Optimal learning of quantum Hamiltonians from high-temperature Gibbs states

Jeongwan Haah[*]   Robin Kothari[†]   Ewin Tang[‡]

March 17, 2023

Abstract

We study the problem of learning a Hamiltonian H to precision ε, supposing we are given copies of its Gibbs state $\rho = \exp(-\beta H)/\operatorname{Tr}(\exp(-\beta H))$ at a known inverse temperature β. Anshu, Arunachalam, Kuwahara, and Soleimanifar [AAKS21] recently studied the sample complexity (number of copies of ρ needed) of this problem for geometrically local N-qubit Hamiltonians. In the high-temperature (low β) regime, their algorithm has sample complexity $\operatorname{poly}(N, 1/\beta, 1/\varepsilon)$ and can be implemented with polynomial, but suboptimal, time complexity.

In this paper, we study the same question for a more general class of Hamiltonians. We show how to learn the coefficients of a Hamiltonian to error ε with sample complexity $S = O(\log N/(\beta\varepsilon)^2)$ and time complexity linear in the sample size, $O(SN)$. Furthermore, we prove a matching lower bound showing that our algorithm's sample complexity is optimal, and hence our time complexity is also optimal.

In the appendix, we show that virtually the same algorithm can be used to learn H from a real-time evolution unitary e^{-itH} in a small t regime with similar sample and time complexity.

Contents

[*]Microsoft Quantum and Microsoft Research, Redmond, WA, USA.

[†]Microsoft Quantum and Microsoft Research, Redmond, WA, USA.

[‡]University of Washington, Seattle, WA, USA. Some of this work was performed while E.T. was a research intern at Microsoft Quantum.

1 Introduction

In this paper we study a problem that is at the intersection of quantum many-body physics and machine learning: learning the Hamiltonian of a quantum system from copies of its Gibbs state. This problem has recently received much attention in the quantum community [BAL19, QR19, EHF19, BGP+20], and the classical analogue of this task, learning undirected graphical models or Markov random fields (MRFs), is well studied in the machine learning community [KS01, AKN06, SW12, BMS13, Bre15, VMLC16, KM17].

Motivation. This problem has a straightforward physical motivation. The Hamiltonian H of a quantum system is an operator that tells us how the constituents of the system interact with each other and how the system evolves in time, which is governed by the Schrödinger equation. The Hamiltonian also tells us what the equilibrium state of the quantum system will be if it is in contact with the environment at a particular temperature and reaches thermal equilibrium. This state, which is a function of the temperature and the Hamiltonian, is called the Gibbs state. Formally, for a Hamiltonian H and inverse temperature β (i.e., temperature $1/\beta$), the Gibbs state is $\rho = \exp(-\beta H)/\operatorname{Tr}(\exp(-\beta H))$.

    In the Hamiltonian learning problem, we imagine that we have a system governed by an unknown Hamiltonian H from a known class of physically reasonable Hamiltonians, such as geometrically local Hamiltonians, and we have access to copies of its Gibbs state at a known inverse temperature β. These copies, for example, result from leaving the system to interact with the environment at a known temperature and stabilize: eventually, the system is described by the Gibbs state. Our goal is to learn the Hamiltonian H, from the assumed class of Hamiltonians, while minimizing the number of copies of ρ required and the running time of the algorithm. These are called the sample complexity and time complexity of the algorithm.

    The classical analogue of Hamiltonian learning is the problem of learning undirected graphical models or Markov random fields. This is, in fact, a special case of Hamiltonian learning where everything is classical, which means that the Hamiltonian is a diagonal operator, and consequently the Gibbs state is a diagonal density operator, which is just a sample from a classical probability distribution. The goal is again to learn the parameters of the classical Hamiltonian from these

samples. This classical problem has been studied for over 50 years, usually in the harder setting of learning the terms and parameters of the classical Hamiltonian, starting with the work of [CL68], to more recent works that provide nearly-sample-optimal algorithms, with time-efficient implementations [SW12, BMS13, Bre15, VMLC16, KM17]. Markov random fields find applications in a variety of areas including computer graphics, vision, economics, sociology, and biology [KS80, Cli90, Lau96, JEMF06, KF09, Li09], so studying its quantum generalization is well-motivated independently from its physical motivation.

Problem statement.   The formal statement of the Hamiltonian learning problem is as follows. Consider a quantum system of N qubits and a Hamiltonian $H = \sum_{a=1}^{M} \lambda_a E_a \in \mathbb{C}^{2^N \times 2^N}$ consisting of M terms, where the operators $E_a \in \mathbb{C}^{2^N \times 2^N}$ are known, distinct, non-identity Pauli operators[1] and the coefficients satisfy $\lambda_a \in [-1, 1]$ for all $a \in [M] = \{1, \ldots, M\}$. We assume the Hamiltonian has no identity term since the Gibbs state is invariant under adding multiples of the identity matrix to the Hamiltonian. We assume terms are distinct because identical terms can be merged. Further suppose that this Hamiltonian H is low-intersection (a constraint defined below).

Given copies of the Gibbs state of this unknown low-intersection Hamiltonian H and known inverse temperature β, our goal is to learn the coefficients λ_a to additive error ε, or equivalently, to learn the vector of coefficients to error ε in ℓ_∞ norm. Previous work on the problem has also considered the goal of learning this vector to ℓ_2 norm, so we study this version of the problem as well.

We define the class of *low-intersection Hamiltonians* to be the set of Hamiltonians where each operator E_a is supported on a constant number of qubits, meaning that it acts as the identity operator on all but a constant number of qubits (which are its support), and for each operator E_a, there are only a constant number of other operators E_b such that E_a and E_b's supports have nontrivial intersection.

Notice that this definition has no geometric constraints. Instead, it generalizes geometrically local Hamiltonians in fixed-dimensional Euclidean spaces, which is the class of physically motivated Hamiltonians with geometric constraints considered in prior work [AAKS21]. In such Hamiltonians, each operator E_a is only supported on a constant number of qubits that are adjacent in the underlying geometry (e.g., a 2-dimensional grid). Since the dimension is fixed and interactions must be local, each operator E_a can only act nontrivially on a constant number of qubits, and furthermore each qubit can only be nontrivially involved in a constant number of operators E_a. So, a geometrically local Hamiltonian in any constant-dimensional space is always low-intersection.

The converse is not true, though. For example, if we arrange qubits on the vertices of a constant-degree expander graph, and let edges denote 2-qubit interaction terms, such a Hamiltonian would be a low-intersection Hamiltonian, but not a geometrically local Hamiltonian in any constant-dimensional Euclidean space. In this introduction we will assume that we have a low-intersection Hamiltonian whose degree is a constant independent of other parameters, although our general algorithm can also handle growing degree.

Prior work.   We first discuss the complexity of the classical problem to understand the best we could do, since classical Hamiltonians, also known as Markov random fields, are a special case of quantum Hamiltonians. The classical problem is then the parameter learning of low-intersection MRFs to ℓ_∞ error ε. The sample complexity and time complexity of this problem are

$$\frac{2^{O(\beta)} \log N}{\beta^2 \varepsilon^2} \quad \text{and} \quad \frac{2^{O(\beta)} N \log N}{\beta^2 \varepsilon^2}, \tag{1}$$

---

[1]It is not essential that these are Pauli operators, but it is convenient for us that the entries of Pauli matrices are small integers, which allows us compute quantities of interest exactly and not worry about numerical precision.

respectively. The sample complexity is optimal up to the constant in the exponent [SW12], so the time complexity, which is the time needed to read all of the samples, is also optimal. This result appears to be folklore, so in Appendix B, we give a simple algorithm demonstrating this result.

Most of the classical literature focuses on the harder task of *structure learning*, which is learning the terms of the Hamiltonian in addition to the coefficients, for Ising models, which are classical Hamiltonians with only pairwise interactions. For structure learning in the Ising model, the same sample complexity bound can be achieved in time only polynomially worse than the sample complexity times the size of each sample [VMLC16, KM17].

Notice that the problem becomes harder as $\beta \to 0$ and as $\beta \to \infty$. This is intuitive because the state at $\beta = 0$ is the maximally mixed state (or the uniform distribution in the classical case), which contains no information about the Hamiltonian. When β tends to ∞, the state tends to the ground state of the Hamiltonian, which does not have enough information to reconstruct the entire Hamiltonian.

The quantum version of this problem for geometrically local Hamiltonians was recently studied. The algorithm in [BAL19] allows us to learn Hamiltonians from stationary states of Hamiltonian dynamics (which include Gibbs states) or from the dynamics itself by measuring local observables and solving a system of linear equations; however, it was unclear how the algorithm would perform in the worst-case. More recently, Anshu, Arunachalam, Kuwahara, and Soleimanifar [AAKS21] was the first to rigorously establish sample complexity upper bounds for this problem in the full range of parameters, and in particular, for all inverse temperatures β. They showed that a geometrically local Hamiltonian in a constant-dimensional space can be learned to ℓ_∞ error ε using

$$O\left(\frac{2^{\mathrm{poly}(\beta)}N^2 \log N}{\beta^c \varepsilon^2}\right) \tag{2}$$

samples[2], for some constant $c > 4$. Note that for a geometrically local Hamiltonian, the number of terms $N = \Theta(M)$, so we have expressed the bound in terms of N.

This upper bound is worse than the classical sample complexity in Eq. (1) in several regards. First, it has worse dependence on β both in the numerator and the denominator, which means it is worse in the high-temperature and low-temperature regime. Second, the dependence on N is quadratic, whereas the dependence on N is logarithmic in the classical upper bound[3]. This leaves two natural open questions: Can we solve the quantum problem with sample complexity matching Eq. (1)? And what about time complexity?

The question of time complexity is not explicitly addressed in [AAKS21], but they note that the problem can be solved in polynomial time in the high-temperature regime, by combining their algorithm with the polynomial-time algorithm for computing partition functions at high temperatures due to [KKB20]. We discuss this approach further in the "Comparison with previous quantum algorithms" section, but in brief, this approach leaves significant room for improvement in both sample complexity and time complexity.

Our results. We study the Hamiltonian learning problem in the high-temperature regime, and we are able to obtain an algorithm with optimal sample complexity and optimal time complexity. The high-temperature regime is where we know β is smaller than some fixed constant called the critical inverse temperature, β_c. This constant β_c depends only on the constant in the definition of a low-intersection Hamiltonian, and not on N or M.

---

[2]Actually, they claim a slightly weaker statement: learning to ℓ_2 error ε using N times the expression in Eq. (2) many samples. We derive the version stated here in Remark 4.5.

[3]In the ℓ_2 error setting, though, this classical upper bound has a factor of N, and so the quantum bound is polynomially close to the classical bound in the high-temperature setting.

Our main algorithmic result is the following. A more precise version can be found in Section 4.

Theorem 1.1 (Algorithm). *Let H be a low-intersection Hamiltonian on N qubits, $\varepsilon > 0$, and $\beta < \beta_c$. Then we can learn the coefficients of H with ℓ_∞ error ε and failure probability δ using $O\left(\frac{1}{\beta^2\varepsilon^2}\log\frac{N}{\delta}\right)$ samples. Consequently, we can learn the coefficients of H with ℓ_2 error ε and failure probability δ using $O\left(\frac{N}{\beta^2\varepsilon^2}\log\frac{N}{\delta}\right)$ samples. In both cases, the time complexity is linear in the sample size, which is the sample complexity multiplied by N, the size of each sample.*

Our upper bound improves on the sample complexity of [AAKS21] and indeed matches the sample complexity of the classical algorithm in the high-temperature regime, where the $2^{O(\beta)}$ term can be dropped since it is constant. Furthermore, our algorithm has optimal time complexity.

Along the way, we show that the log-partition function is $(\frac{\beta^2}{2})$-strongly convex in the high-temperature regime; this is the main quantity bounded by [AAKS21] to achieve their sample complexity result, and our analysis improves this strong convexity parameter to within a constant factor of its true value. As observed in [AAKS21], this strong convexity bound implies a lower bound on the variance of macroscopic observables in thermal equilibrium. Specifically, for a local operator $\sum_{a=1}^M v_a E_a$, its variance with respect to the Gibbs state is $v^\dagger(\nabla^{\otimes 2}\mathcal{L})v = \Omega(\beta^2\|v\|_2^2)$, where $\nabla^{\otimes 2}\mathcal{L}$ is the Hessian of the log-partition function, improving on the bound $\Omega(\beta^c\|v\|_2^2/N)$ implied by [AAKS21].

We also prove a matching lower bound on the sample complexity showing that our algorithm's sample complexity cannot be improved. Our lower bound significantly improves on the lower bound shown in [AAKS21] (displayed below in Eq. (3)), holds for the full range of β, and matches our algorithm's complexity in the high-temperature regime. A more formal version appears as Theorem 5.3 and Theorem 5.5.

Theorem 1.2 (Lower bound). *For any $\varepsilon \in (0, 1/2]$, any $\beta > 0$, and any N, there exists a 2-local Hamiltonian on N qubits such that the sample complexity of learning its coefficients to ℓ_∞ error ε and failure probability δ is $\Omega\left(\frac{\exp(\beta)}{\beta^2\varepsilon^2}\log\frac{N}{\delta}\right)$, and the sample complexity of learning its coefficients to ℓ_2 error ε and constant failure probability is $\Omega\left(\frac{\exp(\beta)N}{\beta^2\varepsilon^2}\right)$.*

The Hamiltonians used in our lower bound are extremely simple 2-local Hamiltonians, where each term acts nontrivially only on 2 qubits and each qubit is involved in only 1 term. This shows that although our algorithms apply to a more general class of Hamiltonians than considered by [AAKS21], restricting our attention to a simpler class of Hamiltonians will not allow us to improve on the sample complexity compared to our algorithm.

This improves significantly on the lower bound given in [AAKS21], which states that any algorithm that learns a Hamiltonian to ℓ_2 error ε has sample complexity

$$\Omega\left(\frac{\sqrt{N}+\log(1-\delta)}{\beta\varepsilon}\right). \tag{3}$$

In addition, we observe that virtually the same algorithm can be used to learn a low-intersection Hamiltonian H, given black-box access to its real-time evolution unitary e^{-itH}, provided t is known and smaller than some critical time that is a constant in the definition of a low-intersection Hamiltonian, which does not depend on N or M.

Theorem 1.3 (Real-time dynamics). *Let H be a low-intersection Hamiltonian on N qubits and let $U = e^{-itH}$ be a blackbox unitary with $t < t_c$. Then we can learn the coefficients of H to ℓ_∞ error ε with success probability $1 - \delta$, using U $O\left(\frac{1}{t\varepsilon^2}\log\frac{N}{\delta}\right)$ times, with time complexity $O\left(\frac{N}{t\varepsilon^2}\log\frac{N}{\delta}\right)$.*

We report this observation in Appendix A. Prior work on this task [BAL19, ZYLB21] uses measurements of short-time evolutions, i.e. time resolution $t = O(\varepsilon)$, to estimate time derivatives, which gives a sample complexity scaling as $1/\varepsilon^4$. We improve this quadratically, and since we only apply U for t as large as constant, we improve the time resolution to constant.

High-level overview of techniques. Our algorithm in Theorem 1.1 proceeds in two steps. First, we notice that, for sufficiently small (but constant) β, the Taylor series expansion of the expectation $\text{Tr}(E_a \rho)$ in β converges. This observation follows from the cluster expansion techniques from [KS20], which essentially describes and bounds the coefficients in the Taylor series expansion of the log-partition function, $\log \text{Tr} \exp(-\beta H)$. We reproduce these (lengthy but elementary) calculations here, amending some minor issues in their presentation. The log-partition function is related to the expectation $\text{Tr}(E_a \rho)$ (in fact, $\frac{\partial}{\partial \lambda_a} \log \text{Tr} \exp(-\beta H) = -\beta \text{Tr}(E_a \rho)$), so we can use these results on convergence of the log-partition function to get convergence of the expectation.

One notable difference from prior work is our results showing how to efficiently compute the Taylor series expansion described above (Proposition 3.13). Prior work asserted such computation was possible [KKB20], but did not provide an explicit algorithm. We provide an algorithm (Algorithm 2), and because the E_a's are Pauli operators in our setting, the matrices are composed of small integers and so this algorithm works in exact arithmetic.

This shows that we can approximate $\text{Tr}(E_a \rho)$, an expression that can be easily estimated from copies of ρ, by a polynomial in $\{\lambda_b\}$, the parameters we wish to estimate. This polynomial is the one we get from truncating the Taylor series expansion of $\text{Tr}(E_a \rho)$. Once we approximate these $\text{Tr}(E_a \rho)$'s, we are left with the task of solving the system of polynomial equations defined by these truncated Taylor series expansions. By bounding the $\infty \to \infty$ norm of the inverse Jacobian of this system, we immediately get a bound on the sample complexity (Theorem 4.2). By performing the Newton–Raphson method for root-finding, we can invert this system efficiently, only needing to compute the (first-order) Jacobian for $O(\log \frac{1}{\beta \varepsilon})$ iterations. In fact, the Newton–Raphson method performs so efficiently that its runtime is dominated by the runtime of simply reading in the input, making the algorithm as a whole run in linear time (Theorem 4.6).

For our lower bound, we use information-theoretic techniques to show that without sufficiently many Gibbs states, the coefficient vector cannot be determined to ε error. In particular, we use Fano's lemma to establish a lower bound from a KL-divergence computation, similarly to prior classical work for lower bounds of learning undirected graphical models [SW12]. This immediately gives the lower bound in the ℓ_∞ case (Theorem 5.3), and a simple argument with error correcting codes bootstraps this to an ℓ_2 bound (Theorem 5.5).

Comparison with previous quantum algorithms. We now provide a comparison of our algorithm's sample and time complexity bounds compared to that of prior work. We consider the task of learning a geometrically local Hamiltonian for sufficiently small β with success probability 0.9. One algorithm to compare to is a naive "state tomography" strategy that one can derive from [KKB20, Theorems 2 and 11]. These results show that to estimate a Hamiltonian coefficient λ_a to ε error, for sufficiently high temperature, it suffices to know the Gibbs state ρ on a "patch", a ball of radius $O(\log \frac{1}{\beta \varepsilon})$ around the support of E_a. So, one can perform state tomography to learn the patch of ρ in time exponential in the number of qubits in the ball, giving an algorithm for Hamiltonian learning with quasi-polynomial sample and time complexity

$$\tilde{O}\Big(e^{\log^c(\frac{1}{\beta \varepsilon})} \log(N)\Big) \text{ and } \tilde{O}\Big(e^{\log^c(\frac{1}{\beta \varepsilon})} N \log(N)\Big), \tag{4}$$

where the factor of c comes from the ambient dimension, i.e. a ball of radius r on the lattice of qubits is size $O(r)^c$.

