)$ is the (estimated) state ρ_x at

time t . As the term dependent on $\bar{K}_x^{n,m}(t)$ in Eq. 6 disappears in the limit $\bar{S}_x(t) \rightarrow S_x(t)$, it may be practical when near the global minimum of $C[H]$ to approximate

$$\frac{d^2 C[H]}{d\bar{h}_n dh_m} = \frac{\delta_{nm}}{w_n^2} - \sum_{x,t} \frac{\bar{J}_x^n(t) \bar{J}_x^m(t)}{\sigma_{x,t}^2}, \quad (11)$$

which may be obtained at no extra cost to the gradient estimation (assuming $\bar{S}_x(t)$ and $\bar{J}_x^n(t)$ are measured to the same relative precision). This is important, as we can approximate the covariance matrix Σ of our final estimation of the $\{\bar{h}_n\}$ as

$$\Sigma = [\nabla_{\bar{h}}^2 C]^{-1}. \quad (12)$$

III. QUANTUM AND CLASSICAL LEARNABILITY

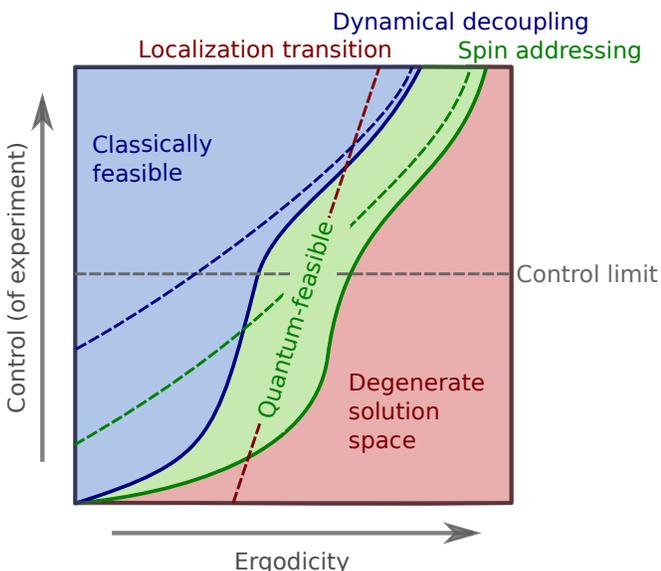


FIG. 1. Cartoon of the “phases of learnability” of a quantum Hamiltonian. In the red region, the set of experimental data is insufficient to distinguish candidate Hamiltonians. This becomes increasingly likely for systems beyond the localization transition (red dashed line), as after this point experiments such as local correlators tend to provide little information about the system structure. In the blue region, the experimental data is sufficient to learn the system’s structure, but the data processing may be achieved classically, rendering a quantum computer unnecessary. This classical processing may either be achieved by the system being well-approximated by a classically-computable model, or by the experiment having sufficient control (grey dashed line) to isolate smaller subsystems. At the limit of good control, an experiment has the ability to dynamically decouple individual terms (blue dashed line), after which techniques such as those in Refs. [41, 43] can be used to estimate terms individually. A simpler limit is to simply have the ability to address individual (or small frequency regions of) spins in the system (green dashed line), which renders the system more learnable.

In the following sections, we will describe how the above may be developed into a complete quantum-assisted algorithm to learn Hamiltonian parameters via gradient optimization. However, there are two issues that limit the usefulness of our proposal to implement these algorithms on a quantum computer. We summarize these issues and give a sketch of the region in parameter space where the techniques we develop in this work are relevant in Fig. 1

The first issue is that the experimental data taken may not contain enough information to learn the desired couplings at all. A large amount of nature is chaotic, or ergodic — where the system tends to approximately explore its entire phase space. In such ergodic systems, properties such as local correlation functions die off quickly without any dependence on the internal structure [11, 46]. When this is the case the set of e.g. local spin-spin correlations generated by two different Hamiltonians may be indistinguishable up to corrections smaller than any experimental noise. However, this may be mitigated somewhat given the ability to perform more complicated experiments (e.g. magic angle spinning [64, 65], decoupling pulses [69], or other composite pulse sequences [72]). We summarize these notions going from left to right in Fig. 1: as the system crosses the localization transition it quickly becomes unlearnable via quantum or classical means unless one has greater control over the experiment (allowing more complex correlations to be measured). Here, the x-axis denotes a rough measure of ergodicity of an arbitrary system (e.g. the localization length left of the localization transition, and the rate of entanglement growth on the right). This can in principle be changed by experimental control (i.e. by suppressing the dipolar interaction strength), which suggests that the localization transition should not be a vertical line.

The second issue we face is that a quantum computer may not be required to interpret the spectrum of a given Hamiltonian: it may be entirely possible to solve the learning problem with a classical device. While this reduction in complexity is a good thing for the experiment in question; nevertheless it limits the utility of using quantum computers in that context. The Hamiltonian learning problem may be solved classically for two reasons. Firstly, the Hamiltonian itself may be classically simulatable, or approximately simulatable, for example due to the system being strongly localized (making perturbation theory sufficiently accurate) or otherwise nearly integrable. The second reason why the Hamiltonian learning problem may be solved classically is if the experiment is controllable enough to isolate smaller subsystems or otherwise simplify the system. This can for instance be achieved if one has the ability to spatially resolve individual spins with a magnetic field. In this case it is possible to apply dynamic decoupling pulse sequences that isolate local Hamiltonian terms while cancelling out the remainder in an experiment [41–43], making local characterization possible. We summarize these notions

along the y-axis of Fig. 1; as one gains more control over an experiment it becomes possible to learn Hamiltonians in more systems classically, until the barrier of dynamic decoupling is reached and all Hamiltonians are classically learnable.

These two concerns leave our quantum Hamiltonian learning algorithm with only a ‘‘Goldilocks’’ zone of applicability. We are interested in those experiments where we have some control over our input state and Hamiltonian, but not those where enough control is available to isolate individual terms. We are also interested in those experiments where our system is somewhat ergodic, but not completely. Experimentally, this can be summarized by saying that we can study those systems where some signal can be extracted, but where that signal is complicated by complicated features that cannot be easily understood perturbatively.

A. Robustness of learning

Assuming access to the derivatives in Sec. II, one may in principle solve the Hamiltonian learning problem using many well-known gradient-based or Hessian-based optimization techniques. (Using the approximate Hessian in Eq. 11 for minimization results in the well-known Levenberg-Marquardt algorithm [73, 74].) One may ask whether this can be made robust, to avoid being stuck in local minima. Similar questions have been asked and answered previously for single parameter estimation [41], quantum phase estimation [75], and device calibration [76, 77]. We give a sketch of an argument here for the robustness of our algorithm given a time-independent Hamiltonian ($H_x(t) = 0$) that works under the assumption that we start with an initial guess $h_n^{(0)}$ sufficiently close to our true parameters h_n . We stress that this is not a proof, and examining the landscape around the global minimum of our learning problem is a clear target for future work.

In this case, we may work in the eigenbasis $|\xi_a\rangle$ of the system Hamiltonian $H|\xi_a\rangle = E_a|\xi_a\rangle$. Inserting two resolutions of the identity, our signal then takes the form $S_x(t) = \sum_{a,b} s_x^{a,b}(t)$, where

$$s_x^{a,b}(t) = \langle \xi_a | \rho_x | \xi_b \rangle \langle \xi_b | O_x | \xi_a \rangle e^{i(E_a - E_b)t}. \quad (13)$$

If our estimate deviates by some parameter $h_n \rightarrow \bar{h}_n = h_n + \delta$, then to lowest order in perturbation theory our

estimated signal takes the form

$$\bar{S}_x(t) = \sum_{a,b} \bar{s}_x^{a,b}(t) + \delta X, \quad (14)$$

$$\bar{s}_x^{a,b}(t) = s_x^{a,b}(t) e^{it\delta(\langle \xi_a | V_n | \xi_a \rangle - \langle \xi_b | V_n | \xi_b \rangle)} \quad (15)$$

$$X = \sum_{a,b,c,a \neq c} \frac{e^{-iE_b t} (e^{iE_a t} - e^{iE_c t})}{E_a - E_c} \times \langle \xi_c | \rho_x | \xi_b \rangle \langle \xi_b | O_x | \xi_a \rangle \langle \xi_a | V_n | \xi_c \rangle + \sum_{a,b,c,b \neq c} \frac{e^{iE_a t} (e^{-iE_b t} - e^{-iE_c t})}{E_b - E_c} \times \langle \xi_a | \rho_x | \xi_c \rangle \langle \xi_c | V_n | \xi_b \rangle \langle \xi_b | O_x | \xi_a \rangle. \quad (16)$$

Note here that X is independent of δ . The second term in our cost function (Eq. 4) takes the form

$$\sum_{x,t} \frac{(\bar{S}_x(t) - S_x(t))^2}{2\sigma_{x,t}^2} = \sum_{x,t} \frac{1}{\sigma_{x,t}^2} \left[\delta X + \sum_{a,b} s_x^{(a,b)}(t) \left(1 - e^{it\delta(\langle \xi_a | V_n | \xi_a \rangle - \langle \xi_b | V_n | \xi_b \rangle)} \right) \right]^2. \quad (17)$$

This oscillates as a function of δ with a frequency bounded by $4t \max \langle \xi_a | V_n | \xi_a \rangle$, which is independent of the system size. This implies that we know that local minima in our parameter space must be separated by at least $[4t \max \langle \xi_a | V_n | \xi_a \rangle]^{-1} \leq [4t \max_n \|V_n\|]^{-1}$. Flipping this around, let us suppose we know our initial guess $h_n^{(0)}$ of our parameters h_n lies within some $\delta \leq \sum_n |h_n - h_n^{(0)}|$, we can perform robust Hamiltonian learning by first learning H from only experiments at times

$$t < t_{\max} = \frac{\pi}{4\delta \max_n \|V_n\|}. \quad (18)$$

After converging on this data we may estimate the variance of our parameter guess (using Eq. 6), refine our estimate of δ , and increase the range of allowed t . Assuming that estimation at each t_{\max} yields an error $\delta_{\text{new}} \leq c\pi/(4 \max_n \|V_n\| t_{\max})$ for some $c < 1$, repeating this procedure over multiple orders will converge to some final error ϵ in $\mathcal{O}(\log(1/\epsilon))$ iterations of this procedure. If our initial uncertainty $\sum_n |h_n - h_n^{(0)}|$ grows with the system size (which can simply be due to the increase in the number of parameters), this will necessarily shrink the initial choice of t_{\max} by the same amount. However, for a local Hamiltonian we expect the number of parameters and the magnitude of terms to only grow linearly in N .

The ability to estimate out to large times allows us to avoid learnability issues in localized systems where small long-range couplings contribute mostly to the off-diagonal part of the Hamiltonian (and thus do not affect the eigenstructure significantly). However, it is of no help in ergodic systems where signals $S_x(t)$ disappear quickly with t . As we will see, in those (ergodic) systems learning anything would be a significant challenge.

IV. QUANTUM ALGORITHMS

In this section, we outline algorithms to efficiently calculate our cost function $C[H]$ and its first and second derivatives using a quantum computer. Quantum algorithm optimization is significantly different when targeting noisy near-term vs fault-tolerant long-term devices; we will present algorithms for both situations. In either case, we will discretize the integral in Eq. 8

$$\int_0^t ds f(s) \sim \sum_{i=0}^{I-1} z_i f(s_i), \quad (19)$$

where the weights $z_i > 0$ are chosen such that $\sum_i z_i = t$. (A 2-dimensional discretization is similarly possible for the integral $K_x^{n,m}(t)$.) There are many possible methods for choosing both the weights and the points s_i ; one may use a simple trapezoidal or midpoint rule, or more complicated Gaussian quadrature methods, or one may take a Monte Carlo approach and choose the points at random. Each method incurs a discretization error that goes to 0 as $I \rightarrow \infty$. For both the FT and NISQ quantum methods that we propose, cost depends primarily on $[\sum_i z_i] = t$, and has only at most a logarithmic dependence on I . This allows us to choose our method of integration for simplicity or ease in circuit design.

A. Algorithms for near-term quantum computers

In the current NISQ era we want to run the shortest quantum circuits possible for any application. To this end, we propose estimating the signal $\bar{S}_x(t)$, and the integrands $\bar{j}_x^n(t)$ and $\bar{k}_x^{n,m}(t)$ on a quantum computer, and performing the integration and summation to yield $\frac{dC[H]}{dh_n}$ and $\frac{d^2C[H]}{dh_n dh_m}$ classically. The signal $\bar{S}_x(t)$ is already in the form where it can be read as the expectation value of a quantum state following the application of a unitary circuit (Fig. 2, top). This requires that it take the form $\text{Trace}[\mathcal{U}\Xi\mathcal{U}^\dagger\mathcal{M}]$ for a quantum state Ξ , unitary \mathcal{U} , and hermitian operator \mathcal{M} . (This can be compared directly to Eq. 1.) The integrands $\bar{j}_x^n(t)$ (Eq. 8) and $\bar{k}_x^{n,m}(t)$ (Eq. 10) are not quite of this form. However, the integrands may be put in the correct form by adding control qubits to enlarge the Hilbert space, and then using the so-called generalized Hadamard test; this results in the middle and bottom circuits of Fig. 2. This circuit construction uses the identity

$$\begin{aligned} & i\text{Trace}[O[U, \rho]] \\ &= 2\text{Trace}\left[(c-U)(|+\rangle\langle+| \otimes \rho)(c-U^\dagger)(Y \otimes O)\right], \quad (20) \end{aligned}$$

to transform a commutator into the above form (with $Y \otimes O = \mathcal{M}$, $c-U = \mathcal{U}$ and $|+\rangle\langle+| \otimes \rho = \Xi$). Here the symbol $c-U := I \oplus U$ denotes the unitary U controlled by the control qubit. These circuits require only local control of the V_n unitary (see below) and uncontrolled time

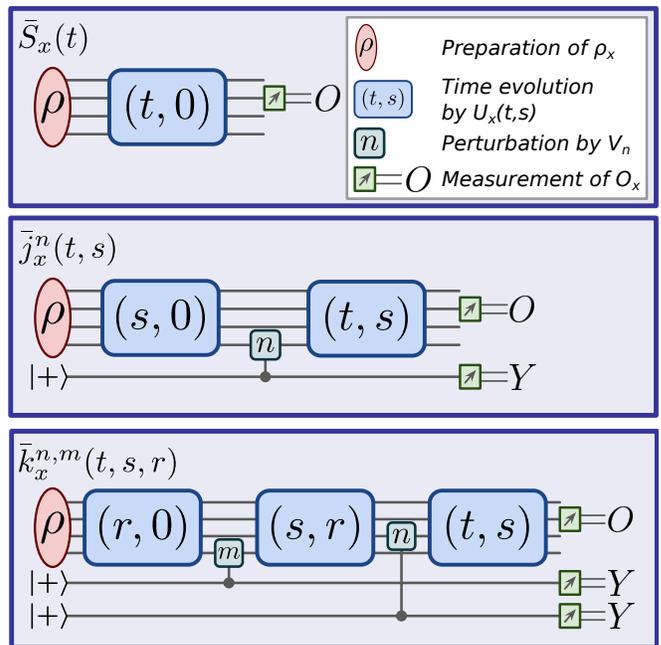


FIG. 2. Circuits to estimate the signal $\bar{S}_x(t)$ (top) and the integrands $\bar{j}_x^n(t, s)$ (middle) and $\bar{k}_x^{n,m}(t, s, r)$ (bottom) that are required to calculate the first and second derivatives of our cost function $C[H]$ (Eq. 4). For ease of viewing, we suppress labels in the circuits themselves (see legend for clarification). Circuits assume access to a preparation of ρ , and a means to simulate $U_x(t, s)$ (without control) and to implement controlled perturbations V_n . The desired integrand can be found to be the expectation value of the product of the indicated operators, which (in NISQ) must be read out by repeated preparation and measurement.

evolution, making them rather NISQ friendly (using e.g. the randomized Trotterization methods of Ref. [78] to simulate the time evolution). We assume in Fig. 2 that the terms V_n are unitary (i.e., tensor products of Pauli operators for pairs of spins), but if this is not the case one may write V_n as a linear combination of unitary operators, execute the circuits for each unitary component separately, and sum the resulting expectation values to yield the desired result [79].