As for [AAKS21], we already gave the sample complexity bound in Eq. (2). As for time complexity, [AAKS21] describe an stochastic gradient descent (SGD) approach that solves the task of Hamiltonian learning, assuming one can evaluate the log partition function. This subroutine is hard in general, but [AAKS21] acknowledge that, for the high-temperature regime, it can be done efficiently [KKB20] to get a time-efficient algorithm for Hamiltonian learning. This resulting algorithm will inherit the sub-optimal sample complexity of [AAKS21], and though [AAKS21] don't give a precise time complexity, we can conclude that at best its time complexity will be linear in the sample size, which is the sample complexity times N. So, the sample and time complexity of [AAKS21] in the high temperature regime is something like

$$O\Big(\frac{N^2 \log(N)}{\beta^c \varepsilon^2}\Big) \text{ and } O\Big(\frac{N^3 \log(N)}{\beta^c \varepsilon^2}\Big). \tag{5}$$

The prior work achieves optimal scaling in either N or $1/\varepsilon$. We get a sample and time complexity,

$$O\Big(\frac{\log(N)}{\beta^2 \varepsilon^2}\Big) \text{ and } O\Big(\frac{N \log(N)}{\beta^2 \varepsilon^2}\Big), \tag{6}$$

that is simultaneously optimal in all parameters and improves at least polynomially over prior results. In fact, in certain regimes (like when $\beta \varepsilon \approx \exp(-\log^{1/c}(N))$), we give a super-polynomial improvement in sample complexity over prior work.

Our algorithm can be viewed as a refinement of the patching argument described in [KKB20]; cluster expansion is the technique used to prove the results there, and we use it in a similar way to argue that it suffices to only consider $O(\frac{1}{\beta \varepsilon})$ terms that are within $\log(\frac{1}{\beta \varepsilon})$ of the support of E_a. Our contribution is that we use this expansion algorithmically via the Newton–Raphson method to improve the quasi-polynomial time from the naive algorithm to polynomial time.

[AAKS21] proceeds by establishing that the log partition function is strongly convex. This actually immediately gives a sample complexity bound, but [AAKS21] instead provide a concrete algorithm, stochastic gradient descent (SGD), that solves the task of Hamiltonian learning, assuming one can evaluate the log partition function (a hard problem in general). Specifically, this means that they need to lower bound the smallest eigenvalue of the Hessian of the log partition function, or equivalently, upper bound the spectral norm of the inverse of this matrix.

At its core, our strategy is similar to that of [AAKS21]. Like in [AAKS21], we also work with the inverse of the Hessian of the log-partition function (or an approximation of it). Since we want to solve the problem with ℓ_∞ error ε, instead of upper bounding the spectral norm of this matrix, we upper bound its $\infty \to \infty$ norm. Our bound also yields an upper bound on the spectral norm that is tighter than the bound in [AAKS21]. The improvement is due to the more precise characterization of this matrix via the series expansion described above. To get a time-efficient algorithm, we then use the Newton–Raphson method, whose analysis also requires us to understand a higher order derivative of the log partition function than is needed for bounding the sample complexity. Our characterization through the series expansion is able to provide this higher order information, which allows us to bound the running time of the Newton–Raphson method and show it to be time efficient.

Comparision with previous classical algorithms. One might wonder why classical techniques for solving the Hamiltonian learning problem do not apply to quantum Hamiltonians. Our algorithm and the [AAKS21] algorithm do not use strategies that are common in the classical literature.

The reason is that classical algorithms for learning Hamiltonians rely on a property of the classical Gibbs state called the Markov property. To understand this property, partition the set of bits into 3 disjoint parts A, B, and C, such that there is no term in the Hamiltonian that has a bit

from A and C. The sets A and C only interact through B. Now it is not hard to show that if we condition the classical Gibbs distribution of this Hamiltonian on the values taken by bits in B, the resulting distribution on A and C is independent. In fact, the Hammersley–Clifford theorem shows that this is not just a property of Gibbs states, but this property characterizes Gibbs states [HC71]. This property fails to hold in general for quantum Gibbs states, although it can hold for special classes of Hamiltonians, such as commuting Hamiltonians [BP12]. For high-temperature Gibbs states, this holds only approximately, as cluster expansion formalizes: roughly, in this setting, A and C can be treated as independent subsystems provided B is "wide" enough.

Classical algorithms can efficiently perform structure learning by treating it as parameter learning on the full space of k-local Pauli matrices. This would naively take exponential time, but algorithms are still able to use the low-intersection guarantee, despite not knowing anything else about the terms [KM17]. It's not clear how to show a similar statement in the quantum setting; we can apply our algorithm, but it only works for β smaller than $1/\operatorname{poly}(N)$.

2 Preliminaries

Throughout the paper all the exponential and logarithm functions $(\exp, \log)$ are with natural base $e = \sum_{k=0}^{\infty} \frac{1}{k!} \approx 2.718$. For a vector v, $\|v\| = \sum_i |v_i|^2$ denotes the Euclidean (or ℓ_2) norm. For a matrix M, we use $\|M\| = \max_{v \neq 0} \frac{\|Mv\|}{\|v\|}$ to denote the operator norm (also known as the spectral norm, $2 \to 2$ norm, or the Schatten ∞-norm).

2.1 Notations and conventions

Definition 2.1 (Hamiltonian). A *Hamiltonian* is a collection of tuples (a, E_a, λ_a), where a is an index ranging over some finite set of M elements, which we usually take to be $[M] = \{1, 2, \ldots, M\}$; the *Hamiltonian term* $E_a \in \mathbb{C}^{\mathsf{D} \times \mathsf{D}}$ is a Hermitian operator with $\|E_a\| \leq 1$ acting on a Hilbert space of dimension D; and the *Hamiltonian term coefficient* $\lambda_a \in [-1, 1]$ is a real number. We use the notation $\lambda = (\lambda_1, \ldots, \lambda_M)$ for the vector of coefficients. The associated Hamiltonian operator H is defined to be $H = \sum_a \lambda_a E_a$.

Our full algorithm will require that E_a are distinct, non-identity Pauli matrices. This assumption is neither essential nor too constraining. Since a Hamiltonian is Hermitian, we definitely want E_a to be Hermitian. Requiring that $\operatorname{Tr} E_a = 0$ is simply a shift in the eigenspectrum of the Hamiltonian. Since Pauli operators form an orthonormal basis for operators, we can always write any Hamiltonian term as a sum of Pauli operators. The dual interaction graph degree may increase in this rewriting by a factor that is at most the exponential of the number of qubits in the support of the Hamitonian terms; however, in the arguably most important scenario where a term acts on a constant number of qubits, this blowup is a constant multiplicative factor.

We always assume that some system of N qubits comprises the Hilbert space, so $\mathsf{D} = 2^N$ is some power of two. Upon introducing this decomposition of the Hilbert space into qubits, we can define the support $\operatorname{Supp}(P)$ of an operator P. The support is the minimal set of qubits such that P can be written as $P = O_{\operatorname{Supp}(P)} \otimes I_{\operatorname{Supp}(P)^c}$ for some operator O. (The superscript c here means the complement.)

Definition 2.2 (Dual interaction graph). For any Hamiltonian $\{(a, E_a, \lambda_a) : a \in [M]\}$, there is an associated undirected *dual interaction graph* $\mathfrak{G}$ with vertex set $[M]$ and an edge between a and b if and only if $a \neq b$ and

$$\operatorname{Supp}(E_a) \cap \operatorname{Supp}(E_b) \neq \varnothing. \tag{7}$$

We denote by $\mathfrak{d}$ the maximum degree of the graph $\mathfrak{G}$ over all vertices.

Note that we have defined a Hamiltonian in such a way that it is possible that $E_a = E_b$ for $a \neq b$. If this is the case, then there is an edge in $\mathfrak{G}$ between a and b. In our learning algorithm we will require that E_a's are distinct nonidentity Pauli operators, but this definition is sufficient for our series expansion of log-partition functions.

Although we do not specify how $\mathfrak{d}$ depends on M, we focus on the case when $\mathfrak{d}$ is a constant independent of M. This case encompasses most Hamiltonians classes discussed in the literature. For example, if every Hamiltonian term E_a acts on a constant number of qubits and every qubit is involved in a constant number of terms, all with respect to M, then $\mathfrak{d}$ will be constant as well. More concretely, if we have a directed graph G with a qubit on each vertex and a two-qubit Hermitian operator for every edge (which may require direction on each edge), then the vertices of $\mathfrak{G}$ correspond to the edges of G and the dual interaction graph has $\mathfrak{d} \leq 2(d-1)$, where d is the degree (in-degree plus out-degree) of G.

As another example, an important class of Hamiltonians is the class of geometrically local Hamiltonians on, say, Euclidean space $\mathbb{R}^d$. There are some constant number of qubits on each point of the lattice $\mathbb{Z}^d \subset \mathbb{R}^d$, and a Hermitian operator is defined for each unit hypercube. Here, the dual interaction graph has $\mathfrak{d} \leq 3^d - 1$, which is again independent of M.

Definition 2.3 (Gibbs state)**.** The *Gibbs state* of the Hamiltonian $\{(a, E_a, \lambda_a)\}$ at inverse temperature $\beta > 0$ is given by

$$\frac{\exp(-\beta H)}{\operatorname{Tr} \exp(-\beta H)} = \exp\left(-\beta \sum_a \lambda_a E_a\right) \Big/ \operatorname{Tr} \exp\left(-\beta \sum_a \lambda_a E_a\right). \tag{8}$$

It is a trivial but important fact that the exponential of the Hamiltonian operator always makes sense for any $\beta \in \mathbb{C}$, not just positive β, since the norm of the Hamiltonian is upper bounded by M.

2.2 Time complexity

When we discuss the time complexity of our algorithms, we will usually do so in the standard *word RAM model*, where operations on *words*, integers of w bits with $w \geq \log_2(N + M + \varepsilon_{\text{machine}}^{-1})$, take unit time. This word size is defined so that an index into qubits, an index into terms, and β can all be stored in one word. With this model, the input to the Hamiltonian learning problem (the terms $\{E_a\}_a$ and β) can be given in $O(LM)$ words, where L is the maximum support of all the terms E_a. This requires representing a term E_a by its support ($\leq L$ words) and the non-identity Pauli operator that E_a performs on each qubit in its support ($\leq 2L$ bits).

Remark 2.4. Our algorithms assume that the input to the Hamiltonian learning problem also contains an adjacency-list representation of the dual graph $\mathfrak{G}$ corresponding to the input Hamiltonian. That is, we want to query any node $b \in [M]$ to receive a list of its neighbors in $\mathfrak{G}$ in unit time, where the list is given as a random-access array.

Producing this adjacency-list representation requires only $O(LM\mathfrak{d} \log \mathfrak{d})$ time: first, for each qubit i, produce a list of the terms that have that qubit in its support, $S_i = \{a \in [M] : i \in \operatorname{Supp}(a)\}$; second, sort the S_i's; third, for every term a, produce a list that is the sorted concatenation of every qubit in its support, $\cup_{i \in \operatorname{Supp}(a)} S_i$. After removing a itself, this list is the set of neighbors of a in $\mathfrak{G}$.

The first step takes time linear in the number of edges, so $O(LM)$ time. The second step takes $O(N\mathfrak{d} \log \mathfrak{d})$ time, since $|S_i| \leq \mathfrak{d} + 1$. The third step takes $O(L\mathfrak{d}M)$ time, since we can merge sorted lists in linear time, removing duplicates as we find them so that we never merge lists of length larger than $\mathfrak{d} + 1$. So, the total time complexity is $O(N\mathfrak{d} \log \mathfrak{d} + LM\mathfrak{d}) = O(LM\mathfrak{d} \log \mathfrak{d})$, since $N \leq LM$.

As for operations on the quantum computer, we will use the standard model of time complexity (or, rather, gate complexity). However, since the only quantum operations our algorithm will ever perform is measuring a single-qubit Pauli operator on a qubit of an input Gibbs state, it suffices to just assume that this operation takes unit time.

Finally, in this paper we will ignore issues of numerical stability. We can do this comfortably because the only portion of our algorithm that is not exact arithmetic is the Newton's method iterations in Algorithm 3. This algorithm accounts for per-iteration error already, so issues with numerical instability do not arise here.

2.3 Analytic functions and series expansions

We will extensively use infinite series expansions of analytic functions, and here we discuss general principles of handling infinite series. The material here is all standard in complex analysis.

A complex function $f : D \to \mathbb{C}$ is defined to be (complex) *analytic* at $a \in D \subseteq \mathbb{C}$ if, for some $\varepsilon > 0$, the function agrees with a power series

$$f(x) = \sum_{k=0}^{\infty} c_k (x - a)^k \tag{9}$$

for all $x \in D$ such that $|x - a| < \varepsilon$. Here, the coefficients $c_k \in \mathbb{C}$ and $\varepsilon \in \mathbb{R}_{>0}$ may depend on a. Since such a power series converges uniformly (at least on a small neighborhood of a), the infinite sum commutes with taking derivatives, and therefore the coefficients c_k must be those of the Taylor expansion:

$$c_k = \frac{(\partial_x^k f)(a)}{k!}. \tag{10}$$

A complex function is said to be analytic on an open set D if it is analytic at every point in D. A basic theorem in complex analysis is that a complex function is complex differentiable (holomorphic) at a point if and only if it is analytic at that point [Rud87, 10.14, 10.16].

Functions with a power series expansion are "rigid" in the following sense. A power series expansion at $a \in D$ is identically zero if and only if there is an infinite sequence of distinct points $x_1, x_2, \ldots \in D$ such that $\lim_{n \to \infty} x_n = a$ and that the power series is zero at every x_n. This implies a uniqueness theorem of analytic functions [Rud87, 10.18]: if two analytic functions f, g on a connected open domain D agree on a subset that has a limit point within D, then $f = g$ on D. In particular,

Lemma 2.5. *Suppose a complex function f is complex differentiable on E, an open neighborhood of the origin. If its Taylor series at the origin*

$$\sum_{k=0}^{\infty} \frac{f^{(k)}(0)}{k!} z^k \tag{11}$$

uniformly converges on $D = \{z \in \mathbb{C} : |z| < r\}$ for some r, then it must converge to $f(z)$ for all z on the connected component of the origin in $D \cap E$.

This basically says that the Taylor series of a holomorphic function can only converge to its function value. Our typical use of the lemma will be as follows. We will show that f is complex differentiable on an open set E containing $\mathbb{R}$ and consider its Taylor expansion at the origin and lower bound the radius of convergence. Then, the lemma will imply that the series converges to the function value in the real domain of convergence of the Taylor series. The way we show a function is complex differentiable is simply by noting that it is a composition of functions (such as addition, multiplication, exponentiation, etc.) that are themselves complex differentiable.

Proof. Let g be the function on D defined by the series. By uniform convergence, the series is differentiable term by term, and hence g is complex differentiable everywhere in D, and hence is analytic on D. Since f is analytic on E, on some tiny open neighborhood $U \subset D \cap E$ of the origin it is represented by its Taylor series, which is the same as g. Since U has a limit point within $D \cap E$, we must have $f = g$ on the connected component of $D \cap E$ containing the origin. $\qquad\square$

Lemma 2.5 is not true if we only assume that f is *real* infinitely differentiable, as witnessed by the well-known function

$$h(x) = \begin{cases} \exp(-\frac{1}{x^2}) & (x \neq 0) \\ 0 & (x = 0) \end{cases}. \tag{12}$$

Its Taylor series at the origin is identically zero, and hence converges everywhere on $\mathbb{R}$ but the function is zero only at the origin. Note that if we extend the domain of definition of h from $\mathbb{R}$ to $\mathbb{C}$ by the same formula, then h is not complex differentiable at the origin; h is divergent at the origin along the imaginary axis.

3   Series expansions of expectation values

The main goal of this section is to prove the following theorem. A direct consequence of this theorem is that the Taylor series expansion for expectation values of local operators converges when $2e^2(\mathfrak{d}+1)^2\beta < 1$ (see the discussion above Eq. (98)). A reader who wishes to understand just our learning algorithm may skip the rest of this section on their first reading, since all notions and properties needed for our algorithm and its analysis are contained in the statement of Theorem 3.1.

Theorem 3.1. *Consider a Hamiltonian $\{(a, E_a, \lambda_a) : a \in [M]\}$. Then, for every $a \in [M]$ we have a Taylor series expansion*

$$\frac{\mathrm{Tr}(E_a \exp(-\beta H))}{\mathrm{Tr}\exp(-\beta H)} = \frac{\mathrm{Tr}(E_a)}{\mathsf{D}} + \sum_{m=1}^{\infty} \beta^m p_m(\lambda_1, \ldots, \lambda_M), \tag{13}$$

where equality holds whenever the series converges absolutely. For any $m \in \mathbb{Z}_{>0}$, the following hold:

1. *$p_m \in \mathbb{R}[\lambda_1, \ldots, \lambda_M]$ is a degree m homogeneous polynomial in the Hamiltonian term coefficients.*

2. *p_m involves λ_b only if the distance between a and b on $\mathfrak{G}$, $\mathrm{dist}_{\mathfrak{G}}(a, b)$, is at most m.*

3. *p_m consists of at most $e\mathfrak{d}(1 + e(\mathfrak{d} - 1))^m$ monomials.*

4. *The coefficient in front of any monomial of p_m is at most $(2e(\mathfrak{d}+1))^{m+1}(m+1)$ in magnitude.*

Suppose further that every E_a is a tensor product of Pauli matrices, supported on at most L qubits. Then, after $O(LM\mathfrak{d}\log\mathfrak{d})$ pre-processing time (see Remark 2.4), the following are true for every $m \in \mathbb{Z}_{>0}$.

A. *The list of monomials that appear in p_m can be enumerated in time $O(m\mathfrak{d}C)$, where C is the number of monomials (so, in particular, in time $O(m\mathfrak{d}^2(1 + e(\mathfrak{d} - 1))^m)$).*

B. *The coefficient of any monomial in p_m can be computed exactly as a rational number in $O(Lm^3 + 8^m m^5 \log^2 m) = (8^m + L)\,\mathrm{poly}(m)$ time.*

11

Overview of the proof. The series expansion in Eq. (13) is certainly conceivable. When $\beta = 0$, the numerator is zero because $\operatorname{Tr} P = 0$. The first order term in β comes from the first order term in β of the numerator and the zeroth order term in the denominator. One can keep finding terms order by order in β, but this calculation soon becomes too complicated to be useful in a proof.