The circuits in Fig. 2 assume the ability to prepare the initial states ρ_x . In the applications we consider in this paper, these will be mixed diagonal states in the computational basis. The measurement operators O_x will similarly be diagonal in the computational basis. Preparing mixed states requires that we average over many pure state preparations. Consider the case where we prepare a computational basis state $|n\rangle$, then perform a circuit U and measure the expectation values of a set of O_x in parallel, yielding a set of estimates of $\langle n|U^\dagger O_x U|n\rangle$. If we have repeated this independently for all computational basis states, and each ρ_x is diagonal in the computational

basis, we can write

$$\begin{aligned} \text{Trace}[U\rho_x U^\dagger O_x] &= \sum_{m,n} \langle n|\rho_x|m\rangle \langle m|U^\dagger O_x U|n\rangle \\ &= \sum_n \langle n|\rho_x|n\rangle \langle n|U^\dagger O_x U|n\rangle. \end{aligned} \quad (21)$$

As we know the initial distributions $\langle n|\rho_x|n\rangle$, the estimations of $\langle n|U^\dagger O_x U|n\rangle$ may be used to compute the target trace. In practice we do not need to prepare all states; it suffices to sample from a distribution proportional to $|\langle n|\rho_x|n\rangle|$. In state-of-the-art quantum experiments this presents a small difficulty, as uploading a new pulse sequence to re-prepare each state may be impractical. One solution may be to initially prepare qubits in the $|+\rangle$ basis and measure them prior to performing a simulation, which results in a new preparation each time.

Repeating the above procedure at multiple points s_i and r_i allows for parallel estimation of $\bar{J}_x^n(t)$ or $\bar{K}_x^{n,m}(t)$ via Eq. 19 (for fixed n and m). In principle in NISQ we are free to draw these times at random from the range $[0, t]$ and $[0, s]$ and choose the initial state. (In practice changing evolution times is more difficult than choosing initial states; making this scheme more practical is an important direction for future work.) In this case, each choice of starting state and time produces an independent random variable and Hoeffding's inequality may be applied. For the estimation of $\bar{J}_x^n(t)$, M repetitions of our experiment yields an estimator $\bar{J}_x^n(t)$ that satisfies

$$P\left(\left|\bar{J}_x^n(t) - \overline{\bar{J}_x^n(t)}\right| > \epsilon\right) \leq 2e^{-\frac{M\epsilon^2}{2\|O_x\|^2 t^2}}, \quad (22)$$

and the number of samples required to estimate this at a constant failure rate scales as $M = \tilde{O}(\epsilon^{-2}\|O_x\|^2 t^2)$ (where we use \tilde{O} to denote asymptotic complexity suppressing polylogarithmic factors). Similarly, for the estimation of $\bar{K}_x^{n,m}(t)$, M repetitions of our experiment yields an estimator $\bar{K}_x^{n,m}(t)$ that satisfies

$$P\left(\left|\bar{K}_x^{n,m}(t) - \overline{\bar{K}_x^{n,m}(t)}\right| > \epsilon\right) \leq 2e^{-\frac{M\epsilon^2}{\|O_x\|^2 t^4}}, \quad (23)$$

and the number of samples required to estimate this at a constant failure rate scales as $M = \tilde{O}(\epsilon^{-2}\|O_x\|^2 t^4)$. This scaling in t is to be expected, as the integrals tend to scale as $J_x^n(t) \sim t$, $K_x^{n,m}(t) \sim t^2$. (To see this, note that the diagonal terms of $\int_0^t ds \bar{V}_{n,x}(t, s)$ in the Hamiltonian basis grow linearly in t [80, 81].) This implies that we may estimate $\bar{J}_x^n(t)$ and $\bar{K}_x^{n,m}(t)$ to constant relative error with a number of samples independent of t .

We now give an analysis of the complexity of this sampling approach where we assume quantum simulation is performed under a Hamiltonian query model. The no-fast-forward theorem [82] requires a number of queries of $\Omega(t)$ to execute each of the circuits in Fig. 2, so the total gate count to estimate $\bar{J}_x^n(t)$ and $\bar{K}_x^{n,m}(t)$ to error ϵ using these methods scales at best as $\tilde{O}(\epsilon^{-2}\|O_x\|^2 t^3)$ and $\tilde{O}(\epsilon^{-2}\|O_x\|^2 t^5)$ respectively. By comparison, the total query count to estimate $\bar{S}_x(t)$ to error ϵ using the circuit in Fig. 2 is bounded asymptotically as $\tilde{O}(\epsilon^{-2}\|O_x\|^2 t)$. These estimates need to be combined to estimate the derivative in Eq. 5. To compute the total cost to estimate this to constant error, let us assume that $\bar{J}_x^n(t) \sim t$, $\bar{S}_x(t) \sim 1$, that $\sigma_{x,t} \sim \sigma$ and $\|O_x\| \sim 1$ are independent of x and t , and that we can optimize the number of repetitions of each experiment to minimize the total query count. Let us also assume that each experiment involves preparation and measurement in a commuting basis, and let us assume no covariance between parallel measurements. Then, we find the total number of queries required to estimate single derivative terms to error ϵ is bounded asymptotically by (see App. A for details)

$$\tilde{O}\left(\sigma^4 \epsilon^{-2} N_x \left[\sum_{\text{sampled } t} t^{3/2}\right]^2\right), \quad (24)$$

where N_x is the number of distinct experiments performed. The evaluation of the sum over t depends on whether $S_x(t)$ are sampled logarithmically sparsely (in which case $\sum_t t^{3/2} \sim T^{3/2} \log(T)$), or densely (in which case $\sum_t t^{3/2} \sim T^{5/2}$), where $T = \max(t)$. In the former case, the total number of queries is bounded asymptotically by $\tilde{O}(\sigma^4 \epsilon^{-2} N_x T^3)$, whilst in the latter it is bounded asymptotically by $\tilde{O}(\sigma^4 \epsilon^{-2} N_x T^5)$.

B. Improved estimation of the gradient term on a fault-tolerant quantum computer

In a fault-tolerant cost model, it is preferable to perform the integration, multiplication and summation over t and x in the second term of Eq. 5 entirely coherently. This because on an error-corrected quantum computer the key resource to minimize is the total number of gates (i.e. the sum of the gate count of each circuit applied) rather than just the depth of the longest circuit. We now outline how this may be achieved for the gradient term. Assuming that $\rho_x = W_x|0\rangle\langle 0|W_x^\dagger$, and using the fact that $\text{Trace}(A) \cdot \text{Trace}(B) = \text{Trace}(A \otimes B)$, we have

$$\sum_{x,t} \frac{i}{\sigma_{x,t}^2} \bar{S}_x(t) \bar{J}_x^n(t) \sim \left\langle 0 \otimes 0 \otimes + \left| \sum_{x,t,l} \frac{z_{l,t}}{\sigma_{x,t}^2} \mathcal{U}^0(x,t,s_{l,t}) \right| 0 \otimes 0 \otimes + \right\rangle \quad (25)$$

$$\begin{aligned} \mathcal{U}^0(x,t,s) &= W_x^\dagger U_x^\dagger(t,0) O_x U_x(t,0) W_x \\ &\otimes W_x^\dagger U_x^\dagger(s,0) \left[|0\rangle\langle 0| - |1\rangle\langle 1| \otimes V_n \right] U_x^\dagger(t,s) O_x U_x(t,s) \left[|0\rangle\langle 0| \otimes V_n + |1\rangle\langle 1| \right] U(s,0) W_x, \end{aligned} \quad (26)$$

where the approximation is the approximation from our numerical integration. One can confirm that $\mathcal{U}^0(x,t,s)$ is unitary as long as V_n and O_x are unitaries (and if this is not the case, they may be decomposed as a linear combination of unitaries themselves). The second part of the second term in Eq. 5 requires multiplying by the experimental signal $S_x(t)$. This signal then needs to be loaded onto the device; if done naively this could easily become the dominant cost in our circuit. To lower this cost, we

make the reasonable assumption that $S_x(t)$ consists of a small number $N_\omega \ll T$ of Fourier components

$$S_x(t) = \sum_{k=1}^{N_\omega} a_{x,k} \cos(t\omega_{x,k} + \phi_x), \quad (27)$$

where $\phi_x = 0$ or $\phi_x = \pi/2$ and $a_{x,k} > 0$ is expected from the $t = 0$ behaviour of our signal. Then, writing $\cos(t\omega_{x,k} + \phi_x) = \frac{1}{2}(e^{i(t\omega_{x,k} + \phi_x)} + e^{-i(t\omega_{x,k} + \phi_x)})$, we have

$$\sum_{x,t} \frac{i}{\sigma_{x,t}^2} S_x(t) \bar{J}_x^n(t) \sim \left\langle 0 \otimes +_0 \otimes +_1 \left| \sum_{x,t,l,k} \frac{z_{l,t} a_{x,k}}{\sigma_{x,t}^2} \mathcal{U}^1(x,t,s_{l,t},k) \right| 0 \otimes +_0 \otimes +_1 \right\rangle \quad (28)$$

$$\begin{aligned} \mathcal{U}^1(x,t,s,k) &= e^{-iZ_1(t\omega_{x,k} + \phi_x)} \\ &\times W_x^\dagger U_x^\dagger(s,0) \left[|0_0\rangle\langle 0_0| - |1_0\rangle\langle 1_0| \otimes V_n \right] U_x^\dagger(t,s) O_x U_x(t,s) \left[|0_0\rangle\langle 0_0| \otimes V_n + |1_0\rangle\langle 1_0| \right] U_x(s,0) W_x, \end{aligned} \quad (29)$$

where we have labeled the operations acting on the different control qubits 0 and 1. Under the above assumptions, $\mathcal{U}^1(x,t,s)$ is also a unitary operator. In the above, the $s_{l,t}$ and $z_{l,t}$ points are our integration points and weights respectively (following Sec. IV A), but allowing for the fact that the limits of integration (and thus both the points we should sample over and the total width we need to multiply by) are dependent on t . Both summations may be then block encoded using standard LCU techniques [83]. These require control registers $|x\rangle, |t\rangle, |l\rangle$, and $|k\rangle$ to encode the summation variables (as well as some additional registers we will introduce later), and SELECT and PREPARE unitaries. (We assume here that our times $|t\rangle$ have some finite binary representation.) These registers contain in turn $n_x \sim \log(N_x), n_t \sim \log(KT), n_l \sim \log(L)$ and $n_k \sim \log(N_\omega)$ qubits, where $1/K$ is the precision to which we store our times t ; see Appendix B 2 for a detailed analysis of the truncation and discretization error.

The SEL_0 unitary selects the correct $\mathcal{U}^0(x,t,s_{l,t})$ unitary to implement based on the control register; in other words,

$$\text{SEL}_0 = \sum_{x,t,l} |x\rangle|t\rangle|l\rangle\langle l|\langle t|\langle x| \mathcal{U}^0(x,t,s_{l,t})$$

. Similarly,

$$\text{SEL}_1 = \sum_{x,t,l,k} |x\rangle|t\rangle|l\rangle|k\rangle\langle k|\langle l|\langle t|\langle x| \mathcal{U}^1(x,t,s_{l,t},k)$$

. In Fig. 3, we show how this can be implemented using oracular access to $U_x(t,s)$, O_x , and W_x (which we will shortly give implementations for). The PREP_a unitaries prepare the corresponding control states

$$|\Psi_c^0\rangle = \frac{1}{\sqrt{\lambda_0}} \sum_{x,t,l} \sqrt{\frac{z_{l,t}}{\sigma_{x,t}^2}} |x\rangle|t\rangle|l\rangle, \quad (30)$$

$$\lambda_0 = \sum_{x,t,l} \left| \frac{z_{l,t}}{\sigma_{x,t}^2} \right|, \quad (31)$$

$$|\Psi_c^1\rangle = \frac{1}{\sqrt{\lambda_1}} \sum_{x,t,l,k} \sqrt{\frac{z_{l,t} a_{x,k}}{\sigma_{x,t}^2}} |x\rangle|t\rangle|l\rangle|k\rangle, \quad (32)$$

$$\lambda_1 = \sum_{x,t,l,k} \left| \frac{z_{l,t} a_{x,k}}{\sigma_{x,t}^2} \right|, \quad (33)$$

from an initial state $|0\rangle$ on the control register. (Note that the absolute value in Eq. 31 and Eq. 33 are technically unnecessary as all summands are positive.) We omit garbage registers in these steps for simplicity. Given

these, one can check that

$$\langle 0 | \text{PREP}_0^\dagger \text{SEL}_0 \text{PREP}_0 | 0 \rangle \sim \frac{1}{\lambda_0} \sum_{x,t} \frac{i}{\sigma_{x,t}^2} \bar{S}_x(t) \bar{J}_x^n(t) \quad (34)$$

$$\langle 0 | \text{PREP}_1^\dagger \text{SEL}_1 \text{PREP}_1 | 0 \rangle \sim \frac{1}{\lambda_1} \sum_{x,t} \frac{i}{\sigma_{x,t}^2} S_x(t) \bar{J}_x^n(t), \quad (35)$$

where the circuits act on the combined system and control register set. For $a = 0, 1$, we may then use the overlap estimation algorithm of Ref. [84] to estimate these values to error ϵ_a with confidence $1 - \delta$ using $\mathcal{O}(\log(\delta^{-1})\epsilon_a^{-1})$ queries to PREP_a , SEL_a ; see Appendix B 1 for a review of this algorithm. The number of additional gates used in [84] are negligible compared to the cost of block encoding. Setting $\epsilon_a = \mathcal{O}(\epsilon/\lambda_a)$ allows us to obtain an estimate of $\frac{dC[H]}{dh_n}$ that is within ϵ with confidence $1 - \delta$. $\sum_i z_{i,t} = t$ when we have sampled at time t , so if we assume that $|S_x(t)| \leq 1$ and $\sigma_{x,t} = \sigma$, we have $\lambda_1 = \lambda_2 = \frac{N_x}{\sigma^2} \sum_{\text{sampled } t} t$. The number of oracle calls to PREPARE and SELECT then scales as (in comparison to Eq. 24)

$$\tilde{\mathcal{O}} \left(\sigma^2 \epsilon^{-1} N_x \sum_{\text{sampled } t} t \right). \quad (36)$$

To make a heuristic comparison to the NISQ results, we again consider an oracular model. Our SELECT oracles require time evolution by up to $T = \max(t)$, so a quantum simulation algorithm with linear scaling in the evolution time would make $\mathcal{O}(T)$ queries to the Hamiltonian oracle. Thus, in the sparse sampling case the total number of oracle calls is bounded by $\mathcal{O}(\sigma^2 \epsilon^{-1} N_x T^2)$ (a saving of $\epsilon^{-1} T$), while in the dense sampling case the total number of oracle calls is $\mathcal{O}(\sigma^2 \epsilon^{-1} N_x T^3)$ (a saving of $\epsilon^{-1} T^2$). We expect much larger savings for simulating concrete Hamiltonians using fault-tolerant quantum algorithms, as NISQ approaches [78, 79] typically cannot achieve linear scaling in the simulation time and also have worse scaling in the target precision.