Following the wisdom of statistical mechanics, as done explicitly in e.g. [KKB20, WA22], we examine the logarithmic partition function $\log \operatorname{Tr} \exp(-\beta H)$, and take differentials to arrive at Eq. (13). This basic connection is established in Section 3.1.

To put our series expansion on a rigorous foundation, we will use complex differentiability (holomorphicity) of the function $\beta \mapsto \log \operatorname{Tr} \exp(-\beta H)$. Although we are only interested in the regime where β is positive real, it is quite useful for us to observe that this function behaves nicely in a sufficiently large domain in the complex plane of β. Armed with Lemma 2.5, we fearlessly write infinite Taylor series and resummations thereof, to derive a series expansion of the logarithmic partition function as a multivariate function of $\lambda_1, \ldots, \lambda_M$. This leads to the concept of cluster expansion

$$\sum_{m \geq 0} \sum_{\mathbf{V}} C(m, \mathbf{V}), \tag{14}$$

which is an infinite sum of finite sums over clusters $\mathbf{V}$ (see Section 3.2). In this section we use the convention that boldface uppercase letters refer to clusters. Each $C(m, \mathbf{V})$ has a β^m factor, so in order for the infinite sum $\sum_{m \geq 0}$ to converge for some small enough β, each $\sum_{\mathbf{V}} C(m, \mathbf{V})$ has to be at most exponentially large in m. This exponential bound will occupy us for most of the proof, regarding which we follow many elements from [KS20] and [WA22]. We count the number of summands of $\sum_{\mathbf{V}}$ purely combinatorially in Section 3.3, and, separately, bound the magnitude of $C(m, \mathbf{V})$ for each $\mathbf{V}$. The second step follows the approach in [WA22]. The result will be that $|\sum_{\mathbf{V}} C(m, \mathbf{V})| \leq \operatorname{poly}(\beta, \mathfrak{d}) O(\mathfrak{d}^2 \beta)^m$. In contrast, assertions in [KS20] imply $|\sum_{\mathbf{V}} C(m, \mathbf{V})| \leq \operatorname{poly}(\beta, \mathfrak{d}) O(\mathfrak{d} \beta)^m$, which is quadratically stronger in $\mathfrak{d}$. This difference is because we do not use [KS20, App. D of arXiv-v2], in which the argument appears to have a mathematical gap in [KS20, (D.10) of arXiv-v2]. Recently, another mathematical gap in [KS20, App. C of arXiv-v2] was pointed out by [WA22].

Finally, we take some care to give an algorithm to compute the Taylor series. This is often elided as it is fairly standard, but for completeness, we show how to do this with symbolic computation to get the coefficients with no error. First, we are able to enumerate the list of summands in the series in Section 3.4. This is a breadth-first search, with some care to avoid duplicating work. Second, we compute the coefficient for each summand in Section 3.6. These coefficients are derivatives of the log-partition function $\log \operatorname{Tr} \exp(-\beta H)$ at the origin. With some simple observations (truncating Taylor series and using the definition of a derivative), we conclude that these derivatives are equal to the trace of a matrix polynomial in the terms of the Hamiltonian at zero Eq. (72). Since we assume that these are tensor products of Pauli matrices, these polynomials can be evaluated efficiently, where the final runtime is exponential in the order of the derivative, as one would expect.

This gives the formal guarantees that this computation is bounded in complexity. Practically, one can use any method at hand, not necessarily relying on the specific algorithm we propose. Indeed, there is a large body of classic literature on high temperature expansions (see e.g. the book [Dom96]) which does not always discuss formal convergence.

3.1 The logarithmic partition function

Given a Hamiltonian $\{(a, E_a, \lambda_a) : a \in [M]\}$, our primary object of study is

$$\mathcal{L} = \log \operatorname{Tr} \exp(-\beta H) = \log \operatorname{Tr} \exp\Big(-\beta \sum_{a \in [M]} \lambda_a E_a\Big). \tag{15}$$

The argument of the log function is called the *partition function* in statistical mechanics. Hence, we will refer to this expression $\mathcal{L}$ as the *logarithmic partition function* or *log-partition function* of the Hamiltonian. The quantity $\beta^{-1}\mathcal{L}$ is called (Helmholtz) free energy in statistical mechanics, but we will not use this terminology.

The connection of the logarithmic partition function to Theorem 3.1 is given by the following.

Proposition 3.2. *For any Hamiltonian $\{(a, E_a, \lambda_a) : a \in [M]\}$, $a \in [M]$, and nonzero $\beta \in \mathbb{C}$,*

$$\frac{\operatorname{Tr}(E_a \exp(-\beta H))}{\operatorname{Tr} \exp(-\beta H)} = -\frac{1}{\beta} \frac{\partial}{\partial \lambda_a} \log \operatorname{Tr} \exp(-\beta H). \tag{16}$$

Proof. Since H and E_a have finite norm, the Taylor expansion of $\operatorname{Tr} \exp$ converges absolutely. The claim is proved by

$$\begin{aligned}
\frac{\partial}{\partial \lambda_a} \operatorname{Tr} \exp(-\beta H) &= \sum_{m=0}^{\infty} \frac{1}{m!} \operatorname{Tr}\Big[\frac{\partial}{\partial \lambda_a}(-\beta H)^m\Big] \\
&= \sum_{m=1}^{\infty} \frac{1}{m!} \sum_{k=1}^{m} \operatorname{Tr}[(-\beta H)^{k-1}(-\beta E_a)(-\beta H)^{m-k}] \\
&= -\beta \sum_{m=1}^{\infty} \frac{1}{m!} \sum_{k=1}^{m} \operatorname{Tr}[E_a(-\beta H)^{m-1}] \\
&= -\beta \operatorname{Tr}[E_a \exp(-\beta H)],
\end{aligned} \tag{17}$$

using linearity and the cyclic property of Tr. Finally,

$$\frac{\partial}{\partial \lambda_a} \log \operatorname{Tr} \exp(-\beta H) = \frac{1}{\operatorname{Tr} \exp(-\beta H)} \frac{\partial}{\partial \lambda_a} \operatorname{Tr} \exp(-\beta H) = -\beta \frac{\operatorname{Tr}(E_a \exp(-\beta H))}{\operatorname{Tr} \exp(-\beta H)}, \tag{18}$$

which completes the proof. $\qquad\square$

If we understood the series expansion of the logarithmic partition function well enough, we could prove Theorem 3.1 easily by way of Proposition 3.2. There will be an important advantage (Proposition 3.5 below) in considering the logarithmic partition function, rather than the ratio of two traces as in Theorem 3.1. So, we will study the series expansion of the logarithmic partition function.

3.2 Deriving multivariate Taylor series expansions

In this section, we prove that a series expansion for the logarithmic partition function like the one in Eq. (13) converges in some open neighborhood around the origin.

Though the logarithmic partition function is a complex-valued function of $\beta, \lambda_1, \ldots, \lambda_M$, depending on context, we'll think of it as either a function of a single variable β for a fixed choice of λ_a's, or a function of $(\lambda_1, \ldots, \lambda_M)$ for a fixed β.

Let us first take the first perspective: Fix[4] $\lambda_a \in (-1, 1)$ for all $a \in [M]$ and consider the map $\beta \mapsto \mathcal{L}$. By our convention that Hamiltonian terms E_a are Hermitian with $\|E_a\| \leq 1$, the spectrum of the Hamiltonian operator $H = \sum_a \lambda_a E_a$ is contained in the real interval $(-M, M)$, and further, for any $\beta \in \mathbb{C}$, the spectrum of the operator $-\beta H$ is contained in the complex disk $\{z \mid z \in \mathbb{C}, |z| \leq |\beta| M\}$. Hence, for $\beta \in E$ where

$$E = \left\{ x + iy \;\middle|\; x, y \in \mathbb{R}, \, |y| < \frac{\pi}{2M} \right\}, \tag{19}$$

if $h \in (-M, M)$ is an eigenvalue of H, then the complex number $e^{-\beta h} = e^{-xh} e^{-iyh}$ has positive real part. Therefore, $\operatorname{Tr} \exp(-\beta H)$ is in the right half-plane of the complex plane and the function $\beta \mapsto \log \operatorname{Tr} \exp(-\beta H)$ is complex differentiable on E, using that $\log$ is complex differentiable on the right half-plane and that complex differentiability is closed under composition. So by Lemma 2.5, we are guaranteed that the function $\beta \mapsto \log \operatorname{Tr} \exp(-\beta H)$ has a Taylor series representation in some open neighborhood of the origin in the complex plane of β, although we do not yet know how large the open neighborhood can be.

Note that the same argument shows that the multivariate function

$$\mathbb{C}^{M+1} \ni (\beta, \lambda_1, \ldots, \lambda_M) \mapsto \log \operatorname{Tr} \exp(-\beta H) \tag{20}$$

is complex differentiable in each variable on

$$\tilde{E} = \left\{ (\beta, \lambda_1, \ldots, \lambda_M) \in \mathbb{C}^{M+1} \;\middle|\; |\operatorname{Im}(\beta \lambda_a)| < \frac{\pi}{2M} \right\}. \tag{21}$$

Here, for any $z \in \mathbb{C}$, $\operatorname{Im}(z)$ denotes the imaginary part of z. This set $\tilde{E}$ is an open neighborhood of the real line of β times the real box $(-1, 1)^M$ of λ_a's.

The Taylor expansion at the origin is straightforward to write as $\sum_{m \geq 0} \frac{1}{m!} \beta^m (\partial_\beta^m \mathcal{L}|_{\beta=0})$, but this is not enlightening. Let us make some observations first. We interpret the logarithmic partition function as a function of $z = (z_1, \ldots, z_M) \in \mathbb{C}^M$:

$$\mathcal{L} = \log \operatorname{Tr} \exp\left(-\sum_a z_a E_a\right) \text{ where } z_a = \beta \lambda_a. \tag{22}$$

It follows that

$$\frac{\partial \mathcal{L}}{\partial \beta} = \sum_a \frac{\partial z_a}{\partial \beta} \frac{\partial \mathcal{L}}{\partial z_a} = \sum_a \lambda_a \frac{\partial \mathcal{L}}{\partial z_a}, \tag{23}$$

so

$$\begin{aligned}
\mathcal{L} &= \sum_{m \geq 0} \frac{\beta^m}{m!} \left(\frac{\partial^m \mathcal{L}}{\partial \beta^m} \bigg|_{\beta=0} \right) \\
&= \sum_{m \geq 0} \frac{\beta^m}{m!} \sum_{a_1, a_2, \ldots, a_m} \lambda_{a_1} \cdots \lambda_{a_m} \left(\frac{\partial^m \mathcal{L}}{\partial z_{a_1} \cdots \partial z_{a_m}} \bigg|_{z=(0,\ldots,0)} \right).
\end{aligned} \tag{24}$$

Since $\mathcal{L}$ is complex differentiable in any variable (at least on $\tilde{E}$), it is infinitely differentiable, and hence any two differentiations commute. So, instead of summing over ordered tuples $(a_1, \ldots, a_m)$, we can sum over *multisets*, which are unordered tuples. This particular class of multisets will be used frequently, so let's give a proper definition.

---

[4]We restrict λ_j to the open interval $(-1, 1)$ to avoid the inconvenience of discussing derivatives at the boundary of the domain of $\mathcal{L}$.

Definition 3.3. A *cluster* $\mathbf{V}$ is a set of tuples $\{(a, \mu(a)) \mid a \in [M]\}$ where the function $\mu : [M] \to \mathbb{Z}_{\geq 0}$ maps a to the multiplicity of a. The total weight, denoted $|\mathbf{V}|$, of $\mathbf{V}$ is $\sum_a \mu(a)$. We will write $a \in \mathbf{V}$ if $\mu(a)$ is nonzero, and the *support* of $\mathbf{V}$ is defined to be $\mathrm{Supp}\,\mathbf{V} = \{a \in [M] : \mu(a) \geq 1\}$. We also introduce a combinatorial factor $\mathbf{V}!$ to mean $\prod_a \mu(a)!$.

One may think of a cluster as a function $a \mapsto \mu(a)$ or a monomial in indeterminates $\lambda_1, \ldots, \lambda_M$. Returning to Eq. (24), we have

$$
\begin{aligned}
\mathcal{L} &= \sum_{m \geq 0} \frac{\beta^m}{m!} \sum_{a_1, a_2, \ldots, a_m} \lambda_{a_1} \cdots \lambda_{a_m} \left(\frac{\partial^m \mathcal{L}}{\partial z_{a_1} \cdots \partial z_{a_m}} \Big|_{z=(0,\ldots,0)} \right) \\
&= \sum_{m \geq 0} \beta^m \sum_{\mathbf{V}:|\mathbf{V}|=m} \frac{1}{\mathbf{V}!} \prod_{a \in \mathrm{Supp}\,\mathbf{V}} \lambda_a^{\mu(a)} \left(\prod_{a \in \mathrm{Supp}\,\mathbf{V}} \frac{\partial^{\mu(a)}}{\partial z_a^{\mu(a)}} \right) \mathcal{L} \Big|_{z=(0,\ldots,0)} \\
&= \sum_{m \geq 0} \sum_{\mathbf{V}:|\mathbf{V}|=m} \frac{1}{\mathbf{V}!} \underbrace{\prod_{a \in \mathrm{Supp}\,\mathbf{V}} \lambda_a^{\mu(a)}}_{\lambda^{\mathbf{V}}} \underbrace{\left(\prod_{a \in \mathrm{Supp}\,\mathbf{V}} \frac{\partial^{\mu(a)}}{\partial \lambda_a^{\mu(a)}} \right) \Big|_{\lambda=(0,\ldots,0)}}_{\mathcal{D}_{\mathbf{V}}} \mathcal{L} \\
&= \sum_{m \geq 0} \sum_{\mathbf{V}:|\mathbf{V}|=m} \frac{\lambda^{\mathbf{V}}}{\mathbf{V}!} \mathcal{D}_{\mathbf{V}} \mathcal{L}.
\end{aligned}
\tag{25}
$$

Note that we have introduced the cluster notations

$$
\lambda^{\mathbf{V}} = \prod_{a \in \mathrm{Supp}\,\mathbf{V}} \lambda_a^{\mu(a)} \text{ and } \mathcal{D}_{\mathbf{V}} = \prod_{a \in \mathrm{Supp}\,\mathbf{V}} \frac{\partial^{\mu(a)}}{\partial \lambda_a^{\mu(a)}} \Big|_{\lambda_1 = \cdots = \lambda_M = 0}.
\tag{26}
$$

Eq. (25) is the series expansion of $\mathcal{L}$ that we are going to investigate. It is nothing but the Taylor expansion of $\mathcal{L}$, treating it as a multivariate function $(\lambda_1, \ldots, \lambda_M) \mapsto \mathcal{L}$. Though we could have guessed this expansion from the outset, this derivation is necessary to show that the series converges: we start with a β-series whose validity is guaranteed, albeit on an unspecified small domain, by the complex differentiability with respect to β.

The number of all clusters of weight m is at least $\binom{M}{m}$. This is much larger than what we claim in Item 3.1(3), where the bounds are independent of M. The special structure of the logarithmic partition function will help us show the improved bound.

3.3 Counting connected clusters

The main point of considering the logarithmic partition function is that a cluster has nonzero coefficient in Eq. (25) only if it is *connected*.

Definition 3.4. A cluster $\mathbf{W} = \{(a, \mu(a))\}$ is *connected* if the subgraph of $\mathfrak{G}$ induced by the support of $\mathbf{W}$ is connected.

Proposition 3.5. *Define $\mathcal{Z} = \frac{1}{\mathsf{D}} \mathrm{Tr} \exp(-\beta H)$. If $\mathbf{W}'$ and $\mathbf{W}''$ are both nonempty such that no edge of $\mathfrak{G}$ connects $\mathrm{Supp}\,\mathbf{W}'$ and $\mathrm{Supp}\,\mathbf{W}''$, then $\mathcal{D}_{\mathbf{W}' \cup \mathbf{W}''}\mathcal{Z} = (\mathcal{D}_{\mathbf{W}'}\mathcal{Z})(\mathcal{D}_{\mathbf{W}''}\mathcal{Z})$. In particular, if a cluster $\mathbf{W}$ is not connected, then $\mathcal{D}_{\mathbf{W}}\mathcal{L} = 0$.*

Proof. The two operators $\sum_{a \in \mathrm{Supp}\,\mathbf{W}'} \lambda_a E_a$ and $\sum_{b \in \mathrm{Supp}\,\mathbf{W}''} \lambda_b E_b$ commute with each other since the operators' supports do not overlap. Define $\mathcal{Z}|_{\mathbf{W}} = \frac{1}{\mathsf{D}} \mathrm{Tr} \exp(-\beta \sum_{a \in \mathrm{Supp}\,\mathbf{W}} \lambda_a E_a)$ and similarly $\mathcal{L}|_{\mathbf{W}} = \log(\mathsf{D}\mathcal{Z}|_{\mathbf{W}})$. Then, $\mathcal{Z}|_{\mathbf{W}} = \mathcal{Z}|_{\mathbf{W}'}\mathcal{Z}|_{\mathbf{W}''}$ and $\mathcal{L}|_{\mathbf{W}} = \mathcal{L}|_{\mathbf{W}'} + \mathcal{L}|_{\mathbf{W}''} - \log \mathsf{D}$. The first claim

immediately follows. Since $\mathcal{D}_{\mathbf{W}}$ evaluates the derivative at the origin of the Hamiltonian coefficient space, we see

$$\mathcal{D}_{\mathbf{W}}\mathcal{L} = \mathcal{D}_{\mathbf{W}}(\mathcal{L}|_{\mathbf{w}}) = \prod_{a \in \text{Supp }\mathbf{W}'} \frac{\partial^{\mu(a)}}{\partial\lambda_a^{\mu(a)}} \prod_{b \in \text{Supp }\mathbf{W}''} \frac{\partial^{\mu(b)}}{\partial\lambda_b^{\mu(b)}} (\mathcal{L}|_{\mathbf{w}'} + \mathcal{L}|_{\mathbf{w}''} - \log \mathsf{D})\Big|_{\lambda=(0,\dots,0)} = 0. \quad (27)$$

In other words, $\mathcal{L}$ decomposes into a sum of $\mathcal{L}|_{\mathbf{w}'}$ and $\mathcal{L}|_{\mathbf{w}''}$, which are each a function of a strict subset of the variables λ_a that appear in W, but they are each being differentiated with respect to all λ_a that appear in W, and hence must evaluate to zero. $\square$

Now we bound the number of all connected clusters of a given total weight w. What matters most for us is that this bound is just exponential in the total weight, $\exp(O(w))$ instead of, say, the more naive bound $O(w!)$. We optimize the base of the exponent in our bound, since this affects the eventual algorithm's runtime.

Proposition 3.6. *Let $\mathfrak{G}$ be any graph with maximum degree $\mathfrak{d} \geq 2$. Given any node a of $\mathfrak{G}$ and any weight $w \in \mathbb{Z}_{>0}$, the number of all connected clusters $\mathbf{W}$ such that $a \in \mathbf{W}$ and $|\mathbf{W}| = w$ is at most $e\mathfrak{d}(1 + e(\mathfrak{d} - 1))^{w-1}$. If $\mathfrak{d} = 1$, then we have an upper bound of w.*

Proof. For a fixed degree $\mathfrak{d}$, the number of clusters is maximized when $\mathfrak{G}$ is an infinite $\mathfrak{d}$-regular tree. Since this graph is self-similar, without loss of generality we can think of $a \in \mathfrak{G}$ as being the root of the tree. So, this question reduces to upper-bounding the number of connected rooted subtrees of the infinite $\mathfrak{d}$-regular tree, where nodes of the subtrees are allowed to have multiplicity. If every node in the tree must have multiplicity one (that is, if we disallow multiplicity), we have the following.

Lemma 3.7. *For $n \in \mathbb{Z}_{\geq 0}$, let D_n be the number of all connected rooted subtrees with n nodes in the infinite $\mathfrak{d}$-regular tree. Then*

$$D_n = \binom{n(\mathfrak{d}-1)+1}{n-1} \frac{\mathfrak{d}}{n(\mathfrak{d}-1)+1} \leq e\mathfrak{d}(e(\mathfrak{d}-1))^{n-1}. \quad (28)$$

To count subtrees *with* multiplicity, we must count the number of ways to assign a positive integer to every node of a subtree. If the subtree has k nodes, there are $\binom{(w-k)+(k-1)}{k-1}$ ways to assign multiplicities to these nodes such that the multiplicities sum to w. Hence, the number of weight-w connected rooted clusters of the infinite $\mathfrak{d}$-regular tree is

$$\sum_{k=1}^{w} D_k \binom{w-1}{k-1} \leq \sum_{k=1}^{w} e\mathfrak{d}(e(\mathfrak{d}-1))^{k-1} \binom{w-1}{k-1} \quad (29)$$

$$= e\mathfrak{d}(1 + e(\mathfrak{d}-1))^{w-1}. \quad \square$$

Proof of Lemma 3.7. This can be done with standard manipulations of generating functions, which we detail below. As a reminder, if we have a sequence $\{a_0, a_1, a_2, \dots\}$ of integers, then the generating function corresponding to it is $A(z) = \sum_{i \geq 0} a_i z^i$. We use the notation $[z^i]A(z)$ to refer to the coefficient a_i of $A(z)$.

Let $E(z)$ be the generating function counting subtrees of the infinite $(\mathfrak{d} - 1)$-ary tree, the tree where every node has $\mathfrak{d} - 1$ children.[5] Then the following recursion holds.

$$[z^n](E(z)) = \sum_{\substack{(i_1,\dots,i_{\mathfrak{d}-1}) \in \mathbb{Z}_{\geq 0}^{\mathfrak{d}-1} \\ i_1 + \cdots + i_{\mathfrak{d}-1} = n-1}} \prod_{j=1}^{\mathfrak{d}-1} [z^{i_j}](E(z)) \quad (30)$$

---

[5]For combinatorialists, this is also the generating function for the Fuss-Catalan numbers: $E(z) + 1 = \sum_{n \geq 0} \frac{1}{n(\mathfrak{d}-1)+1} \binom{n(\mathfrak{d}-1)+1}{n} z^n$ [GKP94, 7.5 Example 5].

In words, this describes an n-node rooted subtree of the infinite $(\mathfrak{d}-1)$-ary tree as $\mathfrak{d}-1$ (possibly empty) subtrees corresponding to each child of the root, where the number of nodes in each subtree sum to $n-1$. These subtrees are also rooted subtrees of the infinite $(\mathfrak{d}-1)$-ary tree, allowing the expression above to recurse as stated. One can verify that Eq. (30) is equivalent to the equation

$$E(z) = z(1 + E(z))^{\mathfrak{d}-1}. \tag{31}$$

Let $D(z)$ be the generating function counting subtrees of the infinite $\mathfrak{d}$-regular tree, the tree where every node has $\mathfrak{d}$ neighbors. In particular, it only differs from the $(\mathfrak{d}-1)$-ary tree in the root node, where there are $\mathfrak{d}$ instead of $\mathfrak{d}-1$ many options. Using a similar argument as with $E(z)$, we can conclude that

$$D(z) = z(1 + E(z))^{\mathfrak{d}} = E(z)(1 + E(z)). \tag{32}$$

As an aside, the number of rooted clusters on the $\mathfrak{d}$-regular infinite tree, $C_n = \sum_{k \geq 1} \binom{n-1}{k-1} D_k$, corresponds to the generating function definition $C(z) = D(\frac{z}{1-z})$. Since $\binom{n-1}{k-1} = [z^n]((z + z^2 + \cdots)^k)$, we have

$$[z^n](C(z)) = \sum_{k \geq 0} \binom{n-1}{k-1}[z^k]D(z) = [z^n]D(z + z^2 + \cdots) = [z^n]D(\tfrac{z}{1-z}). \tag{33}$$

Returning to the proof, we use the Lagrange–Bürmann formula to get the series expansion of $D(z)$ from the inverse of $E(z)$. In particular, we use the formulation common in combinatorics [FS09, Thm A.2 (14)],

$$[z^n]H(y(z)) = \frac{1}{n}[u^{n-1}](H'(u)\phi(u)^n), \tag{34}$$

where H is an arbitrary function and $y(z) = z\phi(y(z))$. We set $H(z) = z(1 + z)$, $y(z) = E(z)$, and $\phi(z) = (1 + z)^{\mathfrak{d}-1}$. So

$$[z^n]D(z) = \frac{1}{n}[u^{n-1}]((2u + 1)(1 + u)^{n(\mathfrak{d}-1)}) \tag{35}$$

$$= \frac{1}{n}\left(2\binom{n(\mathfrak{d}-1)}{n-2} + \binom{n(\mathfrak{d}-1)}{n-1}\right)$$

$$= \binom{n(\mathfrak{d}-1)+1}{n-1}\frac{\mathfrak{d}}{n(\mathfrak{d}-1)+1}.$$

Eq. (35) is the desired equality in the lemma statement. As for the inequality, clearly, $[z^n]D(z) \leq e\mathfrak{d}(e(\mathfrak{d}-1))^{n-1}$ for $n = 0, 1$. For $n \geq 2$, we see

$$\binom{n(\mathfrak{d}-1)+1}{n-1}\frac{\mathfrak{d}}{n(\mathfrak{d}-1)+1} \leq \mathfrak{d}\binom{n(\mathfrak{d}-1)}{n-1} \tag{36}$$

$$\leq \mathfrak{d}\left(\frac{en(\mathfrak{d}-1)}{n-1}\right)^{n-1}$$

$$= \mathfrak{d}(e(\mathfrak{d}-1))^{n-1}\left(\frac{n}{n-1}\right)^{n-1}$$

$$\leq e\mathfrak{d}(e(\mathfrak{d}-1))^{n-1}. \qquad \square$$

3.4 Enumerating connected clusters

Enumerating clusters of total weight m at a node $a \in \mathfrak{G}$ can be done with the following algorithm. First, perform a breadth-first search starting at a to produce the disjoint sets V_i for $i \in \{0, 1, \ldots, m-1\}$, where V_i is the set of nodes exactly i away from a in graph distance. Note that this induces a directed tree $G = (V, E)$ with vertices $V = V_0 \sqcup \cdots \sqcup V_{m-1}$ and a directed edge (u, v) occurring if it is an edge in $\mathfrak{G}$ and v is at a lower level than u (so $u \in V_i$ and $v \in V_{i+1}$). Since G is directed, all the neighbors of $u \in V_i$ are in V_{i+1}. For $S \subset V$, we denote $\Gamma(S)$ to be the neighborhood of S in G. If $\mathbf{S}$ is a multiset, $\Gamma(\mathbf{S})$ is defined to be the G-neighborhood of the support, $\Gamma(\mathrm{Supp}\,\mathbf{S})$.

Every cluster can be represented uniquely as a collection of multisets $\mathbf{S}_i$ of V_i for $i \in \{0, 1, \ldots, m-1\}$ satisfying that every node in $\mathbf{S}_{i+1}$ has a parent in $\mathbf{S}_i$ (or equivalently, satisfying that $\mathrm{Supp}\,\mathbf{S}_0 \cup \cdots \cup \mathrm{Supp}\,\mathbf{S}_{m-1}$ is connected in $\mathfrak{G}$).

Because of this characterization, we can enumerate clusters through a recursive function that, given the first i layers of a cluster $\mathbf{S}_0, \ldots, \mathbf{S}_{i-1}$, outputs a list of all possible ways to complete the cluster $\mathbf{S}_i, \ldots, \mathbf{S}_{m-1}$ such that the total weight of the cluster is m. We describe this function in Algorithm 1; it only requires three parameters, the recursion level i, the remaining weight m_i, and the neighborhood C_i, which correspond to i, $m - (|\mathbf{S}_0| + \cdots + |\mathbf{S}_{i-1}|)$, and $\Gamma(\mathbf{S}_{i-1})$ in the above description. To find all the clusters of weight m, run $\mathrm{tails}(0, m, \{a\})$. The function proceeds as follows: at recursion level $i \geq 0$, we loop over all possible nonempty multisets $\mathbf{S}_i$ of $C_i \subset V_i$ of weight $\leq m_i$. For each such multiset $\mathbf{S}_i$ of C_i, we call the recursive function to enumerate all of its possible continuations, with parameters $i+1$, $m_{i+1} = m_i - |\mathbf{S}_i|$, and $C_{i+1} = \Gamma(\mathbf{S}_i)$. Once it returns the possible continuations of this cluster, we add $\mathbf{S}_i$ to every continuation (to make them continuations of $\mathbf{S}_{i-1}$). Upon enumerating all possible $\mathbf{S}_i$'s, return the resulting (now complete) list of continuations of $\mathbf{S}_{i-1}$ as output.

Algorithm 1: $\mathrm{tails}(i, m_i, C_i)$ [Cluster enumeration recursion]

Data: depth i, size m_i, base $C_i \subseteq V_i$
Result: $\mathcal{C}_i$, the collection of all connected multisets of size m_i supported on
$\quad\quad C_i \sqcup V_{i+1} \sqcup \cdots \sqcup V_{m-1}$

1 **if** $m_i = 0$ **then**
2 $\quad$ Return $\{\varnothing\}$;
3 **end**
4 Let out $\leftarrow \varnothing$;
5 **for** $\mathbf{S} \in \left(\!\!\binom{C_i}{1}\!\!\right) \cup \cdots \cup \left(\!\!\binom{C_i}{m_i}\!\!\right)$, where $\left(\!\!\binom{C_i}{j}\!\!\right)$ denotes the set of all multisets of weight j supported on C_i **do**
6 $\quad$ Recurse continuations$(\mathbf{S}) \leftarrow \mathrm{tails}(i+1,\ m_i - |\mathbf{S}|,\ \Gamma(\mathbf{S}))$;
7 $\quad$ Append out $\leftarrow$ out $\cup \{\mathbf{S} \cup \mathbf{T} \mid \mathbf{T} \in \mathrm{continuations}(\mathbf{S})\}$;
8 **end**
9 Return out;

Given the dual interaction graph as a random-access dictionary, (i.e., one can query an arbitrary node and receives its neighbors), the runtime of this algorithm is $O(m\mathfrak{d}\mathcal{C})$, where $\mathcal{C}$ is the number of clusters output. The main cost is computing $\Gamma(\mathbf{S})$ from a given multiset $\mathbf{S}$ on V_i: this takes time $\sum_{v \in \mathrm{Supp}\,\mathbf{S}} |\Gamma(v)| \leq \mathfrak{d}|\mathbf{S}|$, where the inequality uses that G is degree $\leq \mathfrak{d}$. For every cluster $\mathbf{S}_0 \sqcup \cdots \sqcup \mathbf{S}_{m-1}$, this computation occurs once for each of the $\mathbf{S}_i$'s. Since every cluster has a total weight of m, this gives an upper bound of $O(m\mathfrak{d}\mathcal{C})$ for all such computations.

3.5 Estimating cluster derivatives

The goal of this subsection is to prove the following bound on cluster derivatives.

Proposition 3.8. *Consider a Hamiltonian $\{(a, E_a, \lambda_a)\}$. Let $\mathbf{W} = \{(a, \mu(a))\}$ be a cluster of the associated dual interaction graph $\mathfrak{G}$ with total weight $m + 1 \geq 1$. Then*

$$\left| \frac{1}{\mathbf{W}!} \mathcal{D}_{\mathbf{W}} \mathcal{L} \right| \leq (2e(\mathfrak{d} + 1)\beta)^{m+1}. \tag{37}$$

To prove this, we will bound $|\mathcal{D}_{\mathbf{W}} \mathcal{L}|$ by a quantity that depends on a simple graph constructed from $\mathbf{W}$ and the dual interaction graph $\mathfrak{G}$, which we now define.

Recall from Definition 3.3 that given a set S, a multiset S^μ of elements of S is a set $\{(s, \mu(s)) \mid s \in S\}$ where $\mu(s) \in \mathbb{Z}_{\geq 0}$ is the multiplicity of s. We also write $S^\mu = \{(s_i, \mu(s_i))\} = \{\{s_1, \ldots, s_2, \ldots\}\}$ where s_i is repeated exactly $\mu(s_i)$ times. The size of S^μ is $|S^\mu| = \sum_{s \in S} \mu(s)$. The support of S^μ is $\operatorname{Supp} S^\mu = \{s \in S \mid \mu(s) > 0\}$. We write $S^\mu!$ to mean $\prod_{s \in S} (\mu(s)!)$. If S is the node set of a simple graph F, we define a simple graph $\operatorname{Gra}(S^\mu)$ as follows. The set of nodes are

$$\operatorname{Mar}(S^\mu) := \{(s, i) \in (\operatorname{Supp} S^\mu) \times \mathbb{Z}_{>0} \mid 1 \leq i \leq \mu(s)\}; \tag{38}$$

in other words, there are exactly $\mu(s)$ nodes corresponding to s for each $s \in S$, so there are $|S^\mu|$ nodes in total. In $\operatorname{Gra}(S^\mu)$, an edge between (s, i) and (s', i') exists if and only if either $s = s'$ or (s, s') is an edge of F. In particular, the induced subgraph of $\{(s, i)\}_i$ for any given s is a clique.

Since the dual interaction graph $\mathfrak{G}$ serves as the underlying graph for Hamiltonian terms, for any cluster $\mathbf{W}$ of Hamiltonian terms we have a corresponding $\operatorname{Gra}(\mathbf{W})$. In $\operatorname{Gra}(\mathbf{W})$ there are $|\mathbf{W}|$ nodes in total, and an edge between two nodes $(a, i), (a', i')$ exists iff the Hamiltonian terms E_a and $E_{a'}$ have overlapping supports. Note that if all multiplicities of $\mathbf{W}$ are either 0 or 1, then $\operatorname{Gra}(\mathbf{W})$ is an induced subgraph of $\mathfrak{G}$, but is not otherwise. Most of this section will be devoted to proving the following lemma, which implies Proposition 3.8.

Lemma 3.9 ([WA22])**.** *Denote by $\deg(v)$ the number of neighbors of any node v in $\operatorname{Gra}(\mathbf{W})$. Then,*

$$\left| \mathcal{D}_{\mathbf{W}} \mathcal{L} \right| \leq |\beta|^{|\mathbf{W}|} \prod_{v \in \operatorname{Mar} \mathbf{W}} (2 \deg(v)). \tag{39}$$

Proof of Proposition 3.8. It follows from the definition of $\operatorname{Gra}(\mathbf{W})$ that

$$\deg((b, i)) = \mu(b) - 1 + \sum_{a \in \Gamma(b)} \mu(a) \tag{40}$$

for any $b \in \operatorname{Supp} \mathbf{W}$, where $\Gamma(b)$ is the set of all neighbors of b in $\mathfrak{G}$ that appear in $\mathbf{W}$. Further note that

$$\sum_{b \in \operatorname{Supp} \mathbf{W}} \deg((b, 1)) = \sum_{b \in \operatorname{Supp} \mathbf{W}} \left(\mu(b) - 1 + \sum_{a \in \Gamma(b)} \mu(a) \right) \leq m + 1 + \sum_{b \in \operatorname{Supp} \mathbf{W}} \sum_{a \in \Gamma(b)} \mu(a) \leq m + \mathfrak{d}(m+1). \tag{41}$$

because $\mu(a)$ appears in $\sum_b \sum_{a \in \Gamma(b)}$ at most $\mathfrak{d}$ times. We can now apply Lemma 3.9.

$$
\begin{aligned}
\frac{1}{\mathbf{W}!}|\mathcal{D}_{\mathbf{W}}\mathcal{L}| &\leq \frac{(2\beta)^{m+1}}{\mathbf{W}!} \prod_{b \in \mathrm{Supp}\,\mathbf{W}} \prod_{i=1}^{\mu(b)} \deg((b,i)) && \text{by Lemma 3.9} \\
&= (2\beta)^{m+1} \prod_{b \in \mathrm{Supp}\,\mathbf{W}} \frac{1}{\mu(b)!}\Big(\mu(b) - 1 + \sum_{a \in \Gamma(b)} \mu(a)\Big)^{\mu(b)} && \text{by Equation (40)} \\
&\leq (2e\beta)^{m+1} \prod_{b \in \mathrm{Supp}\,\mathbf{W}} \left(\frac{\mu(b) - 1 + \sum_{a \in \Gamma(b)} \mu(a)}{\mu(b)}\right)^{\mu(b)} && \text{because } u! \geq u^u e^{-u} \\
&\leq (2e\beta)^{m+1} \left(\frac{(1+\mathfrak{d})(m+1)}{m+1}\right)^{m+1}.
\end{aligned}
$$

$$\tag{42}$$

The last inequality uses Lemma 3.10 below and Equation (41). $\qquad\square$

Lemma 3.10. *Let $\mu_1, \ldots, \mu_n > 0$ be real numbers, and $y_1, \ldots, y_n \geq 0$ be real numbers. Then*

$$
\left(\frac{y_1}{\mu_1}\right)^{\mu_1} \cdots \left(\frac{y_n}{\mu_n}\right)^{\mu_n} \leq \left(\frac{y_1 + \cdots + y_n}{\mu_1 + \cdots + \mu_n}\right)^{\mu_1 + \cdots + \mu_n},
\tag{43}
$$

where the equality holds when $y_j/\mu_j = (\sum_i y_i)/(\sum_i \mu_i)$ for all j.