The complexity of the SELECT unitaries is dictated by the need to implement the controlled time evolution $\sum_{x,t,l} |x\rangle|t\rangle|l\rangle\langle l| \langle x| U_x(s_{l,t}, 0)$ and $\sum_{x,t,l} |x\rangle|t\rangle|l\rangle\langle l| \langle x| U_x(t, s_{l,t})$. (We will discuss the initial preparation $\sum_x |x\rangle\langle x| W_x$ later.) There are a wide range of fast quantum algorithms for simulating time evolution [82, 85–88]; which choice is optimal depends on the details of the Hamiltonian being studied. Here, we give an example implementation of the SELECT unitaries using higher-order product formulas, following the analyses developed in Ref. [89]. We expect higher-order formulas to provide the fastest approach for simulating many spin Hamiltonians [89, 90], at least asymptotically. We assume for practical purposes here that our time evolution is experiment-independent — $U_x(t, s) = U(t, s) = e^{iH(t-s)}$. This removes the need to consider the $|x\rangle$ register in our implementation of SELECT. Our implementation requires that we

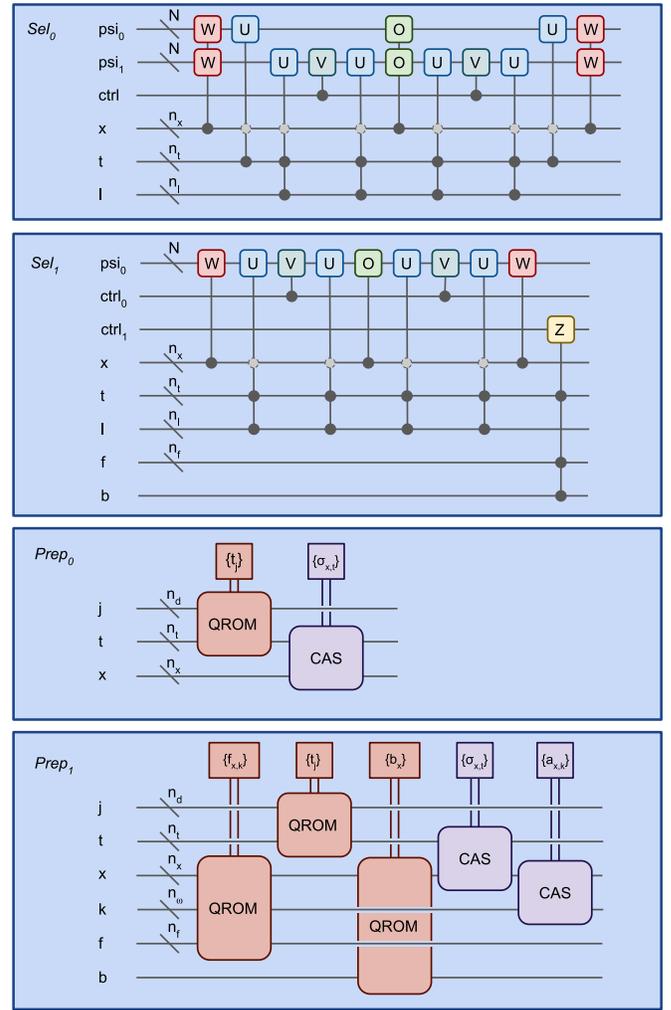


FIG. 3. Circuit diagrams of the fault-tolerant oracles SEL_0 , SEL_1 , PREP_0 and PREP_1 described in this text. See text for details. Black circles on multi-qubit registers denote complex control procedures. Square boxes denote classical input to the system via QROM and coherent alias sampling (CAS) [15]. Subscripts are omitted from gates for ease of reading. The dashed circles on the control for U in the SEL_a circuits indicate control that is only needed if the time evolution during an experiment changes between experiments.

fix the discretization of the integral in Eq. 8. We choose L points $s_{l,t} = \frac{tl}{L}$, with even weights $w_l = \frac{t}{L}$. The error in this approximation can be shown to be bounded by $\mathcal{O}\left(\frac{t^2}{L} \|O_x\| \| [H, V_n] \| \right)$, which is negligible if $L \gg t^2 \|O_x\| \| [H, V_n] \|$; see Appendix B 2 for details. The controlled time evolution part of the SELECT unitary then takes the form

$$c-U(s, 0) = \sum_{l=0}^L \sum_{t=0}^T |l\rangle|t\rangle\langle l| \langle t| e^{-\frac{it}{L} H}, \quad (37)$$

$$c-U(t, s) = \sum_{l=0}^L \sum_{t=0}^T |l\rangle|t\rangle\langle l| \langle t| e^{-\frac{i(L-l)t}{L} H}. \quad (38)$$

We implement this by dividing the total time interval $[0, t]$ into R Trotter steps and implementing a higher-order product formula in each step. To ensure that the simulation has error at most η , we take [89]

$$R = X_1 T (X_2 T / \eta)^{o(1)}, \quad (39)$$

where the lower-case $o(1)$ here represents a constant that can be taken to be arbitrarily small, and the X_1 and X_2 coefficients depend on the system size and graph connectivity. For a linear chain of N qubits, we have $X_1 \sim 1$ and $X_2 \sim N$ [91]. By contrast, assuming a model of clustered Hamiltonians [92]

$$H = \sum_{\mathcal{K}} \sum_{k, k' \in \mathcal{K}} H_{k, k'} + \sum_{\mathcal{K} \neq \mathcal{L}} \sum_{k \in \mathcal{K}, l \in \mathcal{L}} H_{k, l}, \quad (40)$$

where $H_{k, l} \ll H_{k, k'}$ whenever $k, k' \in \mathcal{K} \neq \mathcal{L} \ni l$, and we have $X_1 \sim \Lambda_{\text{ind}}$, $X_2 \sim \Lambda$, where

$$\Lambda_{\text{ind}} = \max_{\mathcal{L}} \max_{l \in \mathcal{L}} \sum_{\mathcal{K}} \sum_{k \in \mathcal{K}} \|H_{k, l}\|, \quad (41)$$

$$\Lambda = \sum_{\mathcal{K}, \mathcal{L}} \sum_{k \in \mathcal{K}, l \in \mathcal{L}} \|H_{k, l}\|. \quad (42)$$

Alternatively, we can apply a partial Trotter decomposition without splitting the terms within each cluster [92]. This reduces the Trotter error to instead scale with $X_2 \sim \Lambda_{\text{int}}$, where

$$\Lambda_{\text{int}} = \sum_{\mathcal{K} \neq \mathcal{L}} \sum_{k \in \mathcal{K}, l \in \mathcal{L}} \|H_{k, l}\| \ll \Lambda. \quad (43)$$

Each cluster can then be simulated using either product formulas or more advanced quantum algorithms. We expect that such a hybrid approach can improve the runtime of our approach, but a detailed study of such an improvement is out of the scope of the present paper and will be left as a subject for future work.

To analyze how the error of quantum simulation affects the accuracy of the overlap estimation, we use the block-diagonal structure of Eqs. 37 and 38. We see that the controlled time evolution has error at most η provided that quantum simulation is performed with accuracy η . To achieve an accuracy of ϵ in the estimate of overlap, we set $\eta_a = \mathcal{O}(\epsilon/\lambda_a)$ for $a = 0, 1$, respectively. This sets the minimum number of Trotter steps R in Eq. 39.

We now explain how to add the double-control by the $|l\rangle$ and $|t\rangle$ registers to a general product formula S_p . (The single-control by the $|t\rangle$ register also required for the SEL_a oracle can be implemented by the following techniques as well.) The near-linear dependence of the evolution time in R and requirement that $L \gg \mathcal{O}(t^2)$ imply that we need $L \gg R$. This in turn implies that lt/L is not necessarily an integer multiple of t/R . For simplicity, we assume that L, R are powers of 2, and write $lt/L = rt/R + q$ for $0 \leq q < t/R$. We write $q' = qR/t < 1$, and then $lR/L = r + q'$ gives the number of integer (r) Trotter steps and the fractional remainder

q' . Because L and R are powers of two, these integers r and q' are already stored in the first $\log(R)$ and the last $\log(L) - \log(R)$ bits of the l register, and may be identified by renaming l as (r, q') . The integer part (r) determines the number of times for which $S_p(t/R)$ needs to be applied: controlling $S_p^{2^{b_r}}(t/R)$ by the b_r -th bit of the $|r\rangle$ control register and the $|t\rangle$ register implements the unitary

$$\sum_{r, t} |r\rangle |t\rangle \langle t| \langle r| S_p^r(t/R), \quad (44)$$

where here the $|t\rangle$ register dictates the angle of rotation of each component of the product formula. For example, we would directly implement $\sum_t |t\rangle \langle t| e^{it\theta Z_i}$ bit-wise, using the b_t th bit of the t register to control a rotation by $e^{i2^{b_t}/2^{n_t}\theta Z_i}$. This has a gate complexity polylogarithmic in the input parameters, and so we neglect it. We can similarly implement the final fractional Trotter step controlled by the $|q'\rangle$ register; i.e, we implement the unitary

$$\sum_{q', t} |q'\rangle |t\rangle \langle t| \langle q'| S_p(tq'/R). \quad (45)$$

This also has a similar cost that is polylogarithmic in the input parameters. The final scaling of our doubly-controlled time evolution is then identical up to logarithmic factors to the cost of implementing the Trotter evolution without control.

Our above implementation of controlled quantum simulation is developed and optimized specifically for product formulas. Another possible circuit implementation that works for not only product formulas but also more advanced quantum simulation algorithms is to use a binary representation of the evolution time and simulate for time 2^k with integer k ; see [93] for details. In any case, the complexity only scales logarithmically with the input parameters and the overhead is negligible. This justifies the comparison between the oracular models in Eq. 36 and Eq. 24, as long as the PREPARE and controlled- W_x circuits have lower costs than SELECT.

A naive implementation of the Trotter steps requires that we exponentiate all the terms in the Hamiltonian. For the clustered model in Eq. 40, this implies a gate complexity of $\mathcal{O}(N^2)$ to implement each Trotter step. However, this may be improved by truncating Hamiltonian terms with very small magnitudes or by switching to an advanced quantum simulation algorithm. We also need to synthesize the rotation gates with respect to a fault-tolerant gate set, but the overhead in the circuit synthesis is asymptotically negligible.

The PREP_0 and PREP_1 oracles require loading the coefficients in Eqs. 30 and 32 respectively onto a quantum register. As we have chosen a uniform integration measure, the integration weights for both PREP_0 and PREP_1 are independent of the value of the l -register, which may be prepared by simply applying a Hadamard gate to all qubits. The remainder of our PREPARE oracles relies heavily on the QROM and the coherent

alias sampling (CAS) technique of Ref. [15], which can be used to perform the mappings $|j\rangle|0\rangle \rightarrow |j\rangle|a_j\rangle$ and $\frac{1}{\sqrt{N_d}} \sum_j |j\rangle \rightarrow \frac{1}{\sqrt{\sum_j a_j}} \sum_j \sqrt{a_j} |j\rangle$ with $\tilde{O}(N_d)$ Toffoli gates, where N_d is the number of unique datapoints or indices j . This is important as we do not assume that our times t are chosen uniformly, so preparing the $|t\rangle$ register is non-trivial. If we index our times by some uniform index j , i.e. writing $t = t_j$, we can use QROM to map $|j\rangle|0\rangle \rightarrow |j\rangle|t_j\rangle$. (This requires the $|j\rangle$ register be of size $n_d = \log(N_d)$.) Coherent alias sampling allows us to prepare the state $\frac{1}{\sqrt{\lambda_1}} \sum_{j,x} \frac{\sqrt{t_j}}{\sigma_{x,t_j}} |j\rangle|t_j\rangle|x\rangle$, with a cost equal to the number of unique datapoints. Combining this with the prepared $|l\rangle$ register above yields the PREP_0 oracle. If $\sigma_{x,t} = \sigma_t$ (i.e. all separate experiments are performed with the same error, which is a reasonable assumption), this has an identical cost of N_d . We assume that N_d scales at worst linearly in T (i.e. for dense sampling), and so the cost of implementing the PREP_0 oracle is bounded by $\tilde{O}(T)$ and dominated in the block encoding by the additive cost of the SEL_0 oracle. The PREP_1 oracle differs from the PREP_0 oracle only by the additional amplitudes $a_{x,k}$. These may be mapped onto the device using coherent alias sampling at a cost scaling as $N_x N_\omega$. This cost is additive to the N_d cost above, and as we expect $N_x N_\omega \ll N_d$, we expect this oracle to also be dominated by the cost of SEL_1 .

As an additional part of the SEL_1 subroutine, we need to implement the controlled Z rotation $e^{-iZ_2(t\omega_{x,k} + \phi_x)}$. The classical values $\omega_{x,k}$ and ϕ_x here need to be loaded onto the quantum device. In order to do this, we rely on the QROM technique of Ref. [15]. Given the set of N_ω classical datapoints $f_{x,k} = 2\pi\omega_{x,k}$ to some fixed precision, QROM can be used to perform the mapping $|k\rangle|x\rangle|0\rangle \rightarrow |k\rangle|x\rangle|f_{x,k}\rangle$, using $\tilde{O}(N_\omega N_x)$ Toffoli gates. We can similarly map ϕ_x onto a single qubit, $|x\rangle|0\rangle \rightarrow |x\rangle|b_x\rangle$ where $b_x = 0$ if $\phi_x = 0$ and $b_x = 1$ if $\phi_x = \pi/2$, at a cost of $\tilde{O}(N_x)$ Toffoli gates. These mappings are more appropriate to implement during the PREP_1 step, so we shall insert them there. Then, in the SEL_1 subroutine, we may assume access to these registers, in which case the controlled Z rotation may be implemented by arithmetic of the same form as in product formulas at a cost polylogarithmic in the size of the $|f_{x,k}\rangle$ and $|t\rangle$ registers. We may alternatively use the phase gradient method described in [94].

It remains to describe a preparation scheme for $\rho_x = W_x|0\rangle\langle 0|W_x^\dagger$. This is especially important to consider as ρ_x is not a pure state, so it is impossible to prepare it from an initial register with a unitary operation on an N -qubit quantum register. We require W_x to be a unitary operation for the expectation value estimation algorithm, as it requires repeated access to W_x and W_x^\dagger (or equivalently the ability to reflect around ρ_x). To solve this problem, we expand the size of our system register, and prepare a purified state $|\psi_x\rangle = W_x|0\rangle$ such that for all observables O within our original system,

$\text{Trace}[O\rho_x] = \text{Trace}[O \otimes I|\psi_x\rangle\langle\psi_x|]$. This requires that we at most double the number of qubits of our system N , and all operations other than W_x and W_x^\dagger we can ignore the additional qubits. As we will see below, for applications in NMR we are mostly interested in preparing states such as

$$\rho_x = \frac{1}{2}(I + Z_{j_x}), \quad (46)$$

which is the maximally-mixed state on all qubits except qubit j_x . To achieve this with an additional N -qubits, we begin with N copies of the Bell state $\frac{1}{\sqrt{2^N}}(|00\rangle + |11\rangle)^{\otimes N}$, which can be prepared using only Clifford gates. We then perform a Toffoli gate with x and j_x -th qubit as controls and $j_x + N$ as target, followed by a Hadamard gate on the j_x -th qubit controlled by x . Each controlled Hadamard can be implemented using a single Toffoli gate [4, FIG. 17]. This prepares the state

$$|x\rangle|\psi_x\rangle = \frac{1}{\sqrt{2^{N-1}}} |0_{j_x} 0_{j_x+N}\rangle \prod_{j \neq j_x} \left(|0_j 0_{j+N}\rangle + |1_j 1_{j+N}\rangle \right), \quad (47)$$

which has our desired properties. More generally, any thermal state of the classical $1D$ Ising model can be prepared as a $2N$ -qubit thermofield double state with perfect fidelity using a depth $N/2$ circuit [95].