Proof. If any of y_i is zero, the inequality is trivial. Assume $y_i > 0$ for all i. Taking log of both sides and dividing by $\sum_i \mu_i$, we have

$$
\sum_{i=1}^n \frac{\mu_i}{\sum_j \mu_j} \log\left(\frac{y_i}{\mu_i}\right) \leq \log\left(\frac{y_1 + \cdots + y_n}{\mu_1 + \cdots + \mu_n}\right).
\tag{44}
$$

This is Jensen's inequality applied to a concave function log. $\qquad\square$

Proof of Lemma 3.9.

We will adopt the approach in [WA22], making some short-cuts.[6] Our combinatorics will be self-contained; we do not assume any prior knowledge of Tutte polynomials or chromatic polynomials, which were used in [WA22]. However, in essence, the proof here is due to [WA22].

We will use multisets of clusters. All the general remarks above on multisets continue to apply. Consider a multiset of clusters $\mathbf{W}_i$, which we denote $\mathsf{P} = \{\{\mathbf{W}_1, \ldots\}\} = \{(\mathbf{W}, \mu(\mathbf{W}))\}$. We write $|\mathsf{P}| = \sum_{\mathbf{W}} \mu(\mathbf{W})$, $\mathrm{Supp}\,\mathsf{P} = \{\mathbf{W} \mid \mu(\mathbf{W}) > 0\}$, and $\mathsf{P}! = \prod_{\mathbf{W} \in \mathrm{Supp}\,\mathsf{P}} (\mu(\mathbf{W})!)$. The set (not multiset) of all *connected* clusters defines a simple graph where there is an edge between $\mathbf{W}$ and $\mathbf{W}'$ if and only if their multiset union $\mathbf{W} \cup \mathbf{W}'$ (obtained by summing the multiplicities) is a connected cluster. With this, a multiset P of connected clusters defines a simple graph $\mathrm{Gra}(\mathsf{P})$: there are $|\mathsf{P}|$ nodes in total, and an edge between two nodes corresponding to $\mathbf{W}$ and $\mathbf{W}'$ exists if and only if $\mathbf{W} \cup \mathbf{W}'$ is connected.

---

[6] We thank the authors of [WA22] for pointing out a problem in our earlier version of this proof, which traces back to [KS20, App. C].

Counting partitions of a cluster Consider a cluster $\mathbf{W}$. From an (unordered) partition $\{W_1, \ldots, W_p\}$ of the *graph node set* $\mathrm{Mar}(\mathbf{W})$, we can get an (unordered) partition of the *cluster* $\mathbf{W}$ simply by forgetting all of the labels on the terms in $\mathbf{W}$. By a "partition" of a cluster, we mean a multiset $\mathsf{P} = \{\{\mathbf{W}_1, \ldots, \mathbf{W}_p\}\}$ such that their multiset union is $\mathbf{W}$. Conversely, given a cluster partition P, the number of graph partitions that get mapped to P by "forgetting" is

$$\frac{\mathbf{W}!}{\mathsf{P}! \prod_{\ell=1}^{p} (\mathbf{W}_\ell!)}. \tag{45}$$

Here, $\mathbf{W}!$ is the number of ways to assign labels to P if we give an arbitrary ordering to both the clusters and the terms within the clusters; $\prod_{\ell=1}^{p}(\mathbf{W}_\ell!)$ addresses the overcounting from ordering each term $\mathbf{W}_\ell$, since swapping labels within a cluster doesn't change the cluster; and $\mathsf{P}!$ addresses the overcounting from ordering P, since swapping the labels across two identical clusters doesn't change the cluster partition. For example, consider a Hamiltonian with two terms E_a, E_b with overlapping support. Then the cluster partition $\mathsf{P} = \{\{a, b\}, \{a, b\}, \{b, b\}\}$ of the connected cluster $\mathbf{W} = \{(a, 2), (b, 4)\}$ corresponds to 12 graph partitions, using the reasoning above.

$$(((a, f), (b, h)), ((a, g), (b, i)), ((b, j), (b, k))) \qquad \{f, g\} = [2], \{h, i, j, k\} = [4],\ 2!4! = 48 \text{ choices}$$
$$\mapsto (\{(a, f), (b, h)\}, \{(a, g), (b, i)\}, \{(b, j), (b, k)\}) \quad \text{overcounts by a } (1!1!)(1!1!)(2!) \text{ factor, 24 choices}$$
$$\mapsto \{\{(a, f), (b, h)\}, \{(a, g), (b, i)\}, \{(b, j), (b, k)\}\} \qquad \text{overcounts by a } 2!1! \text{ factor, 12 choices}$$

We write $\mathrm{PaC}(F)$ for a graph F (not necessarily simple) to mean the collection of all graph partitions of F into *connected* induced subgraphs. For any integer $n \geq 1$, let $\chi^*(n, F)$ denote the number of all node colorings (two end nodes of an edge having different colors) on F using exactly n colors.

Lemma 3.11. *For any nonempty cluster $\mathbf{W}$ we have*

$$|\mathcal{D}_\mathbf{W}\mathcal{L}| \leq \beta^{|\mathbf{W}|} \sum_{\mathfrak{P} \in \mathrm{PaC}(\mathrm{Gra}(\mathbf{W}))} \left| \sum_{n=1}^{|\mathfrak{P}|} \frac{(-1)^{n-1}}{n} \chi^*(n, \mathrm{Gra}(\mathfrak{P})) \right|. \tag{46}$$

This is a repackaging of [WA22, App. B.1].

Proof. Recall that $\mathcal{Z} = \frac{1}{\mathsf{D}} \mathrm{Tr} \exp(-\beta H)$. There is a formal cluster expansion for $\mathcal{Z}$ (which is meaningful in view of Lemma 2.5):

$$\mathcal{Z} = 1 + \sum_{k \geq 1} \sum_{\mathbf{W}: \text{ weight } k} \frac{\lambda^\mathbf{W}}{\mathbf{W}!} \mathcal{D}_\mathbf{W}\mathcal{Z} \tag{47}$$

where $\mathbf{W}$ is not always connected. The cluster derivative $\mathcal{D}_\mathbf{W}\mathcal{Z}$ factorizes if $\mathbf{W}$ is not connected (Proposition 3.5). Let $\mathsf{P}_{\max}(\mathbf{W})$ be the set of maximal connected subclusters of $\mathbf{W}$. Then, using $\log(1 + x) = \sum_{n=1}^{\infty} \frac{(-x)^{n-1}}{n}$ for small x (which is again meaningful in view of Lemma 2.5),

$$\mathcal{Z} = 1 + \sum_{k \geq 1} \sum_{\substack{\mathbf{W}: \\ |\mathbf{W}|=k}} \prod_{\mathbf{V} \in \mathrm{Supp}\, \mathsf{P}_{\max}(\mathbf{W})} \frac{\lambda^\mathbf{V}}{\mathbf{V}!} \mathcal{D}_\mathbf{V}\mathcal{Z}, \tag{48}$$

$$\log \mathcal{Z} = \sum_{n=1}^{\infty} \frac{(-1)^{n-1}}{n} \sum_{k_1, \ldots, k_n \geq 1} \sum_{\substack{\mathbf{W}_1, \ldots, \mathbf{W}_n: \\ |\mathbf{W}_i|=k_i}} \prod_{i=1}^{n} \left(\prod_{\mathbf{V}_i \in \mathrm{Supp}\, \mathsf{P}_{\max}(\mathbf{W}_i)} \frac{\lambda^{\mathbf{V}_i}}{\mathbf{V}_i!} \mathcal{D}_{\mathbf{V}_i}\mathcal{Z} \right)$$

Now, we rearrange this sum as a β-power series; $\mathcal{D}_{\mathbf{V}}\mathcal{Z}$ carries $\beta^{|\mathbf{V}|}$. The double product can be written as a product over a multiset $\mathsf{P} = \{\{\mathbf{V}_{1,1}, \mathbf{V}_{1,2}, \ldots, \mathbf{V}_{n,1}, \mathbf{V}_{n,2}, \ldots\}\}$ of connected clusters. This multiset P of connected clusters obeys a special property that each cluster $\mathbf{V}_{i,j}$ is assigned a label i, and among those of a given label i no two clusters become connected by taking the multiset union. That is, the labels give a unique node coloring on $\mathrm{Gra}(\mathsf{P})$ with exactly n colors.

Conversely, with a node coloring on $\mathrm{Gra}(\mathsf{P})$ with exactly n colors, we see that the collection of nodes of a given color i defines $\mathbf{W}_i$. But this converse direction is many-to-one: a cluster in P with multiplicity ≥ 2 gives two or more nodes in $\mathrm{Gra}(\mathsf{P})$ that have all different colors, and permuting colors among these "duplicate" nodes of $\mathrm{Gra}(\mathsf{P})$ gives the same $\mathbf{W}_i$'s. We see that precisely $\mathsf{P}!$ different colorings give the same $\mathbf{W}_i$'s. The multiset P is a cluster partition of the multiset union $\mathbf{W} = \bigcup_i \bigcup_j \mathbf{V}_{i,j}$, and $|\mathbf{W}| = k = k_1 + \cdots + k_n$ is the order in β of the double product. Hence, letting ℓ assume all tuples (i,j) that index $\mathbf{V}_{i,j}$, we see that

$$\log \mathcal{Z} = \sum_{k=1}^{\infty} \sum_{\substack{\mathbf{W}: \\ |\mathbf{W}|=k, \\ \mathbf{W}\ \mathrm{connected}}} \sum_{\substack{\mathsf{P}=\{\{\mathbf{V}_\ell\}\}: \\ \bigcup_\ell \mathbf{V}_\ell = \mathbf{W}, \\ \mathbf{V}_\ell\ \mathrm{connected}}} \left(\sum_{n=1}^{|\mathsf{P}|} \frac{(-1)^{n-1}}{n} \frac{\chi^*(n, \mathrm{Gra}(\mathsf{P}))}{\mathsf{P}!} \right) \prod_{\ell=1}^{|\mathsf{P}|} \frac{\lambda^{\mathbf{V}_\ell}}{\mathbf{V}_\ell!} \mathcal{D}_{\mathbf{V}_\ell}\mathcal{Z}. \qquad (49)$$

Next, we rewrite the sum over P as a sum over graph partitions $\mathfrak{P}$ of $\mathrm{Gra}(\mathbf{W})$ into connected induced subgraphs. By Eq. (45), for each P there are exactly $\frac{\mathbf{W}!}{\mathsf{P}! \prod_{\ell=1}^{|\mathsf{P}|}(\mathbf{V}_\ell!)}$ different graph partitions $\mathfrak{P}$ that give P. Each induced subgraph in $\mathfrak{P}$ is connected iff the corresponding cluster is connected. Therefore,

$$\log \mathcal{Z} = \sum_{\substack{\mathbf{W}: \\ \mathrm{connected} \\ \mathrm{nonempty}}} \frac{1}{\mathbf{W}!} \sum_{\mathfrak{P}\in\mathrm{PaC}(\mathrm{Gra}(\mathbf{W}))} \left(\sum_{n=1}^{|\mathsf{P}|} \frac{(-1)^{n-1}}{n} \chi^*(n, \mathrm{Gra}(\mathsf{P})) \right) \prod_{\ell=1}^{|\mathsf{P}|} \lambda^{\mathbf{V}_\ell} \mathcal{D}_{\mathbf{V}_\ell}\mathcal{Z}. \qquad (50)$$

where $\mathfrak{P} \mapsto \mathsf{P} = \{\{\mathbf{V}_\ell\}\}$ is implicit. Now we can read off

$$\mathcal{D}_{\mathbf{W}}\mathcal{L} = \mathcal{D}_{\mathbf{W}} \log \mathcal{Z} = \sum_{\mathfrak{P}\in\mathrm{PaC}(\mathrm{Gra}(\mathbf{W}))} \left(\sum_{n=1}^{|\mathsf{P}|} \frac{(-1)^{n-1}}{n} \chi^*(n, \mathrm{Gra}(\mathsf{P})) \right) \prod_{\ell=1}^{|\mathsf{P}|} \mathcal{D}_{\mathbf{V}_\ell}\mathcal{Z} \qquad (51)$$

$$= \sum_{\mathfrak{P}\in\mathrm{PaC}(\mathrm{Gra}(\mathbf{W}))} \left(\sum_{n=1}^{|\mathfrak{P}|} \frac{(-1)^{n-1}}{n} \chi^*(n, \mathrm{Gra}(\mathfrak{P})) \right) \prod_{\ell=1}^{|\mathfrak{P}|} \mathcal{D}_{\mathbf{V}_\ell}\mathcal{Z}.$$

The proof is completed by bounding $\mathcal{D}_{\mathbf{V}}\mathcal{Z}$ as follows. If $\mathbf{V} = \{\{a_1, \ldots, a_k\}\}$, then

$$\mathcal{D}_{\mathbf{V}}\mathcal{Z} = \mathcal{D}_{\mathbf{V}} \frac{1}{\mathsf{D}} \mathrm{Tr}(e^{-\beta H}) = \mathcal{D}_{\mathbf{V}} \sum_{n=0}^{\infty} \frac{(-\beta)^n}{n!} \frac{1}{\mathsf{D}} \mathrm{Tr}(H^n)$$

$$= \mathcal{D}_{\mathbf{V}} \frac{(-\beta)^k}{k!} \frac{1}{\mathsf{D}} \mathrm{Tr}(H^k) = \frac{(-\beta)^k}{k!} \sum_{\sigma\in\mathrm{S}_k} \frac{1}{\mathsf{D}} \mathrm{Tr}(E_{a_{\sigma(1)}} \cdots E_{a_{\sigma(k)}}). \qquad (52)$$

Since $\|E_a\| \leq 1$, we see $|\mathcal{D}_{\mathbf{V}}\mathcal{Z}| \leq |\beta|^k$. $\qquad\qquad \square$

Lemma 3.9 is proved by Lemma 3.11 and a combinatorial estimate in Lemma 3.12 below. We recall some elements of graph combinatorics.

Let F be a possibly nonsimple graph with self-loops and multiple edges. Given an edge e (not a self-loop) of a graph F the *contraction* of the edge e, denoted by F/e, is the graph obtained by

removing the edge e and merging the two end points of e into one vertex; and the *deletion* of e, denoted by $F \setminus e$, is the graph obtained by removing the edge e from F. The following identities are standard. For any graph F (not necessarily simple),

$$\chi^*(k, F \setminus e) = \chi^*(k, F) + \chi^*(k, F/e), \tag{53}$$

$$\tau(F) = \tau(F \setminus e) + \tau(F/e) \tag{54}$$

where $\tau(F)$ is the number of all spanning trees of F. If F is disconnected, $\tau(F) = 0$. For Eq. (53), a coloring of $F \setminus e$ either colors the endpoints of e the same or different: the colorings where they are colored differently are exactly the set of colorings of F, and the colorings where they are colored the same correspond to the set of colorings of F/e. For Eq. (54), a spanning tree of $\tau(F)$ either contains the edge e or it does not: the spanning trees that do not contain e are exactly the spanning trees of $\tau(F \setminus e)$, and the spanning trees that do contain e correspond to spanning trees of F/e.

Lemma 3.12. *Let G be a nonempty connected graph with n vertices. Then,*

$$\sum_{\mathfrak{P} \in \mathrm{PaC}(G)} \left| \sum_{k=1}^{|\mathfrak{P}|} \frac{(-1)^{k-1}}{k} \chi^*(k, \mathrm{Gra}(\mathfrak{P})) \right| \leq 2^{n-1} \tau(G). \tag{55}$$

This can be proved via perhaps more canonical approach using Tutte polynomials and chromatic polynomials [WA22], but we directly use the founding principle of these polynomials—the deletion-contraction recurrence.

Proof. For any nonempty (not necessarily simple) graph F on n vertices, we define a rational number $\eta(F)$, which will turn out to be a nonnegative integer, by

$$\eta(F) = \sum_{k=1}^{n} \frac{(-1)^{n-k}}{k} \chi^*(k, F) \tag{56}$$

Now, consider an edge $e = (a, b)$ where $a \neq b$. Since F/e has $n - 1$ vertices, Eq. (53) gives

$$\begin{aligned}
\eta(F/e) + \eta(F \setminus e) &= \sum_{k=1}^{n-1} \frac{(-1)^{n-1-k}}{k} \chi^*(k, F/e) + \sum_{k=1}^{n} \frac{(-1)^{n-k}}{k} \chi^*(k, F \setminus e) \\
&= -\sum_{k=1}^{n} \frac{(-1)^{n-k}}{k} \chi^*(k, F/e) + \sum_{k=1}^{n} \frac{(-1)^{n-k}}{k} \chi^*(k, F \setminus e) \\
&= \eta(F) \tag{57}
\end{aligned}$$

where the second equality is because $\chi^*(n, F/e) = 0$. Notice that this is the same recursive formula as that of τ in Eq. (54). Our goal will be to show $\eta(F) \leq \tau(F)$: by applying Eq. (57) and Eq. (54) to reduce to cases with fewer edges via deletion and contraction, it suffices to show this for F with no edges, and only self-loops.

We first show that, for S the graph of n isolated vertices with no edges, $\eta(S) = \delta_{n,1}$ where δ is the Kronecker delta. The number of all proper colorings of S using k or fewer colors is $\chi(k, S) = k^n = (x\partial_x)^n x^k|_{x=1}$, where x is an indeterminant. By inclusion-exclusion, we have

$\chi^*(k, T) = \sum_{j=0}^k \binom{k}{j}(-1)^{k-j}\chi(j, T)$. Hence,

$$\begin{aligned}
\eta(T) &= \sum_{k=1}^n \frac{(-1)^{n-k}}{k} \sum_{j=0}^k \binom{k}{j}(-1)^{k-j}(x\partial_x)^n x^j \Big|_{x=1} \\
&= (-x\partial_x)^n\Big|_{x=1} \sum_{k=1}^n \frac{(-1)^k}{k}(x-1)^k \\
&= (-x\partial_x)^{n-1}\Big|_{x=1} x \sum_{k=1}^\infty (-1)^{k-1}(x-1)^{k-1} \\
&= (-x\partial_x)^{n-1}\Big|_{x=1} 1 = \delta_{n,1}.
\end{aligned}$$ (58)

So, $\eta(S) = \tau(S) = \delta_{n,1}$ for graphs S without any edges or loops. For graphs F with some loops but without any edges, $\eta(F) = 0$ since a self-loop prohibits any proper coloring, but $\tau(F)$ may be positive if $n = 1$. Hence, $\eta(F) \leq \tau(F)$ for F without any edges.

We conclude that $0 \leq \eta(F) \leq \tau(F)$ for any nonempty graph F.[7] It follows that

$$\left| \sum_{k=1}^{|\mathfrak{P}|} \frac{(-1)^{k-1}}{k}\chi^*(k, \mathrm{Gra}(\mathfrak{P})) \right| = \eta(\mathrm{Gra}(\mathfrak{P})) \leq \tau(\mathrm{Gra}(\mathfrak{P})).$$ (59)

The proof is completed by observing the following [WA22, Lem. 21]. Consider a spanning tree T of G and a subset E_0 of some edges of T. Let E_1 be the set of all edges of T not in E_0; $E_0 \sqcup E_1$ is the total edge set of T. We obtain a graph partition $\mathfrak{P} \in \mathrm{PaC}(G)$ of G into connected subgraphs, namely the connected components of E_0-deleted subgraph of T. We also obtain a spanning tree of $\mathrm{Gra}(\mathsf{P})$, obtained by contracting all edges of E_1 from T. (Any contraction on a tree is a tree.) Hence, given G, we have a map from pairs (T, E_0) of a spanning tree and its subset of edges to pairs $(\mathfrak{P} \in \mathrm{PaC}(G), T')$ of a graph partition $\mathfrak{P}$ and a spanning tree of $\mathrm{Gra}(\mathfrak{P})$. This map is surjective: by choosing a spanning tree in each party of $\mathfrak{P}$ we have E_0, and by choosing some edges, one among those that would merge to an edge in T' upon contraction of the spanning trees of the parties of $\mathfrak{P}$, we construct a spanning tree T of G whose edge set contains E_0. This surjection gives

$$\sum_{\mathfrak{P} \in \mathrm{PaC}(G)} \tau(\mathrm{Gra}(\mathfrak{P})) \leq 2^{|G|-1}\tau(G).$$ (60)

where $|G| - 1$ is the number of all edges in any spanning tree T of G. □

Combining Lemmas 3.11 and 3.12, we have

$$|\mathcal{D}_\mathbf{E}\mathcal{L}| \leq \sum_{\mathfrak{P} \in \mathrm{PaC}(G)} \left| \sum_{k=1}^{|\mathfrak{P}|} \frac{(-1)^{k-1}}{k}\chi^*(k, \mathrm{Gra}(\mathfrak{P})) \right| \leq 2^{n-1}\tau(G).$$ (61)

where $n = |\mathbf{E}| = |G|$. It remains to show that $\tau(G)$, the number of all spanning trees of G, is at most the product of degrees of nodes. Fix a root, an arbitrary node of G. Given a spanning tree of G, we choose a unique edge attached to each node different from the root along which the unique shortest path to the root from the node traverses. Each nonroot node a have $\deg(a)$, the degree of a, choices at most, implying $\tau(G) \leq \prod_{a \in G} \deg(a)$. This complete the proof of Lemma 3.9.

---

[7][WA22] shows that $\eta(F)$ equals the value of the Tutte polynomial T_F at $(1, 0)$, and quotes the facts that $T_F(1, 0) \leq T_F(1, 1)$ and that $T_F(1, 1) = \tau(F)$.

3.6 Computing cluster derivatives

Proposition 3.13. *Consider a Hamiltonian $\{(a, E_a, x_a)\}$ where every term E_a is a tensor product of Pauli matrices and is supported on at most L qubits. Then, there is a deterministic algorithm running in time $(8^m + L)\operatorname{poly}(m)$ such that for every cluster $\mathbf{W}$ of total weight $m+1$ it outputs $\frac{1}{\beta^{m+1}\mathbf{W}!}\mathcal{D}_{\mathbf{W}}\mathcal{L}$ exactly as a rational number.*

Proof. Let $\mathbf{W} = \{(a, \mu_a)\}$, and suppose there are n distinct a in $\mathbf{W}$. By assumption, $n \leq m+1$. Without loss of generality, we relabel the indices so that a iterates over $[n]$: $\mathbf{W} = \{(1, \mu_1), \ldots, (n, \mu_n)\}$. The expression $\mathcal{D}_{\mathbf{W}}\mathcal{L}$ depends only on $\{x_1, \ldots, x_n\}$, since it evaluates the derivative at the origin. Hence, to evaluate it, we may assume that our Hamiltonian is simply $H = \sum_{a=1}^n x_a E_a$. This restricted Hamiltonian has operator norm at most n.

We can further simplify because $\mathcal{D}_{\mathbf{W}}\mathcal{L}$ depends only on the β^{m+1} term of the Taylor expansion of $\mathcal{L}$. So, we can freely truncate and avoid worrying about higher-order terms. In particular, we will truncate the Taylor expansion of the exponential in $\mathcal{L}$ to get

$$\mathcal{D}_{\mathbf{W}}\mathcal{L} = \beta^{m+1}\mathcal{D}_{\mathbf{W}}\log\operatorname{Tr}\exp\Big(-\sum_{j=1}^n x_j E_j\Big) = \beta^{m+1}\mathcal{D}_{\mathbf{W}}\log r. \tag{62}$$

for r the polynomial function[8]

$$r(x_1, \ldots, x_n) = \sum_{k=0}^{m+1} \frac{(-1)^k}{k!}\frac{1}{\mathsf{D}}\operatorname{Tr}\Big[\Big(\sum_{j=1}^n x_j E_j\Big)^k\Big] \in \mathbb{C}[x_1, \ldots, x_n], \tag{63}$$

$$r(0, \ldots, 0) = 1. \tag{64}$$

This expression normalizes by $\frac{1}{\mathsf{D}}$, which vanishes after taking $\mathcal{D}_{\mathbf{W}}\log$. To compute the derivative of $\log r$, we first make a general observation, whose proof will be given later.

Lemma 3.14. *Let $f : \mathbb{R}^n \to \mathbb{R}$ be a smooth (infinitely differentiable in any variable in any order) function. Then, $f(x)$ for $x = (x_1, \ldots, x_n)$ satisfies*

$$\frac{\partial^\mu f}{\partial x_1^{\mu_1}\cdots\partial x_n^{\mu_n}} = \lim_{\alpha\to 0}\frac{1}{\alpha^\mu}\sum_{z_1=0}^{\mu_1}\cdots\sum_{z_n=0}^{\mu_n}(-1)^{|z|+\mu}\binom{\mu_1}{z_1}\cdots\binom{\mu_n}{z_n}f(x+\alpha z) \tag{65}$$

where $\mu = \mu_1 + \cdots + \mu_n$, $z = (z_1, \ldots, z_n) \in \mathbb{Z}_{\geq 0}^n$ and $|z| = z_1 + \cdots + z_n$.

Lemma 3.14 implies that there are some integers c_z depending on z such that we can write

$$\mathcal{D}_{\mathbf{W}}\log r = \lim_{\alpha\to 0}\frac{1}{\alpha^{m+1}}\sum_{\substack{z \in \mathbb{Z}_{\geq 0}^n \\ z_j \leq \mu_j\,\forall j\in[n]}} c_z \log g(\alpha; z) \tag{66}$$

$$\text{where } g(\alpha; z) = r(z_1\alpha, \ldots, z_n\alpha) \in \mathbb{C}[\alpha]. \tag{67}$$

We treat $g(\alpha; z)$ as a univariate function parametrized by an integer vector $z \in \mathbb{Z}^n$. There are at most 2^{m+1} summands in Eq. (66). Further, the limit in Eq. (66) can be computed by L'Hospital's

---

[8]In this proof of Proposition 3.13, Eq. (63) is the only place we use the fact that E_a are Pauli operators. If we used an arithmetic model of computation, the operators E_a can be more general.

rule:

$$\mathcal{D}_\mathbf{W} \log r = \frac{1}{(m+1)!} \sum_{\substack{z \in \mathbb{Z}_{\geq 0}^n \\ z_j \leq \mu_j \, \forall j \in [n]}} c_z \Big[\partial_\alpha^{m+1} \log g(\alpha; z) \Big]_{\alpha=0} \tag{68}$$

$$= \frac{1}{(m+1)!} \sum_{\substack{z \in \mathbb{Z}_{\geq 0}^n \\ z_j \leq \mu_j \, \forall j \in [n]}} c_z h_m(\alpha; z)|_{\alpha=0} \tag{69}$$

where we define $h_m(\alpha; z) := g(\alpha; z)^{m+1} \partial_\alpha^{m+1} \log g(\alpha; z)$. The above equality holds because $g(0; z) = 1$. The function $h_m(\alpha; z)$ is in fact a polynomial and satisfies a straightforward recursion:

$$\begin{aligned}
h_0 &= g \, \partial_\alpha \log g \\
&= \partial_a g \\
h_t &= g^{t+1} \partial_\alpha^{t+1} \log g \\
&= \partial_\alpha(g^t \partial_\alpha^t \log g)g - t(g^t \partial_\alpha^t \log g)\partial_\alpha g \\
&= \partial_\alpha(h_{t-1})g - t h_{t-1}\partial_\alpha g
\end{aligned} \tag{70} \tag{71}$$

In summary, we have reduced the problem to finding the constant term of h_m, a recursively-defined polynomial, for various choices of z:

$$\frac{1}{\beta^{m+1}\mathbf{W}!} \mathcal{D}_\mathbf{W}\mathcal{L} = \frac{1}{\mathbf{W}!(m+1)!} \sum_z c_z \frac{h_m(0; z)}{g(0; z)^{m+1}} = \frac{1}{\mathbf{W}!(m+1)!} \sum_z c_z h_m(0; z). \tag{72}$$

For a fixed z, the function $g(\alpha; z)$ is a polynomial in α. Using the notation that $[\alpha^k]g(\alpha; z)$ is the coefficient of α^k in $g(\alpha; z)$, we have that

$$g(\alpha; z) = 1 + [\alpha]g(\alpha; z) \cdot \alpha + \cdots + [\alpha^{m+1}]g(\alpha; z) \cdot \alpha^{m+1} \tag{73}$$

$$[\alpha^k]g(\alpha; z) = \frac{(-1)^k}{k!}\frac{1}{\mathsf{D}} \mathrm{Tr}\left[\left(\sum_{j=1}^n z_j E_j\right)^k\right] \tag{74}$$

Time complexity. Now that we've established the form of the expression that we will compute, we will now discuss the time complexity necessary to compute it. First, we consider computing the coefficients $[\alpha^k]g(\alpha; z)$ as seen in Eq. (74).

A binary representation of the n Pauli operators E_a and Gauss elimination reveals a minimal set (multiplicative basis) of Pauli operators which can generate all the n operators by multiplications together with phase factors $\pm 1, \pm i$ [AG04]. Since the set of all those n Pauli operators are supported on at most nL qubits, It takes time $O(n^3 L)$ to find a multiplicative basis. Once we have a multiplicative basis, we can find another set of Pauli operators $\tilde{E}_1, \ldots, \tilde{E}_n$ on n qubits, preserving all the pairwise commutation relations and the multiplicative independence. The procedure is simple: For the first basis element, choose $\tilde{E}_1 = Z_1$. If the second basis element commutes with the first, choose $\tilde{E}_2 = Z_2$, or otherwise, choose $\tilde{E}_2 = X_1 Z_2$. Inductively, for t-th basis element we choose $\tilde{E}_t$ to be Z_t multiplied by an appropriate Pauli operator to preserve all the commutation relations with $\tilde{E}_{t'}$ where $t' < t$. Thus, it takes time $O(n^3 L)$ to find a faithful representation $\tilde{E}_a$ of n Pauli operators E_a.

Equipped with a faithful representation $\tilde{E}_a$, Eq. (74) is evaluated by powering a matrix $P = \sum_j z_j \tilde{E}_a$ of dimension 2^n. The normalization constant also changes from $\frac{1}{\mathsf{D}}$ to $\frac{1}{2^n}$. The matrix

P has $O(n)$ entries in each column and in each row, and they are in $\mathbb{Z}[\sqrt{-1}]$ since the E_a's are Pauli operators. Hence, multiplying a 2^n-dimensional matrix by P takes $O(4^n n)$ integer arithmetic operations. Since we raise P to k-th power, we can compute $\mathrm{Tr}(P^k)$ in $O(4^n nk)$ integer arithmetic operations, and we can compute the trace for all $k \in [m+1]$ in $O(4^n nm)$ operations. This is the numerator of $g(;z)_k$, and we maintain $g(;z)_k$ as a rational number by maintaining its numerator and denominator. By multiplying both by $(k+1)(k+2)\cdots(m+1)$, we can standardize these rational number representations of $g(;z)_k$ to have be some integer in the numerator and $(m+1)!2^n$ in the denominator.

Once we have representations for $g(;z)_k$, we can compute $h_m(0;z)$, the constant term of $h_m(\alpha;z)$, via the recursive formula Eq. (71). This takes $O(m^3)$ integer arithmetic operations, since $[\alpha^k]h_t$, the coefficient of α^k in h_t, satisfies

$$[\alpha^k]h_0 = (k+1)[\alpha^{k+1}]g \tag{75}$$

$$[\alpha^k]h_t = \sum_{j=0}^{k}\Big[[\alpha^j]h'_{t-1}[\alpha^{k-j}]g - t \cdot [\alpha^j]h_{t-1}[\alpha^{k-j}]g'\Big]$$

$$= \sum_{j=0}^{k}\Big[(j+1)[\alpha^{j+1}]h_{t-1}[\alpha^{k-j}]g - t(k-j+1)[\alpha^j]h_{t-1}[\alpha^{k-j+1}]g\Big]. \tag{76}$$

Because our goal is to compute $[\alpha^0]h_m$, we only need to compute the $[\alpha^k]h_t$'s for t from 0 to m and k from 0 to $m-t$. Since we can compute $[\alpha^k]h_t$ with $O(k)$ integer arithmetic operations, we can compute $[\alpha^0]h_m$ with $O(m^3)$ integer operations. Since there are at most 2^{m+1} summands in Eq. (66), computing $\frac{1}{\beta^{m+1}\mathbf{W}!}\mathcal{D}_{\mathbf{W}}\mathcal{L}$ requires $O(2^m(4^n nm + m^3)) = O(8^m m^3)$ integer operations. Since the integer operations in question (addition, subtraction, and multiplication) can be performed in $O(b \log b)$ time [HvdH21], where b is the length of the integer in bits, it suffices to show that, throughout this procedure, we always work with integers that are $\mathrm{poly}(m)$ bits long.

When computing coefficients of g, note that the magnitude of the integers in P^k is bounded by $\left\|P^k\right\| = O(m^k) = O(m^m)$, so the trace (and consequently, the numerator of $[\alpha^k]g$ can be represented with $O(m \log m)$ bits. The denominator is $m!2^n$, which can also be represented in $O(m \log m)$ bits.

As for h_t, we proceed by giving upper bounds on the coefficients. We claim that for any $k \le m+1$ and any $t \le m$

$$\left|[\alpha^k]h_t\right| \le 2^t(m+1)^{3t+1}e^{(t+1)(m+1)}. \tag{77}$$

The case of $t = 0$ is shown below.

$$\left|[\alpha^k]g\right| \le \frac{1}{k!}\Big(\sum_{j=1}^{n} z_j\Big)^k \le \frac{1}{k!}\Big(\sum_{j=1}^{n}\mu_j\Big)^k = \frac{1}{k!}(m+1)^k \le \sum_{s=0}^{\infty}\frac{(m+1)^s}{s!} = e^{m+1} \tag{78}$$

$$\left|[\alpha^k]h_0\right| = \left|[\alpha^k]g'\right| \le \frac{(m+1)^{k+1}}{k!} \le (m+1)e^{m+1}. \tag{79}$$

Since taking derivative brings a factor at most $(m+1)$ to the polynomial coefficient of k-th order

term where $k \leq m + 1$, the induction hypothesis implies that

$$\left|[\alpha^k]h_t\right| = \left|\sum_{j=0}^{k}\Big[(j+1)[\alpha^{j+1}]h_{t-1}[\alpha^{k-j}]g - t(k-j+1)[\alpha^j]h_{t-1}[\alpha^{k-j+1}]g\Big]\right| \tag{80}$$

$$\leq (m+1)\Big[(m+1)2^{t-1}(m+1)^{3t-2}e^{t(m+1)}e^{m+1}$$

$$+ (m+1)^2 2^{t-1}(m+1)^{3t-2}e^{t(m+1)}e^{m+1}\Big]$$

$$= 2^{t-1}(m+1)^{3t}e^{(t+1)(m+1)}(m+2)$$

$$\leq 2^t(m+1)^{3t+1}e^{(t+1)(m+1)}.$$

This implies that every coefficient of $h_t(\alpha; z)$ up to the $(m+1)$-th order is at most $\exp(O(m^2))$ in magnitude.[9] Because all coefficients of g are represented with a denominator of $(m+1)!2^n$, when computing $[\alpha^k]h_t$, the denominator is always $((m+1)!)^{t+1}2^{n(t+1)} = \exp(O(m^2 \log m))$. So, along with the magnitude bound, the numerator of these coefficients is always $\exp(O(m^2 \log m))$. Thus, we are always working with integers of $O(m^2 \log m)$ digits, as desired.