V. APPLICATION TO NUCLEAR MAGNETIC RESONANCE SPECTROSCOPY

We now focus on the application of our quantum learning algorithm to NMR spectroscopy. In an NMR experiment, a sample of a molecule, crystal, or other material is placed in a strong magnetic field, which interacts with the spin of any spinful nuclei. These also couple to each other through dipole-dipole and electron-mediated interactions, making the full Hamiltonian [24]

$$H = \sum_i \chi_i \mathbf{S}_i \cdot \mathbf{B} + \sum_{i \neq j} J_{i,j} \mathbf{S}_i \cdot \mathbf{S}_j + \sum_{i \neq j} \Gamma_{i,j} \left[\mathbf{S}_i \cdot \mathbf{S}_j - \frac{3}{|\mathbf{r}_{i,j}|^2} (\mathbf{S}_i \cdot \mathbf{r}_{i,j})(\mathbf{S}_j \cdot \mathbf{r}_{i,j}) \right]. \quad (48)$$

Here, $\mathbf{r}_{i,j} = \mathbf{r}_i - \mathbf{r}_j$ is the vector between the i th and j th nuclei, \mathbf{S}_i is the spin vector, \mathbf{B} is the external magnetic field, χ_i is the onsite magnetic susceptibility, $\Gamma_{i,j} = \frac{\mu_0 \gamma_i \gamma_j \hbar}{8\pi |\mathbf{r}_{i,j}|^3}$ is the dipole-dipole interaction strength, and $J_{i,j}$ is the electron-mediated coupling strength. Assuming that our nuclei are spin-1/2, $\mathbf{S}_i = \frac{1}{2}(X_i, Y_i, Z_i)$. As the coupling constants $\Gamma_{i,j}$ depend directly on the physical distance $|\mathbf{r}_{i,j}|$ between the molecular spins; knowledge of these distances is sufficient to infer the molecular geometry (modulo global translations, rotations and reflections) [97]. An NMR spectroscopy experiment follows the protocol outlined in Sec. II. The system begins at

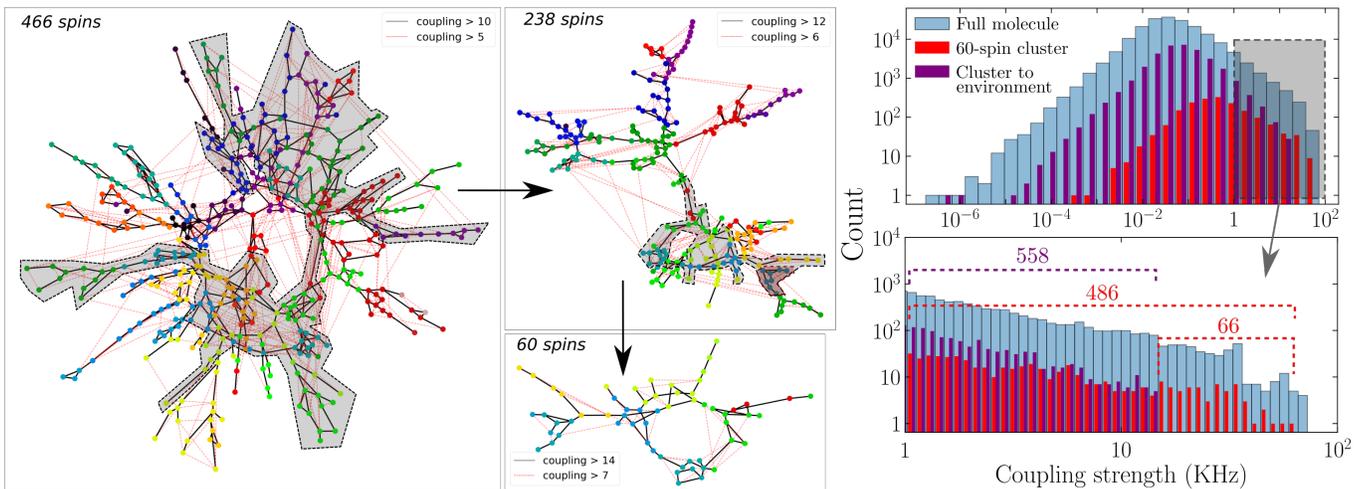


FIG. 4. Finding spin clusters within the ubiquitin protein using data from Ref. [96]. (left): a 466-spin cluster connected at a coupling of 10KHz (the total protein has 692 H atoms). (middle-top) a 238-spin subset of this cluster, connected at a coupling of 12KHz. (middle-bottom) a smaller 60-spin subset of this cluster, connected at a coupling of 14KHz. Black lines and red dashed lines between spins indicate strong and medium-strength couplings within the cluster. Red dashed region in all three clusters is an 8-spin sub-cluster studied in this text. (right-top) histogram of couplings within the 60-spin cluster and to the environment. (right-bottom) a zoom-in on the histogram tail to show the distribution of the dominant couplings.

thermal equilibrium, is perturbed by one or more magnetic field pulses, and has its free-induction decay read out (which is equivalent to measuring $\sum_i X_i$). In this formalism the starting state ρ_x and measurement operator O_x are the same for all experiments (which differ in their choice of $H_x(t)$). However, the external perturbation is often chosen to polarize the initial state and flip spins in the end, making the signal $S_x(t)$ a local spin-spin correlation measurement. (We will discuss how to implement this in a strongly correlated system shortly.)

Whether a NMR Hamiltonian lies within the ergodic, classically-feasible, or quantum-feasible regimes in Fig. 1 depends on the relative energy scales of the terms in Eq. 48. These are typically

$$\beta^{-1} \gg B\chi \gg B|\chi_i - \chi_j| \sim \Gamma_{i,j} \gg J_{i,j}, \quad (49)$$

where β is the inverse temperature of the system, and $\chi = \frac{1}{N} \sum_i \chi_i$ is the mean susceptibility. Only terms coupled to the dipolar term ($\Gamma_{i,j}$) and the electron-mediated interaction ($J_{i,j}$) generate classically challenging dynamics, so one of them must be large for us to lie outside the classically feasible region of Fig. 1. However, in solution a molecule tumbles rapidly, averaging out the dipolar term to 0. In a strong magnetic field, the remaining Hamiltonian can be treated perturbatively and solved classically, rendering our quantum learning algorithm unnecessary. Previous proposals [57] that suggested using quantum computers to learn the structure of molecules in solution suffer from this applicability issue. (Note that in zero- or ultra-low- field NMR [31–34] Eq. 49 does not hold; in these experiments $B\chi \lesssim J_{i,j}$ and quantum computers may have a role to play in learning structure.)

We propose to instead study systems where molecules are out of solution and not free to move. Indeed, NMR is

used in a wide range of systems out-of-solution: gels [27], surfaces [28], proteins in membranes [29, 30] and in the solid-state [24]. In these systems the uniform magnetic field $B\chi$ is still the dominant energy scale, and terms that do not commute with this cancel out, leaving

$$H = B \sum_i \chi_i Z_i + \sum_{i \neq j} J_{i,j} \mathbf{S}_i \cdot \mathbf{S}_j + \sum_{i \neq j} \Gamma_{i,j} (3 \cos^2(\phi_{i,j}) - 1) [\mathbf{S}_i \cdot \mathbf{S}_j - 3Z_i Z_j], \quad (50)$$

where $\phi_{i,j}$ is the angle between \mathbf{B} and $\mathbf{r}_{i,j}$. (This is commonly known as the secular approximation.) In solid state NMR experiments one typically removes the dipolar coupling by magic angle spinning; spinning at a high frequency around an angle $\theta = 54.74^\circ$ to the magnetic field. This spinning suppresses the dipolar term by a factor $(3 \cos^2(\theta) - 1) \sim 0$ as long as the spinning frequency is much higher than the dipolar coupling strength. Combining this with frequency-selective dipolar recoupling [98] has achieved remarkable success in biochemistry, (see e.g. Ref. [99]). However, this is not typically achievable to high precision for proton NMR, where dipolar couplings are typically of the order of 30 – 40 KHz [100] (by comparison, the highest frequency centrifuges are around 100 KHz [29]). Even when magic-angle spinning is combined with decoupling pulse schemes, significant residual coupling in proton-NMR spectra can be observed [29, 101, 102]. (Moreover, suppressing the dipolar coupling term removes valuable structural information about a system.) Proton NMR in solutions has achieved great success in biochemistry, but the folding of proteins can be quite different in vitro versus in vivo. For example, a large number of proteins in cell walls and mem-

branes are folded precisely according to this external environment, and lose their shape in solution [29, 30]. This suggests that learning the structure of proteins in membranes via proton NMR is a potentially valuable beyond classical quantum computing application.

Proposing to study systems with strong dipolar coupling presents a challenge in designing realistic state preparation and measurement schemes. As the temperature is larger than all energy scales in our experiment, our initial state is a thermal state

$$\rho_{\text{th}} = e^{-\beta H} \sim 1 - \sum_i \frac{\chi_i}{\beta} Z_i, \quad (51)$$

and we must perturb this in order to generate any signal at all. (The approximation here is quite good, as even for a 24.5 T background magnetic field $\beta B\chi \sim 10^{-4}$.) We have two external handles to perturb our system: the ability to apply time-dependent RF pulses to modulate the background magnetic field B , and magic-angle spinning [64, 65]. The latter is a double-edged sword; we cannot alter the direction of the magnetic field nor the spinning sample on the timescale of our system, so it is only practical to use this to suppress the dipolar coupling by a fixed amount for the entire experiment. However this may be adjusted during the experiment via dipolar decoupling and recoupling pulse schemes [66–69, 98]. As a simple example, consider the 4-step WAHUA scheme [103, 104], which consists of (1) a rest of time dt and a $\frac{\pi}{2}$ pulse around the x axis, (2) a rest of time dt and a $-\frac{\pi}{2}$ pulse around the y axis, (3) a rest of time $2dt$ and a $\frac{\pi}{2}$ pulse around the y axis, and (4) a rest of time dt , a $-\frac{\pi}{2}$ pulse around the x axis and a final rest of time dt . By performing a Magnus expansion to first order in dt on the above combined unitary scheme, the dipolar term averages to zero. (This is true of any scheme that rotates the X, Y, Z Pauli spins to the z -axis of the magnetic field for equal periods of time.) Once the dipolar field is sufficiently suppressed, individual spins may be flipped by e.g. applying a low-amplitude magnetic field oscillating at a frequency ω which addresses spins for which $B\chi_i = \omega$. This simple scheme likely difficult to achieve sensitivity below $\sim 1\text{KHz}$, but this may be improved by frequency-selective pulsing schemes [66, 105] such as the DANTE (Delays Alternating with Nutations for Tailored Excitation) scheme [106]. After applying such a scheme, the perturbation to the system is roughly

$$\rho - \rho_{\text{th}} = \beta \sum_{|\chi_i B - \omega| < \delta\omega} Z_i, \quad (52)$$

where $\delta\omega$ gives the accuracy of the technique. In an NMR experiment the signal is rescaled by the factor β in Eq. 52. However, in a quantum computation we may divide by β before performing our estimation, which makes the error requirements in our final signal independent of β . Approximations in the above can be accounted for in our learning scheme by adjusting the starting state, or by incorporating the pulse sequence into the unitary U_x , giv-

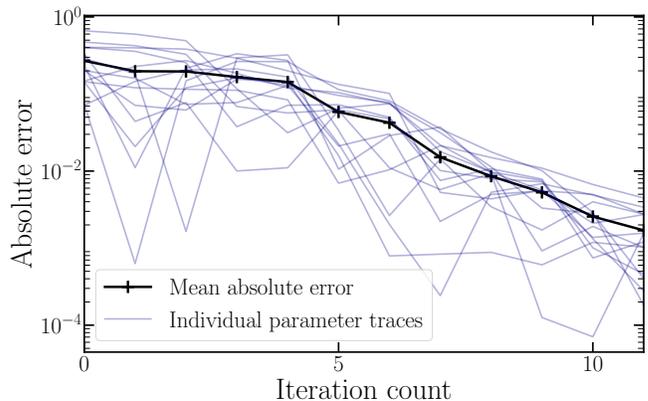


FIG. 5. Learning small couplings of a 6-spin cluster within ubiquitin given a fixed backbone from a noisy dataset with standard deviation 10^{-3} . The system is learned from a dataset of 21 points equally spaced between 0 and 2 ms (assuming a 1 GHz background magnetic field, and a dipolar suppression of $\alpha = 10$). Absolute error in individual parameters (blue faint lines) and the mean absolute error (black line) is plotted at each iteration of the CG algorithm.

ing an additional advantage over classical Hamiltonian inference.

A. Determining the structure of ubiquitin as an example application

A significant body of literature on protein structures already exists, which we can use as a benchmark for our quantum learning algorithm. The protein ubiquitin, named for its relative abundance throughout eukaryotic organisms, was discovered in 1975 [107, 108]. The structure of ubiquitin is well-known, making it a good benchmark. The protein contains over 600 protons, but these tend to cluster (in terms of their relative coupling strengths, which correspond to their location in space); we propose to divide the full molecule into smaller clusters that may be studied individually. This corresponds to assuming our experimental signature

$$S_x(t) = \sum_{\text{cluster } c} S_{x,c}(t), \quad (53)$$

and we can either subtract the contribution of individual pieces from the total signal, or learn this linear combination in a single step. The latter has a linear overhead in the number of pieces, and does not require larger quantum computers or longer circuits. In practice we do not need to determine the clustering ahead of time; clusters will present themselves as minibands in the spectrum that cannot be separated as in Eq. 53, and we may account for these in our learning algorithm by allowing our optimization to adjust the number of spins in any given band.

To investigate this clustering, we write the ubiquitous molecule as a graph (with edges weighted by the dipolar interaction strength). We take the atomic co-ordinates of ubiquitin in solution from previous NMR data (protein data bank ID 1D3Z), and proton chemical shifts from the Biological Magnetic Resonance Data Bank (BMRB), entry 17769 (both in turn taken from Ref. [96]). We then define a cluster as any connected subgraph where all edges are higher weight than any edges (in the larger graph) that point from the subgraph out. (The above definition in principle removes some edges from the subgraph / couplings from the cluster, but upon identification of the cluster we consider all couplings between the spins regardless of their size.) Such subgraphs may be found by thresholding; setting a truncation threshold V_{\min} and eliminating all couplings lower than this threshold in the nuclear spin Hamiltonian cuts the graph of spins into a set of disconnected subgraphs, each of which is a cluster. In the ubiquitin molecule (Fig. 4), by adjusting this V_{\min} we can first split off a 466 spin central core of the molecule (left), and then a 238-spin backbone (middle top), and finally a 60-spin cluster (middle bottom). Given its size, we expect the problem of learning the 60-spin cluster Hamiltonian to lie around the beyond-classical boundary. To investigate the cluster's coupling to its environment, we plot the distribution of coupling strengths (Fig. 4, right) both within the full molecule, within the 60-spin cluster, and between the cluster and the rest of the molecule. We see that the mean coupling within the ubiquitin cluster is around two to three times all couplings to the environment, and the majority of couplings to the environment are more than ten times smaller than the couplings within the cluster. We believe that this is sufficiently weak that these couplings may be treated perturbatively, though verifying this is a clear task for future study.