Altogether, this makes the time complexity of computing a coefficient

$$O(n^3 L + 8^m m^3(m^2 \log m(\log(m^2 \log m)))) = O(Lm^3 + 8^m m^5 \log^2 m). \tag{81}$$

$\square$

We summarize the algorithm in Algorithm 2.

Algorithm 2: Evaluating a cluster derivative

Data: Cluster $\mathbf{W} = \{(a, \mu(a)) : a = 1, 2, \ldots, n\}$ of total weight $\sum_a \mu(a) = m+1$ and associated Pauli operator E_a with $a \in \mathbf{W}$

Result: A rational number $\frac{1}{\beta^{m+1}\mathbf{W}!}\mathcal{D}_{\mathbf{W}}\mathcal{L}$

1 Let out $\leftarrow 0$;
2 Let $\tilde{E}_a$ be a faithful representation on $m+1$ qubits for E_a;
3 **for** $z_1 \in \{0, 1, \ldots, \mu(1)\}, \ldots, z_n \in \{0, 1, \ldots, \mu(n)\}$ **do**
4 $\quad$ Compute coefficients $[\alpha^k]g(\alpha; z)$ of Eq. (74) using $\tilde{E}_a$;
5 $\quad$ **for** $t \in \{0, 1, \ldots, m\}$ **do**
6 $\quad\quad$ Compute $[\alpha^k]h_t(\alpha; z)$ for k from 0 to $m-t$ by Eq. (76);
7 $\quad$ **end**
8 $\quad$ out $\leftarrow$ out $+(-1)^{z_1 + \cdots + z_n + m + 1}\binom{\mu(1)}{z_1} \cdots \binom{\mu(n)}{z_n}[\alpha^0]h_m(\alpha; z)$;
9 **end**
10 Return $\frac{1}{\mathbf{W}!(m+1)!}$ out;

Proof of Lemma 3.14. Without loss of generality, we assume that our derivative is taken at the origin in the domain. Define functions e_k for $k = 0, 1, 2, \ldots, n$ recursively as

$$e_0 = f, \tag{82}$$

$$e_k(x_1, \ldots, x_n) = e_{k-1}(x_1, \ldots, x_{k-1}, x_k, x_{k+1}, \ldots, x_n) - e_{k-1}(x_1, \ldots, x_{k-1}, 0, x_{k+1}, \ldots, x_n). \tag{83}$$

---

[9]If this bound were $\exp(O(m))$, we would have implied Proposition 3.8.

28

The function e_n is a sum of 2^n terms. The mean value theorem implies that for any $\alpha \neq 0$ there exist $\theta_n, \theta_{n-1}, \ldots, \theta_1 \in (0, 1)$ such that

$$
\begin{aligned}
e_n(\alpha, \alpha, \ldots, \alpha) &= \alpha(\partial_n e_{n-1})(\underbrace{\alpha, \ldots, \alpha}_{n-1}, \theta_n \alpha) \\
&= \alpha(\partial_n \alpha \partial_{n-1} e_{n-2})(\underbrace{\alpha, \ldots, \alpha}_{n-2}, \theta_{n-1}\alpha, \theta_n\alpha) \\
&= \alpha^n(\partial_n \partial_{n-1} \cdots \partial_1 e_0)(\theta_1 \alpha, \theta_2 \alpha, \ldots, \theta_n \alpha). \tag{84}
\end{aligned}
$$

This means that

$$
\left.\frac{\partial^n f}{\partial x_1 \cdots \partial x_n}\right|_{x=0} = \lim_{\alpha \to 0} \frac{e_n(\alpha, \ldots, \alpha)}{\alpha^n} = \lim_{\alpha \to 0} \frac{1}{\alpha^n} \sum_{y \in \{0,1\}^n} (-1)^{|y|+n} f(x + \alpha y), \tag{85}
$$

where $|y|$ is the sum of components of y. This completes the proof of the lemma in the case where the derivative is first order in each variable.

Higher order cases are proved by considering the composition h of f and a linear function

$$
\begin{aligned}
h: \quad &(x_{i,j} \mid j = 1, \ldots, \mu_i, \quad i = 1, 2, \ldots, n) \tag{86} \\
&\mapsto \left(x_i = \sum_j x_{i,j} \;\middle|\; i = 1, 2, \ldots, n\right) \\
&\mapsto f(x_1, \ldots, x_n)
\end{aligned}
$$

We see that $\partial_1^{\mu_1} \cdots \partial_n^{\mu_n} f = (\prod_{i=1}^n \prod_{j=1}^{\mu_i} \partial_{x_{i,j}})h$. Let $\mu = \mu_1 + \mu_2 + \cdots + \mu_n$. For any $y \in \{0,1\}^\mu$ we define $z(y) \in \mathbb{Z}_{\geq 0}^n$ to be a vector whose component $z(y)_k$ is

$$
z(y)_k = y_{\mu_1 + \cdots + \mu_{k-1}+1} + y_{\mu_1 + \cdots + \mu_{k-1}+2} + \cdots + y_{\mu_1 + \cdots + \mu_k}. \tag{87}
$$

In other words, we put the components of y into n bins of sizes μ_k and sum the numbers in each bin to make $z(y)$. Then, an arbitrary mixed derivative is expressed as

$$
\left.\frac{\partial^\mu f}{\partial x_1^{\mu_1} \cdots \partial x_n^{\mu_n}}\right|_x = \lim_{\alpha \to 0} \frac{1}{\alpha^\mu} \sum_{y \in \{0,1\}^\mu} (-1)^{|y|+\mu} f(x + \alpha z(y)). \tag{88}
$$

Expressing the summation over y as a summation over z_k, we complete the proof. $\square$

3.7 Proof of Theorem 3.1

To be clear where we are evaluating derivatives, in this proof we let $\mathcal{L} = \mathcal{L}(\lambda)$ be the function of variables λ_a. Let ∂_a denote the derivative $\partial/\partial\lambda_a$ at $\lambda = \xi$. Proposition 3.2 says that the expectation value $\mathrm{Tr}(E_a e^{-\beta H})/\mathrm{Tr}\, e^{-\beta H}$ of E_a is given by the first derivative of the logarithmic partition function:

$$
\left.\frac{\mathrm{Tr}(E_a e^{-\beta H})}{\mathrm{Tr}\, e^{-\beta H}}\right|_{\lambda=\xi} = -\frac{1}{\beta}\partial_a \log \mathrm{Tr} \exp(-\beta H)|_{\lambda=\xi} \tag{89}
$$

$$
= -\frac{1}{\beta}\partial_a|_{\lambda=\xi} \sum_{m=0}^{\infty} \sum_{\mathbf{V}:|\mathbf{V}|=m} \frac{\lambda^{\mathbf{V}}}{\mathbf{V}!}\mathcal{D}_{\mathbf{V}}\mathcal{L} \qquad \text{by Eq. (25)} \tag{90}
$$

29

Recall that $\mathcal{D}_\mathbf{V}\mathcal{L}$ is a constant in λ, since $\mathcal{D}_\mathbf{V}$ is a derivative evaluated at $\lambda = 0$ (Eq. (26)). This multivariate Taylor series is in fact a (disguised) power series in β. Since all the functions here are complex differentiable on an open set that contains $\mathbb{R}$ in the β-complex plane, Lemma 2.5 implies that the equality holds whenever the series is absolutely convergent and $\beta \in \mathbb{R}$, in which case the infinite sum over m can be interchanged with ∂_a. The term with $m = 0$ is eliminated by ∂_a, so we shift the dummy variable m by one. Hence,

$$= -\frac{1}{\beta} \sum_{m=0}^{\infty} \sum_{\mathbf{V}:|\mathbf{V}|=m+1} \frac{\partial_a \lambda^{\mathbf{V}}|_{\lambda=\xi}}{\mathbf{V}!} \mathcal{D}_\mathbf{V}\mathcal{L} \tag{91}$$

$$= -\frac{1}{\beta} \sum_{m=0}^{\infty} \sum_{\substack{\mathbf{V}:|\mathbf{V}|=m+1, \\ a \in \mathbf{V}}} \frac{\partial_a \lambda^{\mathbf{V}}|_{\lambda=\xi}}{\mathbf{V}!} \mathcal{D}_\mathbf{V}\mathcal{L}. \qquad \text{since } \partial_a \lambda^{\mathbf{V}} = 0 \text{ if } a \notin \mathbf{V} \tag{92}$$

For a cluster $\mathbf{V}$ of total weight $m + 1$, the expression $\mathcal{D}_\mathbf{V}\mathcal{L}$ has a factor of β^{m+1}. The overall $\frac{1}{\beta}$ factor reduces the exponent of β by one, so, if we group summands by degree of β, we have

$$\left.\frac{\mathrm{Tr}(E_a e^{-\beta H})}{\mathrm{Tr}\, e^{-\beta H}}\right|_{\lambda=\xi} = p_0 + \beta p_1 + \beta^2 p_2 + \cdots, \tag{93}$$

$$p_m = (-1)^{m+1} \sum_{\substack{\mathbf{V}:|\mathbf{V}|=m+1, \\ a \in \mathbf{V}}} (\partial_a \lambda^{\mathbf{V}}|_{\lambda=\xi}) \frac{\mathcal{D}_\mathbf{V}\mathcal{L}}{\beta^{m+1}\mathbf{V}!}. \tag{94}$$

Note that p_0 is proportional to Eq. (89) with $\lambda = 0$. Since $\mathrm{Tr}\, E_a = 0$, we have $p_0 = 0$. This proves Item 3.1(1) that p_m is a homogeneous polynomial in λ_a of total degree m. Proposition 3.5 says that $\mathbf{V}$ has to be connected and includes a, implying that $\mathbf{V}$ of total weight $m + 1$ includes nodes within $\mathfrak{G}$-distance m from a. This implies Item 3.1(2). Proposition 3.6 bounds the number of clusters to be summed over. This is Item 3.1(3). Proposition 3.8 bounds the magnitude of $\mathcal{D}_\mathbf{V}\mathcal{L}/\mathbf{V}!$. The derivative ∂_a may put an additional factor at most $m + 1$. This proves Item 3.1(4).

We have considered algorithms to enumerate clusters, proving Item 3.1(A). Proposition 3.13 shows Item 3.1(B). This completes the proof of Theorem 3.1.

4 Learning algorithm

In this section, we describe our algorithm for learning the coefficients of a Hamiltonian given copies of its Gibbs state. This section relies on the results of the previous section only through Theorem 3.1.

Unlike Section 3, we only consider Hamiltonians $\{(a, E_a, \lambda_a) : a = 1, 2, \ldots, M\}$ (Definition 2.1) where the E_a's are distinct non-identity tensor products of Pauli matrices, so that they are orthonormal with respect to the normalized Hilbert–Schmidt inner product. That is,

$$\forall a, b \in [M] : \quad \mathrm{Tr}(E_a E_b) = \mathsf{D}\delta_{ab}, \tag{95}$$

where δ_{ab} is the Kronecker delta function.

Our overall strategy for the learning algorithm can be broken down into the following two steps. Let $\rho(\lambda)$ be the Gibbs state with coefficients λ.

1. Find estimates $\hat{E}_a$ for all of the expectation values $\langle E_a \rangle(\lambda) = \mathrm{Tr}(E_a \rho(\lambda))$ that satisfy $\left|\hat{E}_a - \langle E_a \rangle(\lambda)\right| \leq \beta\varepsilon$ for all $a \in [M]$.

2. Then (approximately) invert the function $x \mapsto \langle E_a \rangle(x)$ on these estimates to find an estimate of the coefficients $\hat{x}$.

Step 1 of this plan is the easier step and not too hard to establish.

Lemma 4.1. *Consider a Hamiltonian $\{(a, E_a, \lambda_a) : a \in [M]\}$ on N qubits. We can find estimates $\hat{E}_a$ such that $\left|\hat{E}_a - \langle E_a \rangle(\lambda)\right| \leq \beta\varepsilon$ for all $a \in [M]$, with probability at least $1 - \delta$, using only $O(\frac{\mathfrak{d}}{\beta^2\varepsilon^2} \log(\frac{M}{\delta}))$ copies of the Gibbs state and with time complexity $O(\frac{N\mathfrak{d}}{\beta^2\varepsilon^2} \log(\frac{M}{\delta}))$.*

Proof. Recall the problem of estimating $\mathrm{Tr}(E\rho)$ for $\|E\| \leq 1$ and a quantum state ρ. If we want to estimate this to accuracy ε with success probability at least $1 - \delta$, it is a standard result that this can be done with $O(\log(1/\delta)/\varepsilon^2)$ copies of ρ. Indeed, we measure ρ in the eigenbasis of E and output the corresponding eigenvalue of E on getting that outcome. This is a random variable with expected value $\mathrm{Tr}(E\rho)$. Since $\|E\| \leq 1$, this is a random variable in $[-1, 1]$. Hence by the Chernoff bound we can estimate it to additive error ε with probability at least $1 - \delta$ using $O(\log(1/\delta)/\varepsilon^2)$ copies of ρ.

Now we want to measure all the observables E_a. But not all of these have overlapping support, and we can measure a large number of them simultaneously. Imagine we color the vertices of $\mathfrak{G}$ using $\mathfrak{d} + 1$ colors such that no neighboring pair of nodes have the same color; a greedy coloring algorithm can be used. By definition of the dual interaction graph, all the E_a's of a particular color act on separate qubits. So we can estimate all of the E_a's of a particular color using only $O(\log(1/\delta')/\varepsilon'^2)$ Gibbs state where δ' is the probability that one of estimates has error larger than $\varepsilon' = \varepsilon\beta$.

Since E_a is a Pauli operator (a tensor product of single-qubit Paulis), it suffices to measure individual qubits in some Pauli basis and multiply them (each of which is ± 1) to infer the eigenvalue of E_a. Hence, for a particular color, the time complexity is $O(N \log(1/\delta')/\varepsilon'^2)$. We repeat this for each color, resulting in $\mathfrak{d} + 1$ rounds.

Since we want all M estimates to be correct with probability at least $1 - \delta$, it suffices to set $\delta' = \delta/M$ to apply the union bound. $\square$

The remainder of this section is devoted to implementing Step 2 of the above plan. We start by upper bounding the sample complexity, and then move on to bounding the time complexity of our algorithm.

4.1 Definitions and a sample complexity upper bound

Recall that Theorem 3.1 implies that we can expand $\langle E_a \rangle$ into a Taylor series

$$\langle E_a \rangle(x) = \beta p_1^{(a)}(x) + \beta^2 p_2^{(a)}(x) + \beta^3 p_3^{(a)}(x) + \cdots, \tag{96}$$

where the sum of the absolute values of the coefficients of p_m is bounded by a universal constant that depends only on $\mathfrak{d}$ and m. We call this constant $c_m \in \mathbb{R}_{>0}$, and from Items 3.1(3) and 3.1(4) we have

$$\begin{aligned} c_m &= e\mathfrak{d}(1 + e(\mathfrak{d} - 1))^m (2e(\mathfrak{d} + 1))^{m+1}(m + 1) \\ &= 2e^2\mathfrak{d}(\mathfrak{d} + 1)\tau^m(m + 1), \end{aligned} \tag{97}$$

where

$$\tau = (1 + e(\mathfrak{d} - 1))(2e(\mathfrak{d} + 1)) \leq 2e^2(\mathfrak{d} + 1)^2. \tag{98}$$

Further, $p_k^{(a)}(x)$ only depends on the entries of x whose operators are within $\mathfrak{G}$-distance k from a. The first term $p_1^{(a)}$ can be determined more explicitly by

$$p_1^{(a)}(x) = \frac{\partial}{\partial\beta} \frac{\mathrm{Tr}(E_a \exp(-\beta H))}{\mathrm{Tr}\exp(-\beta H)}\bigg|_{\beta=0,\lambda=x} = \frac{1}{\mathsf{D}} \mathrm{Tr}(E_a(-H))|_{\lambda=x} = -x_a, \tag{99}$$

where we used Eq. (95) in the last equality.

Let $\mathcal{F} : [-1, 1]^M \to \mathbb{R}^M$ be $\langle E_a \rangle(x)$, truncated to order $\mathfrak{m}$ terms ($\mathfrak{m} \geq 1$) and shifted by our known estimates $\hat{E}_a$ of $\langle E_a \rangle(\lambda)$ from Lemma 4.1, which satisfy $|\hat{E}_a - \langle E_a \rangle(\lambda)| \leq \beta\varepsilon$. Thus we have

$$\mathcal{F}_a(x) := \mathcal{F}_a(x_1, \ldots, x_M) = \sum_{k=0}^{\mathfrak{m}} \beta^k p_k^{(a)}(x) = -\hat{E}_a - \beta x_a + \beta^2 p_2^{(a)}(x) + \cdots + \beta^{\mathfrak{m}} p_{\mathfrak{m}}^{(a)}(x), \quad (100)$$

where we defined $p_0^{(a)} = -\hat{E}_a$. Our goal is to find an x such that $\mathcal{F}(x)$ is small, since, as we argue below, such an x will be close to the true coefficient vector λ.

As a warmup for the time complexity upper bound proved in the next section, we will show a sample complexity upper bound. The fundamental idea in both upper bounds is the same: Find an x such that $\|\mathcal{F}(x)\|_\infty = O(\beta\varepsilon)$.

Theorem 4.2. *Consider a Hamiltonian $\{(a, E_a, \lambda_a) : a \in [M]\}$ such that E_a are traceless and orthonormal with respect to the Hilbert-Schmidt inner product. Then, for any β such that*

$$100e^6(\mathfrak{d} + 1)^8 \beta \leq 1, \quad (101)$$

we can find $x \in [-1, 1]^M$, such that $\|x - \lambda\|_\infty \leq \varepsilon$ with probability $\geq 1 - \delta$ using only

$$O\left(\frac{\mathfrak{d}}{\beta^2 \varepsilon^2} \log \frac{M}{\delta}\right) \quad (102)$$

copies of the Gibbs state.

Proof. From Lemma 4.1 we know that $O(\frac{\mathfrak{d}}{\beta^2 \varepsilon^2} \log(\frac{M}{\delta}))$ Gibbs states suffice to estimate $\langle E_a \rangle$ to $\beta\varepsilon$ accuracy for all a with probability $\geq 1 - \delta$.