We now demonstrate our Hamiltonian learning algorithm for a small cluster of 6 spins in the ubiquitin molecule (Fig. 5). We assume that we have access to magic angle spinning or dipolar decoupling techniques to suppress our dipolar field by a factor $\alpha = 10$, and that the background field is 1GHz. To simulate the proposal that we know the protein backbone and are focused on learning long-range couplings, we start from a Hamiltonian (Eq. 50) where all couplings larger than some V_{\min} are known precisely, and set ourselves the task of learning smaller couplings. This leaves 12 couplings to learn (we treat $XX+YY$ and ZZ couplings independently), which we initialise at 0. To simulate sampling noise from the quantum computer, to each query of the device for $S_x(t)$ or the gradient we add a normally-distributed error term with a value of 10^{-3} . Using the conjugate gradient optimization algorithm implemented in `scipy` [109], we find that our learning problem converges to a total error of 0.008 KHz in only 11 iterations (a relative error of 0.2%).

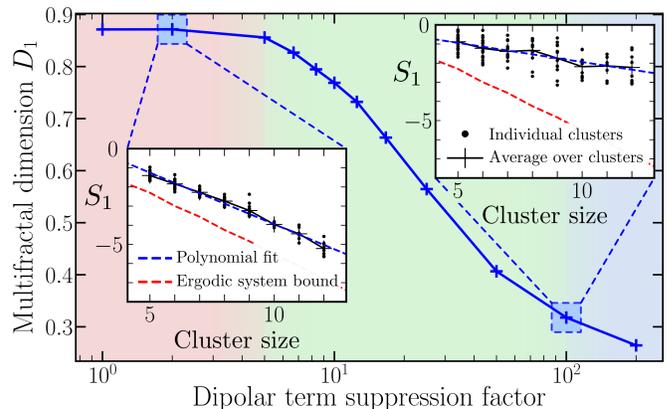


FIG. 6. Plot of the multifractal dimension of small clusters of spins in the ubiquitin molecule as the dipolar term is suppressed in a background 1GHz field. Points in the main plot are extracted from a linear fit of the mean participation entropy of the different clusters as a function of the system size, as demonstrated in the two insets for a suppression factor 2 and 100. From these fits the multifractal dimension can be estimated. The colouring of the plot indicates the expected quantum and classical learnability of the system as the dipolar term is suppressed, corresponding to the different regions in Fig. 1.

B. Learnability of spin clusters in ubiquitin

It remains to demonstrate that spin clusters in ubiquitin in a membrane or cell wall will generate an NMR signal within the region of quantum feasibility in Fig. 1. In order to study this, we investigate the participation entropy [55]

$$S_1 = \sum_{\alpha} |\psi_{\alpha}^2| \log(|\psi_{\alpha}^2|) \quad (54)$$

of eigenstates $|\psi\rangle = \sum_{\alpha} \psi_{\alpha} |\alpha\rangle$ of small clusters in the ubiquitin protein. By adjusting our threshold V_{\min} in the clustering protocol described above, we identify a collection of 127 clusters of $N = 5 - 12$ spins in ubiquitin. We do not expect the spectra of the ubiquitin Hamiltonian to be represented by a disjoint sum of some of these clusters (which would make it classically simulatable): the truncation is simply performed to give an ensemble on which to study, and we expect that the spectra from some of these clusters in the larger spin environment would differ significantly from those of the truncated piece. The dipolar couplings in these Hamiltonians may be suppressed relative to their chemical shifts by magic angle spinning or decoupling pulse schemes; we simulate this numerically by dividing the dipolar term by a variable suppression factor α . For the Hamiltonian of each cluster (Eq. 50) we calculate the mean participation entropy S_1 across the middle half of the spectrum at $P = N/2$ half-filling as we increase the dipolar term suppression by a factor α . (For odd-sized clusters of $N = 5, 7, 9, 11$ spins, we take $P = (N - 1)/2$ -filling.) In an ergodic system

where nearly all computational basis states contribute nearly equally to each eigenstate, the participation entropy scales as $S_1 = -\log[\dim(\mathcal{H})] = -\log\left(\frac{N}{P}\right)$. For the systems considered this is roughly $S_1^{(\text{ergodic})} \sim -0.63N$. By comparison, a completely localized system has a constant (or logarithmically-growing) participation entropy. As we increase α and the system crosses the localization transition, the participation entropy follows a trend $S_1 \sim D_1 S_1^{(\text{ergodic})}$, where the multifractal dimension $D_1 = D_1(\alpha)$ characterizes the fraction of the Hilbert space explored by an eigenstate [55].

In Fig. 6 we plot the multifractal dimension of our 127 spin clusters as we suppress the dipolar term by a factor α . Our quantum-feasible region corresponds to a multifractal dimension $D_1 < 1$, while our classically-feasible region corresponds to a multifractal dimension $D_1 \ll 1$. We see a clear region between a suppression factor of $\alpha \sim 5$ and $\alpha \sim 100$ where our system begins to localize and the multifractal dimension quickly drops, but the system is not completely localized and classically simple. In the thermodynamic limit this transition is discontinuous [55], but as we are interested in finite system sizes, we believe that the observed trend of D_1 as we cross this localization transition is relevant to our situation. On either side of the localization transition, local disorder will make individual clusters either more ergodic or more local than the mean, which implies that the boundaries between degenerate, quantum-feasible and classically-feasible are not sharp as a function of the dipolar suppression. (This can be seen in the insets of Fig. 6.)

The MBL-ergodic phase transition observable in the Hamiltonian eigenstructure maps immediately to the learnability of the NMR spin Hamiltonian. This can be studied in the Hessian of the Hamiltonian learning problem at complete convergence, given by Eq. 11. As the Hessian corresponds to the Fisher information of the system, small eigenvalue-eigenvector pairs (λ, \vec{v}) correspond to ‘floppy modes’ in our parameter space; linear combinations of parameters that may be adjusted in tandem without significantly altering our signal. Large eigenvalue-eigenvector pairs correspond to combinations of parameters that are well-learned. In Fig. 7 we study the typical eigenstructure of the Hessians of clusters of 5 – 8 spins. We see a clear transition that corresponds exactly to the MBL-ergodic phase transition identified in Fig. 6. On the left-hand side, corresponding to the ergodic phase, the typical maximum eigenvalue (Fig. 7, top) of the ensemble shows a clear exponential decay in system size, implying that learning in a large system will be high-impossible. Moreover, the typical participation in these systems,

$$\exp\left(-\sum_j v_j^2 \log v_j^2\right), \quad (55)$$

grows quickly (Fig. 7, bottom), implying that these

modes correspond to global data rather than specific couplings. By contrast, in the strongly localized phase, the largest eigenvalues are roughly constant in the system size, and correspond to linear combinations of only one or two parameters. As the system shifts between these two phases, we see a continuous improvement in learnability, where it appears we can learn some of but not all of the system. This can be observed in the full eigenspectrum data (Fig. 7, top inset). We note that the largest eigenvalues also correspond to smaller typical participation, which suggests that when a system is on the ergodicity boundary we can learn some local couplings rather than just global information. This result demonstrates the importance of having access to the Hessian when solving the learning problem, as it tells which of the converged parameters can be relied upon.

VI. CONCLUSION

In this work we introduced a new method for learning an unknown quantum Hamiltonian of a spin system from time-resolved measurements of the system. We constructed and costed circuits within NISQ and FT frameworks to estimate gradients of the cost function of this learning problem, finding clear asymptotic speedups when one is not constrained by poor coherence in NISQ devices. We outlined an application for these algorithms in classically intractable NMR experiments, and proposed a specific region in the NMR field (when dipolar couplings are strong and cannot be simply removed) as an area where beyond-classical computations may be very useful. Taking the ubiquitin protein as an example, we identified small clusters of spins in the larger ^1H dipolar coupling matrix, demonstrated the convergence of our learning algorithm on a toy example, and investigated the cluster-environment coupling and the learnability of the system as the dipolar coupling is suppressed.

As part of this work, we identified a direct correspondence between the ergodic-MBL localization transition and a transition to learnability. The latter was clearly observable in the structure of the Hessian of the learning problem. We believe this to be a more general result than in the NMR problems that we have studied in this work. As far as we know, this is the first link demonstrated between the onset of wavefunction fractality and the onset of learnability of the generating Hamiltonian.

Our work opens a new field of quantum computing applications, leaving clear directions for future study. Our learnability data suggests that not all dipolar coupling parameters (that encode the 3D protein structure) may be learnable using the simple spin-spin correlations measured in this work. Developing future experiments to target the parameters relevant for structure calculations in NMR systems will be highly relevant in the future.. A related question (which we have not yet determined the answer to) is pertains to the conditions (if any) under which this problem is classically difficult. While we

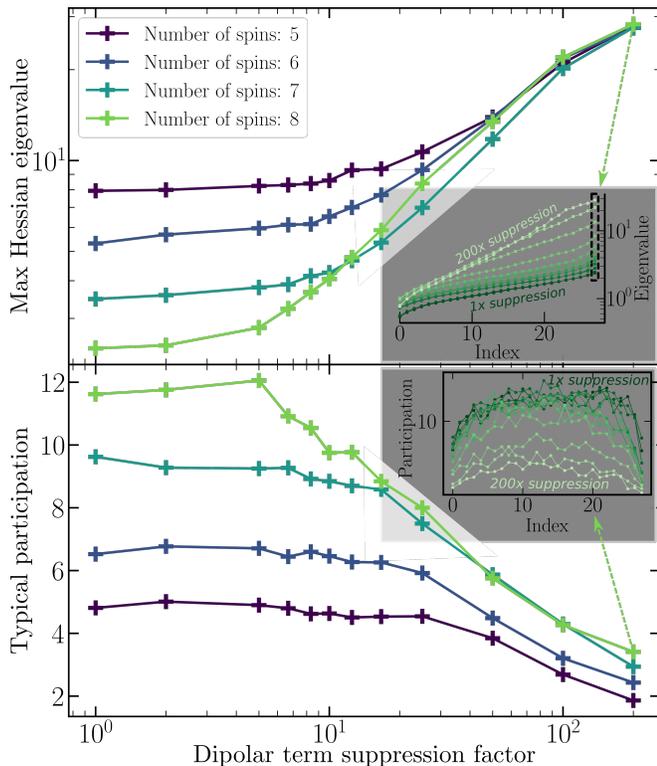


FIG. 7. Eigenvalue and eigenvector participation data for Hessians of the Hamiltonian learning problem for small spin clusters in ubiquitin. (Data is taken as the set of $\langle Z_i(t)Z_j(0) \rangle$ and $\langle X_i(t)X_j(0) \rangle$ correlators using a set of equally spaced times between $t = 0$ and $t = 5$ ms in the absence of sampling noise, implying that units of the Hessian eigenvalues are arbitrary.) (Top) typical (geometric mean) largest Hessian eigenvalue for different system sizes as the dipolar term is suppressed. Inset shows the typical Hessian spectrum for each suppression factor over clusters of 8 spins - the light green line in the main plot is taken from the indicated cut through the inset. (Bottom) typical (geometric mean) eigenvector participation (Eq. 55) for the same dataset. Inset shows the participation across the entire Hessian spectrum for clusters of 8 spins; each datapoint in the light green line in the main plot corresponds to a geometric mean over a single line in the inset.

can rely on the fact that inference of NMR spectra from strongly-correlated problems appears difficult (and the forward problem of generating the spectra is at least DQC1-hard [41, 56]), we do not know of a complexity-theory result explaining when the inverse problem is as difficult. Also, the protocol we propose in Sec. III A can, in principle, avoid trapping in local minima; however this is not optimized or fully costed. We suggest that following Ref. [41], it may be possible to define a protocol using our methods that learns at the Heisenberg limit (when learning is possible). It is also unclear precisely how well our protocol behaves in a system with a large number of unknown couplings, and whether we will have problems with vanishing gradients as our system size grows. (We note however that this is avoided somewhat in NMR problems where we start with a reasonable guess of our Hamiltonian parameters.) Another clear direction for future work is to consider learning Hamiltonian parameters from the Fourier transform of the spectra $S_x(t)$ instead of the time-resolved data, as spectral information is typically more robust to noise than amplitude data. It also remains to determine what the cost is to optimize a real-world NMR experiment from a realistic initial guess. Finally, thanks to the generality of our Hamiltonian learning techniques, it should be possible to extend these methods to the design of new NMR experiments, or to interpret data from other types of experimental procedures used to probe condensed matter, high energy physics, chemistry and materials science. We look forward to further explorations of what new possibilities for quantum experiments and data analysis these techniques can bring to the scientific community.

Acknowledgements

The authors wish to thank Igor Aleiner, Dmitry Abanin, Nicholas Rubin, Jarrod McClean, Joonho Lee, Ashok Ajoy, Joshua Goings, Nathan Wiebe and William Huggins for useful advice and discussions of this work.

-
- [1] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, Elucidating reaction mechanisms on quantum computers, *PNAS* **114**, 7555 (2017).
- [2] D. W. Berry, C. Gidney, M. Motta, J. R. McClean, and R. Babbush, Qubitization of arbitrary basis quantum chemistry leveraging sparsity and low rank factorization, *Quantum* **3** (2019).
- [3] V. von Burg, G. H. Low, T. Häner, D. S. Steiger, M. Reiher, M. Roettler, and M. Troyer, Quantum computing enhanced computational catalysis, *Phys. Rev. Research* **3**, 033055 (2021).
- [4] J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, Even more efficient quantum computations of chemistry through tensor hypercontraction, *PRX Quantum* **2**, 030305 (2021).
- [5] R. Babbush, N. Wiebe, J. McClean, J. McClain, H. Neven, and G. K.-L. Chan, Low-depth quantum simulation of materials, *Phys. Rev. X* **8**, 011044 (2018).
- [6] I. D. Kivlichan, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, W. Sun, Z. Jiang, N. Rubin, A. Fowler, A. Aspuru-Guzik, H. Neven, and R. Babbush, Improved fault-tolerant quantum simulation of condensed-phase correlated electrons via trotterization, *Quantum* **4** (2020).
- [7] Y. Su, D. Berry, N. Wiebe, N. Rubin, and R. Babbush, Fault-tolerant quantum simulations of chemistry in first quantization, *arXiv:2105.12767* (2021).

- [8] J. Preskill, Quantum computing in the nisq era and beyond, *Quantum* **2**, 79 (2018).
- [9] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [10] Y. Wu, W.-S. Bao, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan, M. Gong, C. Guo, C. Guo, S. Guo, L. Han, L. Hong, H.-L. Huang, Y.-H. Huo, L. Li, N. Li, S. Li, Y. Li, F. Liang, C. Lin, J. Lin, H. Qian, D. Qiao, H. Rong, H. Su, L. Sun, L. Wang, S. Wang, D. Wu, Y. Xu, K. Yan, W. Yang, Y. Yang, Y. Ye, J. Yin, C. Ying, J. Yu, C. Zha, C. Zhang, H. Zhang, K. Zhang, Y. Zhang, H. Zhao, Y. Zhao, L. Zhou, Q. Zhu, C.-Y. Lu, C.-Z. Peng, X. Zhu, and J.-W. Pan, Strong quantum computational advantage using a superconducting quantum processor, *ArXiv:2106.14734* (2021).
- [11] X. Mi, P. Roushan, C. Quintana, S. Mandra, J. Marshall, C. Neill, F. Arute, K. Arya, J. Atalaya, R. Babbush, J. C. Bardin, R. Barends, A. Bengtsson, S. Boixo, A. Bourassa, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, Z. Chen, B. Chiaro, R. Collins, W. Courtney, S. Demura, A. R. Derk, A. Dunsworth, D. Eppens, C. Erickson, E. Farhi, A. G. Fowler, B. Foxen, C. Gidney, M. Giustina, J. A. Gross, M. P. Harrigan, S. D. Harrington, J. Hilton, A. Ho, S. Hong, T. Huang, W. J. Huggins, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, C. Jones, D. Kafri, J. Kelly, S. Kim, A. Kitaev, P. V. Klimov, A. N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, E. Lucero, O. Martin, J. R. McClean, T. McCourt, M. McEwen, A. Megrant, K. C. Miao, M. Mohseni, W. Mroczkiewicz, J. Mutus, O. Naaman, M. Neeley, M. Newman, M. Y. Niu, T. E. O'Brien, A. Opremcak, E. Ostby, B. Pato, A. Petukhov, N. Redd, N. C. Rubin, D. Sank, K. J. Satzinger, V. Shvarts, D. Strain, M. Szalay, M. D. Trevithick, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, I. Aleiner, K. Kechedzhi, V. Smelyanskiy, and Y. Chen, Information scrambling in computationally complex quantum circuits, *ArXiv:2101.08870* (2021).
- [12] X. Mi, M. Ippoliti, C. Quintana, A. Greene, Z. Chen, J. Gross, F. Arute, K. Arya, J. Atalaya, R. Babbush, J. C. Bardin, J. Basso, A. Bengtsson, A. Bilmes, A. Bourassa, L. Brill, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, B. Chiaro, R. Collins, W. Courtney, D. Debroy, S. Demura, A. R. Derk, A. Dunsworth, D. Eppens, C. Erickson, E. Farhi, A. G. Fowler, B. Foxen, C. Gidney, M. Giustina, J. A. Gross, M. P. Harrigan, S. D. Harrington, J. Hilton, A. Ho, S. Hong, T. Huang, W. J. Huggins, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, C. Jones, D. Kafri, J. Kelly, S. Kim, A. Kitaev, P. V. Klimov, A. N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, E. Lucero, O. Martin, J. R. McClean, T. McCourt, M. McEwen, K. C. Miao, M. Mohseni, S. Montazeri, W. Mroczkiewicz, O. Naaman, M. Neeley, C. Neill, M. Newman, M. Y. Niu, T. E. O'Brien, A. Opremcak, E. Ostby, B. Pato, A. Petukhov, N. C. Rubin, D. Sank, K. J. Satzinger, V. Shvarts, Y. Su, D. Strain, M. Szalay, M. D. Trevithick, B. Villalonga, T. White, Z. J. Yao, P. Yeh, J. Yoo, A. Zalcman, H. Neven, S. Boixo, V. Smelyanskiy, A. Megrant, J. Kelly, Y. Chen, S. L. Sondhi, R. Moessner, K. Kechedzhi, V. Khemani, and P. Roushan, Observation of time-crystalline eigenstate order on a quantum processor, *ArXiv:2107.13571* (2021).
- [13] J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik, Simulation of electronic structure hamiltonians using quantum computers, *Mol. Phys.* **109**, 735 (2011).
- [14] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum random access memory, *Phys. Rev. Lett.* **100**, 160501 (2008).
- [15] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, Encoding electronic spectra in quantum circuits with linear t complexity, *Phys. Rev. X* **8**, 041015 (2015).
- [16] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, *ArXiv:1411.4028* (2014).
- [17] D. Wecker, M. B. Hastings, and M. Troyer, Towards practical quantum variational algorithms, *Phys. Rev. A* **92**, 042303 (2015).
- [18] W. J. Huggins, J. McClean, N. Rubin, Z. Jiang, N. Wiebe, K. B. Whaley, and R. Babbush, Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers, *ArXiv:1907.13117* (2019).
- [19] X. Bonet-Monroig, R. Babbush, and T. E. O'Brien, Nearly optimal measurement scheduling for partial tomography of quantum states, *Phys. Rev. X* **10**, 031064 (2020).
- [20] A. Zhao, N. C. Rubin, and A. Miyake, Fermionic partial tomography via classical shadows, *ArXiv:2010.16094* (2020).
- [21] I. Rabi, J. Zacharias, S. Millman, and P. Kusch, A new method of measuring nuclear magnetic moment, *Phys. Rev.* **53**, 318 (1938).
- [22] C. Boesch, Nobel prizes for nuclear magnetic resonance: 2003 and historical perspectives, *J. Mag. Res. Imag.* **20**, 177 (2004).
- [23] L. M. Vandersypen and I. L. Chuang, NMR techniques for quantum control and computation, *Rev. Mod. Phys.* **76** (2004).
- [24] A. McDermott and T. Polenova, *Solid State NMR Studies of Biopolymers*, eMagRes Books (Wiley, 2012).
- [25] D. Sakellariou, A. Lesage, P. Hodgkinson, and L. Emsley, Homonuclear dipolar decoupling in solid-state nmr using continuous phase modulation, *Chem. Phys. Lett.* **319**, 253 (2000).

- [26] P. K. Madhu, High-resolution solid-state nmr spectroscopy of protons with homonuclear dipolar decoupling schemes under magic-angle spinning, *Sol. State Nuc. Mag. Res.* **35**, 2 (2009).
- [27] Nonappa and E. Kolehmainen, Solid state nmr studies of gels derived from low molecular mass gelators, *Soft Matter* **12**, 6015 (2016).
- [28] F. Deng, J. Yang, and C. Ye, Solid state nmr characterization of solid surface of heterogeneous catalysts, in *Modern Magnetic Resonance* (Springer, 2008).
- [29] V. Agarwal, S. Penzel, K. Szekely, R. Cadalbert, E. Testori, A. Oss, J. Past, A. Samoson, M. Ernst, A. Böckmann, and B. H. Meier, De novo 3d structure determination from sub-milligram protein samples by solid-state 100 khz mas nmr spectroscopy, *Ang. Chem.* **53**, 12253 (2014).
- [30] D. Sinnaeve, Selective homonuclear 2d j-resolved spectroscopy, in *eMagRes*, Vol. 9 (Wiley, 2020).
- [31] M. Ledbetter, C. Crawford, A. Pines, D. Wemmer, S. Knappe, J. Kitching, and D. Budker, Optical detection of nmr j-spectra at zero magnetic field, *J. Mag. Res.* **199**, 25 (2009).
- [32] T. Theis, P. Ganssel, G. Kervern, S. Knappe, J. Kitching, M. P. Ledbetter, D. Budker, and A. Pines, Parahydrogen-enhanced zero-field nuclear magnetic resonance, *Nat. Phys.* **7**, 571 (2011).
- [33] J. W. Blanchard, M. P. Ledbetter, T. Theis, M. C. Butler, D. Budker, and A. Pines, High-resolution zero-field nmr j-spectroscopy of aromatic compounds, *J. Am. Chem. Soc.* **135**, 3607 (2013).
- [34] D. A. Barskiy, M. C. D. Tayler, I. Marco-Rius, J. Kurhanewicz, D. B. Vigneron, S. Cikrikci, A. Aydogdu, M. Reh, A. N. Pravdivtsev, J.-B. Hövener, J. W. Blanchard, T. Wu, D. Budker, and A. Pines, Zero-field nuclear magnetic resonance of chemically exchanging systems, *Nat. Comm.* **10** (2019).
- [35] M. P. da Silva, O. Landon-Cardinal, and D. Poulin, Practical characterization of quantum devices without tomography, *Phys. Rev. Lett.* **107**, 210404 (2011).
- [36] C. E. Granade, C. Ferrie, N. Wiebe, and D. Cory, Robust online hamiltonian learning, *New J. Phys.* **14** (2012).
- [37] A. Sergeevich and S. D. Bartlett, Optimizing qubit hamiltonian parameter estimation algorithms using pso, *Proc. CEC 2012* (2012).
- [38] A. Valenti, G. Jin, J. Léonard, S. D. Huber, and E. Greplova, Scalable hamiltonian learning for large-scale out-of-equilibrium quantum dynamics, *ArXiv:2103.01240* (2021).
- [39] P. Bienias, A. Seif, and M. Hafezi, Meta hamiltonian learning, *ArXiv:2104.04453* (2021).
- [40] A. A. Gentile, B. Flynn, S. Knauer, N. Wiebe, S. Paesani, C. E. Granade, J. G. Rarity, R. Santagati, and A. Laing, Learning models of quantum systems from experiments, *Nat. Phys.* (2021).
- [41] R. D. Somma and S. Boixo, Parameter estimation with mixed-state quantum computation, *Phys. Rev. A* **77** (2008).
- [42] A. Ajoy and P. Cappellaro, Quantum simulation via filtered hamiltonian engineering: Application to perfect quantum transport in spin networks, *Phys. Rev. Lett.* **110**, 220503 (2013).
- [43] S.-T. Wang, D.-L. Deng, and L.-M. Duan, Hamiltonian tomography for quantum many-body systems with arbitrary couplings, *New J. Phys.* **17**, 093017 (2015).
- [44] E. Bairey, I. Arad, and N. H. Lindner, Learning a local hamiltonian from local measurements, *Phys. Rev. Lett.* **122**, 020504 (2019).
- [45] T. J. Evans, R. Harper, and S. T. Flammia, Scalable bayesian hamiltonian learning, *ArXiv:1912.07636* (2019).
- [46] M. Srednicki, The approach to thermal equilibrium in quantized chaotic systems, *J. Phys. A* **32**, 1163 (1999).
- [47] D. Basko, I. Aleiner, and B. Altshuler, Metal-insulator transition in a weakly interacting many-electron system with localized single-particle states, *Ann. Phys.* **321**, 1126 (2006).
- [48] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, Many-body localization, thermalization, and entanglement, *Rev. Mod. Phys.* **91**, 021001 (2019).
- [49] M. Serbyn, Z. Papić, and D. A. Abanin, Local conservation laws and the structure of the many-body localized states, *Phys. Rev. Lett.* **111**, 127201 (2013).
- [50] B. Bauer and C. Nayak, Area laws in a many-body localized state and its implications for topological order, *J. Stat. Mech.* **2013**, P09005 (2013).
- [51] M. Pino, B. Altshuler, and L. Ioffe, Non-ergodic metallic and insulating phases of josephson junction chains, *Proc. Nat. Acad. Sci.* **113** (2015).
- [52] B. Altshuler, E. Cuevas, L. Ioffe, and V. Kravtsov, Non-ergodic phases in strongly disordered random regular graphs, *Phys. Rev. Lett.* **117**, 156601 (2016).
- [53] B. Altshuler, L. Ioffe, and V. Kravtsov, Non-ergodic delocalized phases on bethe lattice and random regular graphs, *Ann. Phys.* **389**, 141 (2018).
- [54] L. Faoro, M. Feigelman, and L. Ioffe, Non-ergodic extended phase of the quantum random energy model, *Ann. Phys.* **409**, 167916 (2019).
- [55] N. Macé, F. Alet, and N. Laflorencie, Multifractal scalings across the many-body localization transition, *Phys. Rev. Lett.* **123**, 180601 (2019).
- [56] E. Knill and R. Laflamme, Power of one bit of quantum information, *Phys. Rev. Lett.* **81** (1998).
- [57] D. Sels, H. Dashti, S. Mora, O. Demler, and E. Demler, Quantum approximate bayesian computation for nmr model inference, *ArXiv:1910.14221* (2019).
- [58] N. Wiebe, C. Granade, C. Ferrie, and D. Cory, Quantum hamiltonian learning using imperfect quantum resources, *Phys. Rev. A* **89**, 042314 (2014).
- [59] N. Wiebe, C. Granade, C. Ferrie, and D. Cory, Hamiltonian learning and certification using quantum resources, *Phys. Rev. Lett.* **112**, 190501 (2014).
- [60] J. Wang, S. Paesani, R. Santagati, S. Knauer, A. A. Gentile, N. Wiebe, M. Petruzzella, J. L. O'Brien, J. G. Rarity, A. Laing, and M. G. Thompson, Experimental quantum hamiltonian learning, *Nat. Phys.* **13**, 551 (2017).
- [61] H.-Y. Huang, M. Broughton, M. Mohseni, R. Babbush, S. Boixo, H. Neven, and J. McClean, Power of data in quantum machine learning, *Nat. Comm.* **12**, 1 (2021).
- [62] H.-Y. Huang, R. Kueng, and J. Preskill, Information-theoretic bounds on quantum advantage in machine learning, *Phys. Rev. Lett.* **126**, 190505 (2021).
- [63] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, Provably efficient machine learning for quantum many-body problems, *ArXiv:2106.12627* (2021).
- [64] E. Andrew, A. Bradbury, and R. Eades, Nuclear magnetic resonance spectra from a crystal rotated at high

- speed, *Nature* **182**, 1659 (1958).
- [65] I. Lowe, Free induction decays of rotating solids, *Phys. Rev. Lett.* **2**, 285 (1959).
- [66] M. H. Levitt, Composite pulses, *Prog. Nuc. Mag. Res. Spec.* **18**, 61 (1986).
- [67] A. Bielecki, A. Kolbert, and M. Levitt, Frequency-switched pulse sequences: Homonuclear decoupling and dilute spin nmr in solids, *Chem. Phys. Lett.* **155**, 341 (1989).
- [68] A. K. Paravastu and R. Tycko, Frequency-selective homonuclear dipolar recoupling in solid state nmr, *J. Chem. Phys.* **124**, 194303 (2006).
- [69] R. Tycko, Homonuclear dipolar recoupling in solid-state nmr, in *Solid-State NMR Studies of Biopolymers*, edited by A. E. McDermott and T. Polenova (Wiley, 2010) Chap. 10.
- [70] M. Ernst and B. H. Meier, Adiabatic polarization-transfer methods in mas spectroscopy, in *Solid-State NMR Studies of Biopolymers*, edited by A. E. McDermott and T. Polenova (Wiley, 2010) Chap. 12.
- [71] C. P. Jaroniec, Dipolar recoupling: Heteronuclear, in *Solid-State NMR Studies of Biopolymers*, edited by A. E. McDermott and T. Polenova (Wiley, 2010) Chap. 11.
- [72] J. A. Jones, Composite pulses in NMR quantum computation, *J. Ind. Inst. Sci.* **89** (2009).
- [73] K. Levenberg, A method for the solution of certain nonlinear problems in least squares, *Quart. Appl. Math.* **2**, 164 (1944).
- [74] D. Marquardt, An algorithm for least-squares estimation of nonlinear parameters, *SIAM J. Appl. Math.* **11**, 431 (1963).
- [75] S. Kimmel, G. H. Low, and T. J. Yoder, Robust calibration of a universal single-qubit gate-set via robust phase estimation, *Phys. Rev. A* **92**, 062315 (2015).
- [76] C. Neill, T. McCourt, X. Mi, Z. Jiang, M. Y. Niu, W. Mruzckiewicz, I. Aleiner, F. Arute, K. Arya, J. Atalaya, R. Babbush, J. C. Bardin, R. Barends, A. Bengtsson, A. Bourassa, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, J. Campero, Z. Chen, B. Chiaro, R. Collins, W. Courtney, S. Demura, A. R. Derk, A. Dunsworth, D. Eppens, C. Erickson, E. Farhi, A. G. Fowler, B. Foxen, C. Gidney, M. Giustina, J. A. Gross, M. P. Harrigan, S. D. Harrington, J. Hilton, A. Ho, S. Hong, T. Huang, W. J. Huggins, S. V. Isakov, M. Jacob-Mitos, E. Jeffrey, C. Jones, D. Kafri, K. Kechedzhi, J. Kelly, S. Kim, P. V. Klimov, A. N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, E. Lucero, O. Martin, J. R. McClean, M. McEwen, A. Megrant, K. C. Miao, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, M. Newman, T. E. O'Brien, A. Opremcak, E. Ostby, B. Pató, A. Petukhov, C. Quintana, N. Redd, N. C. Rubin, D. Sank, K. J. Satzinger, V. Shvarts, D. Strain, M. Szalay, M. D. Trevithick, B. Villalonga, T. C. White, Z. Yao, P. Yeh, A. Zalcman, H. Neven, S. Boixo, L. B. Ioffe, P. Roushan, Y. Chen, and V. Smelyanskiy, Accurately computing the electronic properties of a quantum ring, *Nature* **594**, 508 (2021).
- [77] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, A. Bengtsson, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, Y.-A. Chen, B. Chiaro, R. Collins, S. J. Cotton, W. Courtney, S. Demura, A. Derk, A. Dunsworth, D. Eppens, T. Ekl, C. Erickson, E. Farhi, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, J. A. Gross, S. Habegger, M. P. Harrigan, A. Ho, S. Hong, T. Huang, W. Huggins, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, C. Jones, D. Kafri, K. Kechedzhi, J. Kelly, S. Kim, P. V. Klimov, A. N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, M. Lindmark, E. Lucero, M. Marthaler, O. Martin, J. M. Martinis, A. Maruszczyk, S. McArdle, J. R. McClean, T. McCourt, M. McEwen, A. Megrant, C. Mejuto-Zaera, X. Mi, M. Mohseni, W. Mruzckiewicz, J. Mutus, O. Naaman, M. Neeley, C. Neill, H. Neven, M. Newman, M. Y. Niu, T. E. O'Brien, E. Ostby, B. Pató, A. Petukhov, H. Putterman, C. Quintana, J.-M. Reiner, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, D. Strain, K. J. Sung, P. Schmitteckert, M. Szalay, N. M. Tubman, A. Vainsencher, T. White, N. Vogt, Z. J. Yao, P. Yeh, A. Zalcman, and S. Zanker, Observation of separated dynamics of charge and spin in the fermi-hubbard model, *ArXiv:2010.07965* (2020).
- [78] E. Campbell, A random compiler for fast hamiltonian simulation, *Phys. Rev. Lett.* **123**, 070503 (2019).
- [79] P. K. Faehrmann, M. Steudtner, R. Kueng, M. Kieferova, and J. Eisert, Randomizing multi-product formulas for improved hamiltonian simulation, *ArXiv:2101.07808* (2021).
- [80] P. Facchi, D. A. Lidar, and S. Pascazio, Unification of dynamical decoupling and the quantum zeno effect, *Phys. Rev. A* **69**, 032314 (2004).
- [81] M. C. Tran, Y. Su, D. Carney, and J. M. Taylor, Faster digital quantum simulation by symmetry protection, *PRX Quantum* **2**, 010323 (2021).
- [82] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, Efficient quantum algorithms for simulating sparse hamiltonians, *Comm. Math. Phys.* **270** (2007).
- [83] A. M. Childs and N. Wiebe, Hamiltonian simulation using linear combinations of unitary operations, *Quant. Inf. Comp.* **12**, 901 (2012).
- [84] E. Knill, G. Ortiz, and R. D. Somma, Optimal quantum measurements of expectation values of observables, *Phys. Rev. A* **75**, 012328 (2007).
- [85] 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).
- [86] G. H. Low and I. L. Chuang, Hamiltonian Simulation by Qubitization, *Quantum* **3**, 163 (2019).
- [87] G. H. Low and N. Wiebe, Hamiltonian simulation in the interaction picture, *ArXiv:1805.00675* (2018).
- [88] J. Haah, M. B. Hastings, R. Kothari, and G. H. Low, Quantum algorithm for simulating real time evolution of lattice hamiltonians, *SIAM Journal on Computing*, **FOCS18** (2021).
- [89] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, A theory of trotter error, *Phys. Rev. X* **11**, 011020 (2021).
- [90] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, Toward the first quantum simulation with quantum speedup, *Proceedings of the National Academy of Sciences* **115**, 9456 (2018).
- [91] A. M. Childs and Y. Su, Nearly optimal lattice simulation by product formulas, *Phys. Rev. Lett.* **123**, 050503 (2019).
- [92] T. Peng, A. Harrow, M. Ozols, and X. Wu, Simulating large quantum circuits on a small quantum computer,

- Phys. Rev. Lett.* **125**, 150504 (2020).
- [93] A. M. Childs, R. Kothari, and R. D. Somma, Quantum algorithm for systems of linear equations with exponentially improved dependence on precision, *SIAM Journal on Computing* **46**, 1920 (2017).
- [94] Y. R. Sanders, D. W. Berry, P. C. S. Costa, L. W. Tessler, N. Wiebe, C. Gidney, H. Neven, and R. Babbush, Compilation of fault-tolerant quantum heuristics for combinatorial optimization, *PRX Quantum* **1**, 020312 (2020).
- [95] J. Wu and T. H. Hsieh, Variational thermal quantum simulation via thermofield double states, *Phys. Rev. Lett.* **123**, 220502 (2019).
- [96] G. Cornilescu, J. Marquardt, M. Ottiger, and A. Bax, Validation of protein structure from anisotropic carbonyl chemical shifts in a dilute liquid crystalline phase, *J. Am. Chem. Soc.* **120**, 6836 (1998).
- [97] G. Crippen, Note rapid calculation of coordinates from distance matrices, *J. Comp. Phys.* **26**, 449 (1978).
- [98] T. Gullion, Rotational-echo, double-resonance nmr, in *Modern Magnetic Resonance*, edited by G. A. Webb (Springer, 2008) pp. 713–718.
- [99] M. T. Colvin, R. Silvers, Q. Z. Ni, T. V. Can, I. Sergeyev, M. Rosay, K. J. Donovan, B. Michael, J. Wall, S. Linse, and R. G. Griffin, Atomic resolution structure of monomeric $\alpha\beta42$ amyloid fibrils, *J. Am. Chem. Soc.* **138**, 9663 (2016).
- [100] E. Vinogradov, P. Madhu, and S. Vega, High-resolution proton solid-state nmr spectroscopy by phase-modulated lee-goldburg experiment, *Chem. Phys. Lett.* **314**, 443 (1999).
- [101] A. Lesage, D. Sakellariou, S. Hediger, B. Eléna, P. Char-mont, S. Steuernagel, and L. Emsley, Experimental aspects of proton nmr spectroscopy in solids using phase-modulated homonuclear dipolar decoupling, *J. Mag. Reson.* **163**, 105 (2003).
- [102] B. Elena, G. de Pa epe, and L. Emsley, Direct spectral optimisation of proton–proton homonuclear dipolar decoupling in solid-state nmr, *Chem. Phys. Lett.* **398**, 532 (2004).
- [103] J. Choi, H. Zhou, H. S. Knowles, R. Landig, S. Choi, and M. D. Lukin, Robust dynamic hamiltonian engineering of many-body spin systems, *Phys. Rev. X* **10**, 031002 (2020).
- [104] J. Waugh, L. Huber, and U. Haeberlen, Approach to high-resolution nmr in solids, *Phys. Rev. Lett.* **20** (1968).
- [105] H. Geen and R. Freeman, Band-selective radiofrequency pulses, *J. Mag. Res.* **93**, 93 (1991).
- [106] G. A. Morris and R. Freeman, Selective excitation in fourier transform nuclear magnetic resonance, *J. Mag. Res.* **29**, 433 (1978).
- [107] G. Goldstein, M. Scheid, U. Hammerling, D. Schlesinger, H. Niall, and E. Boyse, Isolation of a polypeptide that has lymphocyte-differentiating properties and is probably represented universally in living cells, *Proc. Nat. Acad. Sci.* **72**, 11 (1975).
- [108] K. D. Wilkinson and D. Fushman, On the chain gang, *The Scientist*, 40 (2012).
- [109] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. Jarrod Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. Carey, Í. Polat, Y. Feng, E. W. Moore, J. Vand erPlas, D. Lax-alde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and S. . . Contributors, SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, *Nature Methods* **17**, 261 (2020).
- [110] N. C. Rubin, R. Babbush, and J. McClean, Application of fermionic marginal constraints to hybrid quantum algorithms, *New J. Phys.* **20**, 053020 (2018).

Appendix A: Optimizing query counts for NISQ algorithms

In this section we derive Eq. 24 by optimizing the allocation of oracle calls across a set of NISQ experiments. This is a slightly atypical approach; typically in a NISQ experiment one attempts to optimize the number of experiments (or shots) to achieve some target error ϵ [17, 110] rather than working with an oracular model. However, this optimization allows us to make a fair comparison between the results of Sec. IV A and Sec. IV B. Propagating variance through Eq. 5 yields

$$\text{Var} \left[\frac{dC[H]}{dh_n} \right] = \sum_{x,t} \frac{1}{\sigma_{x,t}^4} \left[\text{Var} [\bar{S}_x(t)] \bar{J}_x^n(t)^2 + \text{Var} [\bar{J}_x^n(t)] [\bar{S}_x(t) - S_x(t)]^2 \right]. \quad (\text{A1})$$

As our estimate of $\bar{J}_x^n(t)$ is bounded, assuming $\|O_x\| = 1$ the bound in Eq. 22 yields an estimator for $\bar{J}_x^n(t)$ with variance ϵ^2 using $M = t^2 \epsilon^{-2}$ repetitions of the circuit. We now assume that we may measure $S_x(t)$ and $J_x^n(t)$ for different x in parallel. This is realistic for our NMR application, and will save a factor N_x in the asymptotic scaling. As the number of oracle calls per circuit scales as t , we can achieve a variance $\text{Var} [\bar{J}_x^n(t)] = a_{x,t,J} t^3 C_{t,J}^{-1}$ using $C_{t,J}$ oracle calls (for some t and x -independent constant $a_{x,t,J}$). Estimating $\bar{S}_x(t)$ with to a variance ϵ^2 requires simply repeating the corresponding circuit (Fig. 2, top) ϵ^{-2} times, and so with $C_{t,S}$ oracle calls we achieve a variance $\text{Var} [\bar{S}_x(t)] = a_{x,t,S} t C_{t,S}^{-1}$ (for some t and x -independent constant $a_{x,t,S}$). As $[\bar{S}_x(t) - S_x(t)]$ is independent of t (being bounded by 2 when $\|O_x\| = 1$), and $\bar{J}_x^n(t)$ scales linearly in t (as discussed in the main text), assuming that $\sigma_{x,t} = \sigma$ we have

$$\text{Var} \left[\frac{dC[H]}{dh_n} \right] = \sum_{x,t} \frac{1}{\sigma^4} \left[\frac{a_{x,t,S} t^3}{C_{t,S}} + \frac{a_{x,t,J} t^3}{C_{t,S}} \right], \quad (\text{A2})$$

where we have absorbed the constants of the $\bar{S}_x(t)$ and $\bar{J}_x^n(t)$ scaling into $a_{x,t,S}$ and $a_{x,t,J}$. To optimize this, we adopt the same Lagrangian approach introduced for measurement optimization in Ref. [110]. We write a Lagrangian

$$\mathcal{L} = \lambda \left\{ \sum_{x,t} \frac{1}{\sigma^4} \left[\frac{a_{x,t,S} t^3}{C_{t,S}} + \frac{a_{x,t,J} t^3}{C_{t,J}} \right] - \text{Var} \left[\frac{dC[H]}{dh_n} \right] \right\} + \sum_{x,t} (C_{x,t,J} + C_{x,t,S}), \quad (\text{A3})$$

and then differentiate with respect to our free parameters $C_{x,t,J}$ and $C_{x,t,S}$ and solve for the result being equal to 0.

$$\frac{\partial \mathcal{L}}{\partial C_{t,S}} = 1 - \frac{\sum_x \lambda a_{x,t,S} t^3}{\sigma^4 C_{t,S}^2} = 0 \rightarrow C_{t,S} = \lambda^{\frac{1}{2}} [\text{sum}_x a_{x,t,S}]^{\frac{1}{2}} t^{\frac{3}{2}} \sigma^{-2} \quad (\text{A4})$$

$$\frac{\partial \mathcal{L}}{\partial C_{t,J}} = 1 - \frac{\sum_x \lambda a_{x,t,J} t^3}{\sigma^4 C_{t,J}^2} = 0 \rightarrow C_{t,J} = \lambda^{\frac{1}{2}} \left[\sum_x a_{x,t,J} \right]^{\frac{1}{2}} t^{\frac{3}{2}} \sigma^{-2}. \quad (\text{A5})$$

Substituting into the expression for a variance $\text{Var} \left[\frac{dC[H]}{dh_n} \right] = \epsilon^2$ yields

$$\epsilon^2 = \lambda^{-\frac{1}{2}} \sigma^{-2} \sum_t \left\{ \left[\sum_x a_{x,t,J} \right]^{\frac{1}{2}} t^{\frac{3}{2}} + \left[\sum_x a_{x,t,S} \right]^{\frac{1}{2}} t^{\frac{3}{2}} \right\}, \quad (\text{A6})$$

and rearranging for λ obtains

$$\lambda = \epsilon^{-4} \sigma^{-4} \left\{ \sum_t \left(\left[\sum_x a_{x,t,J} \right]^{\frac{1}{2}} t^{\frac{3}{2}} + \left[\sum_x a_{x,t,S} \right]^{\frac{1}{2}} t^{\frac{3}{2}} \right) \right\}^2. \quad (\text{A7})$$

Finally, we can write the total number of oracle calls as

$$C = \sum_{x,t} (C_{x,t,S} + C_{x,t,J}) = \epsilon^{-2} \sigma^{-4} \left\{ \sum_t \left(\left[\sum_x a_{x,t,J} \right]^{\frac{1}{2}} t^{\frac{3}{2}} + \left[\sum_x a_{x,t,S} \right]^{\frac{1}{2}} t^{\frac{3}{2}} \right) \right\}^2, \quad (\text{A8})$$

as $a_{x,t,J}$ and $a_{x,t,S}$ are constants, we have asymptotically that $\sum_x a_{x,t,J}, \sum_x a_{x,t,S} \sim \mathcal{O}(N_x)$. Substituting this into Eq. A8 yields immediately Eq. 24.

Appendix B: Fault-tolerant approach to the backwards problem

In this appendix, we expand on the circuitry and cost analysis of the fault-tolerant estimation of the gradient of the Hamiltonian learning problem in Sec. IV B.

We consider estimating the following quantity

$$\sum_{x=1}^{N_x} \sum_{j=1}^{N_d} \frac{1}{\sigma_{x,t_j}^2} \left(\text{Trace} (O_x \tilde{\rho}_x(t_j)) - \sum_{m=1}^{N_w} a_{x,m} \cos(t_j \omega_{x,m} + \phi_x) \right) \int_0^{t_j} ds \text{Trace} (O_x [V_n(t_j, s), \tilde{\rho}_x(t_j)]) \quad (\text{B1})$$

on a fault-tolerant quantum computer for the backwards problem. Here, operators V and $\tilde{\rho}_x$ are defined as

$$V_n(t, s) := e^{-i(t-s)H} V_n e^{i(t-s)H}, \quad \tilde{\rho}_x(t) := e^{-itH} |\psi_x\rangle\langle\psi_x| e^{itH} \quad (\text{B2})$$

for some Hamiltonian H , an arbitrary term V_n in H and x -dependent initial state $|\psi_x\rangle$.

Our approach depends on a quantum overlap estimation algorithm which we review in Appendix B 1. To implement this on a fault-tolerant quantum computer, we need to discretize the time integral and re-express the target quantity as a linear combination of quantum overlaps. We also need to truncate real parameters σ_{x,t_j} , t_j , $a_{x,m}$, and $\omega_{x,m}$ to a finite number of digits to construct the quantum circuit. We analyze the truncation and discretization error in Appendix B 2.

1. Quantum overlap estimation

To estimate (B1) on a fault-tolerant quantum computer, we first re-express it as a linear combination of quantum overlaps:

$$\begin{aligned} & \sum_{x,j} \frac{1}{\sigma_{x,t_j}^2} \langle \psi_x | e^{it_j H} O_x e^{-it_j H} | \psi_x \rangle \int_0^{t_j} ds \left(\langle \psi_x | e^{it_j H} O_x e^{-i(t_j-s)H} V_n e^{-isH} | \psi_x \rangle - \langle \psi_x | e^{isH} V_n e^{i(t_j-s)H} O_x e^{-it_j H} | \psi_x \rangle \right) \\ & + \sum_{x,j} \frac{1}{\sigma_{x,t_j}^2} \sum_m a_{x,m} \cos(t_j \omega_{x,m} + \phi_x) \int_0^{t_j} ds \left(\langle \psi_x | e^{it_j H} O_x e^{-i(t_j-s)H} V_n e^{-isH} | \psi_x \rangle - \langle \psi_x | e^{isH} V_n e^{i(t_j-s)H} O_x e^{-it_j H} | \psi_x \rangle \right). \end{aligned} \quad (\text{B3})$$

Then our goal is to find unitary operators SEL_a , quantum states $|\Psi^a\rangle = \text{PREP}_a |0\rangle$, and positive numbers $\lambda_a > 0$ for $a = 0, 1$, such that $\langle \Psi^a | \text{SEL}_a | \Psi^a \rangle$ compute the desired linear combinations up to scaled-down factors of λ_a ; equivalently, we say the target quantities are block-encoded by $|\Psi^a\rangle$ and SEL_a with scaled-down factor λ_a . We describe how to construct such SEL_a and $|\Psi^a\rangle$ in Sec. IV B.

We estimate the quantum overlap $\langle \Psi | \text{SEL} | \Psi \rangle$ using the overlap estimation algorithm of Ref. [84]. Specifically, we consider the two reflections

$$I - 2|\Psi\rangle\langle\Psi|, \quad I - 2\text{SEL}|\Psi\rangle\langle\Psi|\text{SEL}^\dagger. \quad (\text{B4})$$

These reflections keep the two-dimensional subspace $\text{span}\{|\Psi\rangle, \text{SEL}|\Psi\rangle\}$ invariant, on which their product

$$(I - 2|\Psi\rangle\langle\Psi|) (I - 2\text{SEL}|\Psi\rangle\langle\Psi|\text{SEL}^\dagger) \quad (\text{B5})$$

has eigenvalues $e^{\pm i 2 \arccos|\langle \Psi | \text{SEL} | \Psi \rangle|}$. Therefore, we can perform quantum phase estimation on the above operator and take the cosine of the outcome to estimate the amplitude $\alpha \sim |\langle \Psi | \text{SEL} | \Psi \rangle|$. To further retrieve the phase, we introduce an ancilla qubit and estimate

$$\begin{aligned} \beta_0 & \sim |\langle +\Psi | c-(\text{SEL}) | +\Psi \rangle| = \frac{|1 + \langle \Psi | \text{SEL} | \Psi \rangle|}{2}, \\ \beta_{\pi/2} & \sim |\langle +\Psi | (e^{i\frac{\pi}{4}Z} \otimes I) c-(\text{SEL}) | +\Psi \rangle| = \frac{|1 - i \langle \Psi | \text{SEL} | \Psi \rangle|}{2}, \end{aligned} \quad (\text{B6})$$

where $c-(\text{SEL})$ is the controlled operation

$$c-(\text{SEL}) := |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes \text{SEL}. \quad (\text{B7})$$

From these we obtain

$$y = \frac{4\beta_0^2 - \alpha^2 - 1}{2} + i \frac{4\beta_{\pi/2}^2 - \alpha^2 - 1}{2} \sim \langle \Psi | \text{SEL} | \Psi \rangle. \quad (\text{B8})$$

The standard quantum phase estimation outputs an estimate of the eigenphase with accuracy ϵ by making $\mathcal{O}(1/\epsilon)$ queries to the reflections, succeeding with a constant probability greater than $1/2$. The precision parameter ϵ directly translates to a maximal error of $\mathcal{O}(\epsilon)$ in the estimated overlap. To succeed with a higher probability, we can repeat quantum phase estimation and take the median of the outcomes. By Hoeffding's inequality, the success probability can be made arbitrarily close to one with only logarithmic overhead. We thus obtain:

Lemma 1 (Quantum overlap estimation). *Given quantum state $|\Psi\rangle = \text{PREP} |0\rangle$, unitary SEL , $\epsilon > 0$, and $0 < \delta < 1$, there exists a quantum algorithm with output y such that*

$$\mathbb{P}(|y - \langle \Psi | \text{SEL} | \Psi \rangle| \geq \epsilon) < \delta. \quad (\text{B9})$$

This algorithm makes $\mathcal{O}(\log(1/\delta)/\epsilon)$ queries to PREP and SEL (or their controlled version $c-(\text{PREP})$ and $c-(\text{SEL})$), and uses $\mathcal{O}(N \cdot \text{polylog}(1/\epsilon, 1/\delta))$ additional gates, where N is the number of qubits in the target system.

In the description of the above algorithm, we have ignored the normalization factor $\lambda > 0$ introduced by $|\Psi\rangle = \text{PREP} |0\rangle$ and SEL . To get the target quantity, we need to multiply the outcome of quantum overlap estimation by λ .

To ensure that the estimation succeeds with probability $1 - \delta$ and accuracy ϵ , it then suffices to make $\mathcal{O}(\log(1/\delta)\lambda/\epsilon)$ queries to PREP and SEL and use $\mathcal{O}(N \cdot \text{polylog}(\lambda, 1/\epsilon, 1/\delta))$ additional gates.

It is instructive to compare the fault-tolerant approach with the sampling-based approach that is more suitable to implement on near-term quantum devices. That approach uses the generalized Hadamard test which produces an unbiased estimate of the real and imaginary part of $\langle \Psi | \text{SEL} | \Psi \rangle$ with constant variance. By Hoeffding's inequality, it suffices to take $\mathcal{O}(\log(1/\delta)\lambda^2/\epsilon^2)$ samples to estimate with accuracy ϵ and probability $1 - \delta$. Each sample requires constant number of queries to PREP and SEL. Therefore, we get a factor of $\Theta(\lambda/\epsilon)$ saving by switching to the fault-tolerant quantum algorithm. See Sec. IV A and IV B for further discussions of these two approaches.

2. Truncation and discretization error

In this section, we analyze the error due to the truncation of real parameters and discretization of the time integral. We will see in Sec. IV B that this only introduces a logarithmic overhead in the overall cost.

We first consider discretizing the integral as

$$\int_0^{t_j} ds f(s) \approx \frac{t_j}{L} \sum_{\ell=0}^{L-1} f\left(\frac{\ell}{L}t_j\right), \quad (\text{B10})$$

where

$$f(s) := \text{Trace} \left(O_x \left[e^{-i(t_j-s)H} V_n e^{i(t_j-s)H}, e^{-it_j H} |\psi_x\rangle\langle\psi_x| e^{it_j H} \right] \right). \quad (\text{B11})$$

This discretization error can be made arbitrarily small by choosing L sufficiently large. Here, we analyze how the error scales as a function of L . Using the integral expansion

$$\begin{aligned} \int_0^{t_j} ds f(s) - \frac{t_j}{L} \sum_{\ell=0}^{L-1} f\left(\frac{\ell}{L}t_j\right) &= \sum_{\ell=0}^{L-1} \int_{\frac{\ell}{L}t_j}^{\frac{\ell+1}{L}t_j} ds \left(f(s) - f\left(\frac{\ell}{L}t_j\right) \right) \\ &= \sum_{\ell=0}^{L-1} \int_{\frac{\ell}{L}t_j}^{\frac{\ell+1}{L}t_j} ds \int_{\frac{\ell}{L}t_j}^s d\tau f'(\tau), \end{aligned} \quad (\text{B12})$$

we have

$$\left| \int_0^{t_j} ds f(s) - \frac{t_j}{L} \sum_{\ell=0}^{L-1} f\left(\frac{\ell}{L}t_j\right) \right| \leq \sum_{\ell=0}^{L-1} \int_{\frac{\ell}{L}t_j}^{\frac{\ell+1}{L}t_j} ds \int_{\frac{\ell}{L}t_j}^s d\tau \|f'\|_{\max} = \frac{t_j^2}{2L} \|f'\|_{\max}, \quad (\text{B13})$$

where

$$\|f'\|_{\max} := \max_{0 \leq s \leq t_j} |f'(s)| \quad (\text{B14})$$

is the largest derivative of f within the time interval $[0, t_j]$. The derivative $f'(s)$ takes the form

$$f'(s) := \text{Trace} \left(O_x \left[e^{-i(t_j-s)H} [iH, V_n] e^{i(t_j-s)H}, e^{-it_j H} |\psi_x\rangle\langle\psi_x| e^{it_j H} \right] \right), \quad (\text{B15})$$

which gives

$$\|f'\|_{\max} \leq 2 \|[H, V_n]\|. \quad (\text{B16})$$

In the following, we evaluate this bound for a model of clustered Hamiltonians acting on N sites:

$$H := \sum_{\mathcal{K}} H_{\mathcal{K}} + \sum_{\mathcal{K} \neq \mathcal{L}} H_{\mathcal{K}:\mathcal{L}} = \sum_{\mathcal{K}} \sum_{k, k' \in \mathcal{K}} H_{k, k'} + \sum_{\mathcal{K} \neq \mathcal{L}} \sum_{k \in \mathcal{K}, l \in \mathcal{L}} H_{k, l}. \quad (\text{B17})$$

Here, each term from the Hamiltonian acts on at most two sites and the sites are further grouped into clusters. We use calligraphic capital letters such as \mathcal{K} and \mathcal{L} to denote the clusters, and use k to denote an arbitrary single site within \mathcal{K} . Assuming that Hamiltonian terms are normalized $\|H_{k, l}\| \leq 1$, we have

$$\|f'\|_{\max} \leq 4\Lambda_{\text{ind}}, \quad \Lambda_{\text{ind}} := \max_{\mathcal{L}} \max_{l \in \mathcal{L}} \sum_{\mathcal{K}} \sum_{k \in \mathcal{K}} \|H_{k, l}\|. \quad (\text{B18})$$

Altogether, we have discretized (B1) with error at most

$$\sum_{x,j} \frac{1}{\sigma_{x,t_j}^2} \left| \text{Trace}(O_x \tilde{\rho}_x(t_j)) - \sum_m a_{x,m} \cos(t_j \omega_{x,m} + \phi_x) \right| \frac{2t_j^2 \Lambda_{\text{ind}}}{L} = \mathcal{O}\left(\frac{\lambda \Lambda_{\text{ind}} T}{L}\right), \quad (\text{B19})$$

where

$$\lambda := \lambda_0 + \lambda_1, \quad \lambda_0 := \sum_{x,j} \frac{t_j}{\sigma_{x,t_j}^2}, \quad \lambda_1 := \sum_{x,j,m} \frac{a_{x,m} t_j}{\sigma_{x,t_j}^2}, \quad T := \max_j t_j. \quad (\text{B20})$$

To achieve an accuracy of ϵ , it suffices to choose

$$L = \mathcal{O}\left(\frac{\lambda \Lambda_{\text{ind}} T}{\epsilon}\right). \quad (\text{B21})$$

We take T and L to be powers of two to simplify our circuit implementation.

We now consider the error due to the finite-digit truncation of the real parameters σ_{x,t_j} , t_j , $a_{x,m}$, and $\omega_{x,m}$. In general, the error in σ_{x,t_j} can be bounded under certain continuity assumptions with respect to the argument t_j . Here, we take $\sigma_{x,t_j} \equiv \sigma$ to be constant to simplify the analysis. In our circuit implementation, the evolution time will be loaded onto a quantum register using the QROM approach of Ref. [15] as

$$|t_{\log T-1} \cdots t_1 t_0\rangle |t_{-1} t_{-2} \cdots t_{-\log K}\rangle, \quad (\text{B22})$$

where we have again taken K to be a power of two to simplify the implementation. Here, T is the maximum possible time so $\log T$ bits suffice to represent the integer part of t . The length of the decimal part should be chosen large enough to represent the time sufficiently accurately. Specifically, for $|t' - t| \leq 1/K$, we have

$$\|e^{-it'H} - e^{-itH}\| \leq |t' - t| \|H\| = \mathcal{O}\left(\frac{N \Lambda_{\text{ind}}}{K}\right), \quad (\text{B23})$$

which implies

$$\begin{aligned} \left| \int_0^{t'} ds f(t', s) - \int_0^t ds f(t, s) \right| &\leq \left| \int_0^{t'} ds f(t', s) - \int_0^t ds f(t', s) \right| + \left| \int_0^t ds f(t', s) - \int_0^t ds f(t, s) \right| \\ &= \mathcal{O}\left(\frac{N \Lambda_{\text{ind}} t}{K}\right), \\ |g(t') - g(t)| &\leq |t' - t| \max_{\tau} |g'(\tau)| = \mathcal{O}\left(\frac{N \Lambda_{\text{ind}}}{K}\right) \end{aligned} \quad (\text{B24})$$

for

$$\begin{aligned} f(t, s) &:= \text{Trace}\left(O_x \left[e^{-i(t-s)H} V_n e^{i(t-s)H}, e^{-itH} |\psi_x\rangle\langle\psi_x| e^{itH} \right]\right), \\ g(t) &:= \text{Trace}\left(O_x e^{-itH} |\psi_x\rangle\langle\psi_x| e^{itH}\right). \end{aligned} \quad (\text{B25})$$

Similarly, if t'_j and $\omega'_{x,m}$ satisfy $|t'_j - t_j| \leq 1/K$ and $|\omega'_{x,m} - \omega_{x,m}| \leq 1/K$, then

$$|\cos(t'_j \omega'_{x,m} + \phi_x) - \cos(t_j \omega_{x,m} + \phi_j)| \leq |t'_j \omega'_{x,m} - t_j \omega_{x,m}| = \mathcal{O}\left(\frac{T+W}{K}\right), \quad (\text{B26})$$

where $T := \max_j t_j$ and $W := \max_{x,m} |\omega_{x,m}|$. The coefficients in the Hamiltonian can be approximately prepared using the coherent alias sampling approach also described in Ref. [15]. Using that approach with $\log K$ qubits for the inequality test, we have

$$\left| \frac{t'_j/\sigma'^2}{\lambda'_0} - \frac{t_j/\sigma^2}{\lambda_0} \right| \leq \frac{1}{K N_d}, \quad \left| \frac{a'_{x,m} t'_j/\sigma_j'^2}{\lambda'_1} - \frac{a_{x,m} t_j/\sigma_j^2}{\lambda_1} \right| \leq \frac{1}{K N_x N_\omega N_d}. \quad (\text{B27})$$

The total truncation error can now be bounded by

$$\mathcal{O}\left(\frac{\lambda N \Lambda_{\text{ind}}}{K} + \frac{\lambda_0 N \Lambda_{\text{ind}}}{K} + \frac{\lambda_1 (T+W)}{K} + \frac{N_d \lambda_0}{K N_d} + \frac{N_x N_\omega N_d \lambda_1}{K N_x N_\omega N_d}\right) = \mathcal{O}\left(\frac{\lambda (N \Lambda_{\text{ind}} + T + W)}{K}\right). \quad (\text{B28})$$

To ensure that this error is at most ϵ , it suffices to choose

$$K = \mathcal{O}(\text{poly}(\lambda, N, \Lambda_{\text{ind}}, T, W, 1/\epsilon)). \quad (\text{B29})$$