Next, consider $\mathcal{F}$ for $\mathfrak{m} = \infty$, so $\mathcal{F}_a(x) = \langle E_a \rangle(x) - \hat{E}_a$. Notice that this means that $\|\mathcal{F}(\lambda)\|_\infty \leq \beta\varepsilon$, by our assumption about the accuracy of the estimates $\hat{E}_a$. Our algorithm will be to find and output any $x \in [-1, 1]^M$ satisfying $\|\mathcal{F}(x)\|_\infty \leq \beta\varepsilon$. We know one such x must exist, since λ satisfies this equation. It remains to be shown that any such x is also close to λ.

Let ∂_b denote the derivative with respect to x_b and let $J = \mathrm{d}\mathcal{F}$ be the Jacobian of $\mathcal{F}$, so $J_{ab} := \partial_b \mathcal{F}_a$. Then, for each a, by the multivariate mean value theorem, there exists $y(a) \in (-1, 1)^M$ such that

$$\mathcal{F}_a(x) = \mathcal{F}_a(\lambda) + (J|_{y(a)}(x - \lambda))_a. \quad (103)$$

This implies that

$$|x_a - \lambda_a| = |\sum_b (J|_{y(a)}^{-1})_{ab}(\mathcal{F}_b(x) - \mathcal{F}_b(\lambda))|$$

$$\leq \|J|_{y(a)}^{-1}\|_{\infty \to \infty}(\|\mathcal{F}(x)\|_\infty + \|\mathcal{F}(\lambda)\|_\infty) \leq (2\beta^{-1})(2\beta\varepsilon) = 4\varepsilon, \quad (104)$$

where the final inequality uses Lemma 4.3 below, which holds when β is bounded as in Eq. (101). Rescaling $\varepsilon \to \frac{1}{4}\varepsilon$ completes the proof. $\qquad \square$

Lemma 4.3. *For Hamiltonians as in Theorem 4.2, if Eq. (101) holds, then for any $x \in [-1, 1]^M$, we have $\|I + \beta^{-1} J(x)\|_{\infty \to \infty} \leq \frac{1}{2}$ and $\|J(x)^{-1}\|_{\infty \to \infty} \leq 2\beta^{-1}$ for any $\mathfrak{m} \geq 1$.*

In particular, the lemma is true when $\mathfrak{m} = \infty$ and consequently J is also the Jacobian of the function $\mathbb{R}^M \ni x \mapsto (\langle E_a \rangle(x)) \in \mathbb{R}^M$. The proof implicitly uses a band-diagonal property of J: if b and a are distance k apart, then J_{ab} scales as β^{k+1}.

32

Proof. In this proof we suppress the argument x in $J(x)$. If $\|I + \beta^{-1}J\|_{\infty\to\infty} \leq \frac{1}{2}$, then

$$J^{-1} = -\frac{1}{\beta}\frac{I}{I - (I + \beta^{-1}J)} = -\frac{1}{\beta}\sum_{k=0}^{\infty}(I + \beta^{-1}J)^k \tag{105}$$

$$\|J^{-1}\|_{\infty\to\infty} \leq \beta^{-1}\sum_{k=0}^{\infty}\|I + \beta^{-1}J\|_{\infty\to\infty}^k \leq 2\beta^{-1}. \tag{106}$$

Hence, we have to show that $\|\beta I + J\|_{\infty\to\infty} \leq \frac{\beta}{2}$ in the stated range of β to complete the proof. The leading order term of J is $-\beta I$,

$$J_{ab} = \partial_b \mathcal{F}_a = -\beta\delta_{ab} + O(\beta^2); \tag{107}$$

we will bound the rest of J to show that J is close to $-\beta I$. Let $u = (u_1, u_2, \ldots, u_M)$ be such that $|u_b| \leq 1$ for all b.

$$\begin{aligned}((J + \beta I)u)_a &= \sum_b (J + \beta I)_{ab}u_b \\ &= \sum_b u_b\left(\beta^2\partial_b p_2^{(a)}(x) + \cdots + \beta^{\mathfrak{m}}\partial_b p_{\mathfrak{m}}^{(a)}(x)\right) \\ &= \sum_{k=2}^{\mathfrak{m}}\beta^k\sum_{b:\text{dist}(a,b)\leq k}u_b\partial_b p_k^{(a)}(x) \qquad\qquad \text{by Item 3.1(2).} \end{aligned} \tag{108}$$

For each k in the last sum, the index b ranges over at most $1 + \mathfrak{d} + \cdots + \mathfrak{d}^k \leq (\mathfrak{d} + 1)^k$ nodes of $\mathfrak{G}$. Further, Item 3.1(1) says that $p_k^{(a)}$ is a homogeneous polynomial of degree k and the sum of the absolute value of its coefficients is bounded by c_k of Eq. (97). As a result, $|\partial_b p_k^{(a)}| \leq kc_k$ everywhere in the domain of $\mathcal{F}$.

$$|((J + \beta I)u)_a| \leq \sum_{k=2}^{\infty}\beta^k \cdot (\mathfrak{d} + 1)^k \cdot kc_k \tag{109}$$

$$\leq 2e^2(\mathfrak{d} + 1)^2(\beta(\mathfrak{d} + 1)\tau)^2\sum_{k=2}^{\infty}(\beta(\mathfrak{d} + 1)\tau)^{k-2} \cdot k(k+1) \tag{110}$$

$$= 2e^2(\mathfrak{d} + 1)^4\beta^2\tau^2\left(\frac{6 - 6r + 2r^2}{(1-r)^3}\Big|_{r=\beta(\mathfrak{d}+1)\tau}\right) \qquad\qquad \text{if } \beta(\mathfrak{d} + 1)\tau < 1$$

$$\leq 2e^2(\mathfrak{d} + 1)^4\beta^2\tau^2 \cdot \frac{25}{4} \qquad\qquad\qquad\qquad \text{if } \beta(\mathfrak{d} + 1)\tau \leq \frac{1}{100}. \tag{111}$$

Since $u \in [-1, 1]^M$ is arbitrary, the last quantity is an upper bound on $\|J + \beta I\|_{\infty\to\infty}$. The bound on τ, Eq. (98), and the bound on β, Eq. (101), together imply that it is $\leq \frac{\beta}{2}$. $\qquad\square$

With this analysis, we can also deduce a bound on the strong convexity of the log-partition function, as analyzed by [AAKS21], that is optimal up to constants. This is simply a matter of bounding J^{-1} in the usual operator norm, $\|\cdot\|_{2\to2}$, rather than the (in this case, larger) $\|\cdot\|_{\infty\to\infty}$ norm.

Corollary 4.4. *For Hamiltonians as in Theorem 4.2, if Eq. (101) holds, then $\mathcal{L}$ is $(\frac{\beta^2}{2})$-strongly convex, i.e., $\nabla^{\otimes 2}\mathcal{L} - \frac{\beta^2}{2}I$ (whose (a,b)-component is $\partial_a\partial_b\mathcal{L} - \beta^2\delta_{ab}/2$) is positive semidefinite. The strong convexity constant is only a constant factor off from optimal.*

Proof. In this proof we suppress the argument x in $J(x)$, and similar arguments. By Proposition 3.2, $-\beta J$ is the Hessian of $\mathcal{L}$, taking $\mathfrak{m} = \infty$. Since it comes from a Hessian, J is Hermitian. So, it suffices to show that $\|I + \beta^{-1}J\| \leq \frac{1}{2}$, since

$$\|I + \beta^{-1}J\| \leq \frac{1}{2} \implies I + \beta^{-1}J \preceq I/2 \implies \beta^{-1}J \preceq -I/2 \implies \nabla^{\otimes 2}\mathcal{L} \succeq \beta^2 I/2, \tag{112}$$

$$\|I + \beta^{-1}J\| \leq \frac{1}{2} \implies -I - \beta^{-1}J \preceq I/2 \implies -\beta^{-1}J \preceq 3I/2 \implies \nabla^{\otimes 2}\mathcal{L} \preceq \beta^2 3I/2. \tag{113}$$

The second equation above proves optimality, up to a factor of 3. The bound we need follows immediately from Lemma 4.3, since for a Hermitian matrix X, it holds that $\|X\| \leq \|X\|_{\infty \to \infty}$. (For an eigenvector v achieving $Xv = \mu v$ with $|\mu| = \|X\|$, we see $\|X\|_{\infty \to \infty} \geq \|Xv\|_\infty / \|v\|_\infty = \|X\|$.) So,

$$\|I + \beta^{-1}J\| \leq \|I + \beta^{-1}J\|_{\infty \to \infty} \leq \frac{1}{2}, \tag{114}$$

as desired. $\qquad\square$

Remark 4.5. In this remark, we show how to tweak the result in [AAKS21] to get a slightly improved version shown in Eq. (2). We assume knowledge of [AAKS21]. First, if we do not perform the final bound in [AAKS21, Proof of Theorem 28, p.28], we have that $v^\dagger \nabla^{\otimes 2}\mathcal{L}v \geq C\|v\|_\infty^2$ for $C = e^{-O(\beta^c)}\beta^{c'}$. Using Proposition 3.2, we have that $-\beta(v^\dagger J v) \geq C\|v\|_\infty^2$. Consider taking the v that achieves $\|J^{-1}v\|_\infty = \|J^{-1}\|_{\infty \to \infty}\|v\|_\infty$. Then, using that $\|X\|_{2 \to 2} \leq \|X\|_{\infty \to \infty}$ for Hermitian X,

$$C\|J^{-1}\|_{\infty \to \infty}^2 \|v\|_\infty^2 = C\|(-J)^{-1}v\|_\infty^2 \leq \beta(v^\dagger J^{-1}v) \leq \beta\|J^{-1}\|\|v\|_2^2 \leq \beta\|J^{-1}\|_{\infty \to \infty}M\|v\|_\infty^2. \tag{115}$$

So, $\|J^{-1}\|_{\infty \to \infty} \leq \frac{\beta M}{C}$. This can be plugged in directly into, say, Eq. (104) to see that, using this bound, we would need to estimate the marginals to $\varepsilon \frac{C}{\beta M}$ error, giving the bound. Note that the assumption that $\beta = O(1)$ is not needed to achieve this sample complexity bound.

4.2 Time complexity and analysis of the Newton–Raphson method

The goal of this section is to prove the following theorem, which when combined with Lemma 4.1 to get the assumed estimates, gives us the main result (Theorem 1.1).

Theorem 4.6. *Consider a Hamiltonian $\{(a, E_a, \lambda_a) : a \in [M]\}$ such that E_a are traceless and orthonormal with respect to the Hilbert-Schmidt inner product. Suppose $\beta > 0$ satisfies*

$$25e^6(\mathfrak{d} + 1)^{10}\beta \leq 1. \tag{116}$$

Suppose we know estimates $\hat{E} \in [-1, 1]^M$ such that $\left|\hat{E}_a - \langle E_a \rangle\right| \leq \beta\varepsilon$ for all $a \in [M]$. Then we can find an x such that $\|x - \lambda\|_\infty \leq 18\varepsilon$ in time $O\left(\frac{ML}{\varepsilon} \operatorname{poly}(\mathfrak{d}, \log \frac{1}{\beta\varepsilon})\right)$.

Recall that we defined L as the maximum number of qubits that a Hamiltonian term acts on in Theorem 3.1. If L and $\mathfrak{d}$ are constant (as in our definition of a low-intersection Hamiltonian), then our time complexity has linear dependence in M, which is optimal since our output consists of M numbers. In addition, our ε-dependence is better than the ε^{-2} dependence in the sample complexity. There is very mild β-dependence since $\mathcal{F}$ becomes simpler for smaller β. The rest of this section constitutes the proof of this theorem.

Algorithm 3: The Newton–Raphson method

Data: β satisfying Eq. (116), and estimates $\{\hat{E}_a\}_{a\in[M]}$ such that $\left|\hat{E}_a - \langle E_a \rangle\right| \leq \beta\varepsilon$ for all $a \in [M]$ with $\varepsilon < \frac{1}{18}$.

Result: Estimates $\hat{\lambda} \in [-1, 1]^M$ such that $\left|\hat{\lambda}_a - \lambda_a\right| \leq 18\varepsilon$ for all $a \in [M]$

1 Define $T = \Theta(\log(\frac{1}{\beta\varepsilon\mathfrak{d}}))$ (see Eq. (134) for a precise expression) and $K = \lceil\log(\frac{3}{\beta\varepsilon})\rceil$;

2 Initialize $x^{(0)} = \vec{0}$ for $x^{(0)} \in \mathbb{R}^M$;

3 Compute all of the coefficients in the polynomials $\mathcal{F}_a(\cdot)$ for all $a \in [M]$ via Algorithm 2;

4 **for** $t = 0, 1, 2, \ldots, T - 1$ **do**

5 $\quad$ Compute $\mathcal{F}(x^{(t)})$ and (the nonzero entries of) $J(x^{(t)})$;

6 $\quad$ Compute $x^{(t+1)}$ by Eq. (118),

$$x^{(t+1)} = \text{Proj}_{[-1,1]^M}\left[x^{(t)} + \beta^{-1}\sum_{k=0}^{K-1}(I + \beta^{-1}J(x^{(t)}))^k\mathcal{F}(x^{(t)})\right];$$

7 **end**

8 Return $\hat{\lambda} \leftarrow x^{(T)}$;

From this point on, we will fix the point where we truncate $\mathcal{F}$ to be a particular value

$$\mathfrak{m} = \left\lceil \frac{e}{e-1}\frac{1}{\ln\frac{1}{\beta\tau}}\ln\left(\frac{12e^2(\mathfrak{d}+1)^2}{\beta\varepsilon\ln\frac{1}{\beta\tau}}\right)\right\rceil, \tag{117}$$

a choice that is explained in Eq. (125).

To perform the task in the theorem statement, we use Algorithm 3. Our analysis only applies when $\varepsilon \leq \frac{1}{12}$, but when $\varepsilon \geq \frac{1}{18}$, we can simply output $\hat{\lambda} = \vec{0}$ as a sufficient approximation. As in the previous section, the main idea is to find an $x \in [-1, 1]^M$ such that $\|\mathcal{F}(x)\|_\infty = O(\beta\varepsilon)$. We will do this with a version of the *Newton–Raphson method*. Typically, the Newton–Raphson method performs the iteration $x^{(t+1)} = x^{(t)} - (J^{-1}\mathcal{F})(x^{(t)})$ until convergence. However, we want to avoid computing the inverse of J explicitly, so we will perform the iteration

$$x^{(0)} = \vec{0} \qquad x^{(t+1)} = \text{Proj}_{[-1,1]^M}\left[x^{(t)} + \beta^{-1}\sum_{k=0}^{K-1}(I + \beta^{-1}J(x^{(t)}))^k\mathcal{F}(x^{(t)})\right]. \tag{118}$$

This uses the Taylor series approximation for J^{-1} from Eq. (105). We also perform a projection to remain inside our parameter space $[-1, 1]^M$, where $\text{Proj}_{[-1,1]^M}$ is the coordinate-wise application of

$$\text{Proj}_{[-1,1]}(u) = \begin{cases} 1 & \text{if } u \in (1, \infty) \\ u & \text{if } u \in [-1, 1] \\ -1 & \text{if } u \in (-\infty, -1) \end{cases}. \tag{119}$$

In Algorithm 3 it might seem counterintuitive that $T = \Theta(\log(\frac{1}{\beta\varepsilon\mathfrak{d}}))$ decreases as $\mathfrak{d}$ increases when β and ε are held constant; however, due to Eq. (116) our algorithm is not guaranteed to work for arbitrarily large $\mathfrak{d}$ with β and ε fixed.

Time complexity. First, we will show that Algorithm 3 has the time complexity claimed in Theorem 4.6. There are several parameters that appear in the algorithm, and it will be helpful

to upper bound them with simpler expressions now. Note that T, the number of iterations of Newton–Raphson, and K, the number of terms used in the approximation of the inverse of J, are both clearly $O(\ln(\frac{1}{\beta\varepsilon}))$. The other parameter, which is implicit in the definition of $\mathcal{F}$ is $\mathfrak{m}$, which is also $O(\ln(\frac{1}{\beta\varepsilon}))$ due to Eq. (121).

Now let us bound the time complexity of the algorithm line by line. The first line of the algorithm with a nontrivial contribution to time complexity is Line 3. We need to compute all the coefficients in the polynomials representing $\mathcal{F}_a$ for $a \in [M]$ up to truncation order $\mathfrak{m}$. By Item 3.1(3), we know that each polynomial p_m has at most $e\mathfrak{d}(1 + e(\mathfrak{d} - 1))^m$ monomials, and hence the total number of monomials in $\mathcal{F}_a$ is at most C, where $C \leq \sum_{m=1}^{\mathfrak{m}} e\mathfrak{d}(1 + e(\mathfrak{d} - 1))^m \leq (e\mathfrak{d})^{\mathfrak{m}+1}$. By Item 3.1(A), we can enumerate these coefficients in time $O(\mathfrak{d}MC)$. Then by Item 3.1(B), each coefficient can be computed exactly in time $D = (8^{\mathfrak{m}} + L)\operatorname{poly}(\mathfrak{m})$. Finally there are M different $\mathcal{F}_a$ to be computed, and hence we can write down all of $\mathcal{F}$ in $O(\mathfrak{d}\mathfrak{m}CMD)$ time.

Then in Line 5, we can perform evaluations of $\mathcal{F}(x)$ in $O(CM\mathfrak{m})$ time, since there are C monomials in each $\mathcal{F}_a$, and each has up to $\mathfrak{m}$ variables. Now recall that $J_{ab}(x) = \partial_b\mathcal{F}_a(x)$ is a sparse matrix with at most $\mathfrak{d}^{\mathfrak{m}}$ nonzero entries per row or column due to Item 3.1(2). We start by setting all these entries to 0. Then we fill out the nonzero entries of column b of the matrix by enumerating the monomials of $\mathcal{F}_a$, and for those monomials that contain x_b (and hence will contribute to $J_{ab}(x)$), adding the contribution due to this monomial to the memory location for $J_{ab}(x)$. For a given $b \in [M]$, this takes time $O(C\mathfrak{m})$, and so we can compute $J(x)$ in $O(CM\mathfrak{m})$ time.

Finally, in Line 6, we need to compute the power $(I + \frac{1}{\beta}J)^k\mathcal{F}$, which can be done by starting with $\mathcal{F}$ and multiplying by $I + \frac{1}{\beta}J$ k times, where each matrix–vector product takes time linear in the number of nonzero entries in $I + \frac{1}{\beta}J$, which is $O(M\mathfrak{d}^{\mathfrak{m}})$. So, the total runtime is

$$O\left(\underbrace{\mathfrak{d}\mathfrak{m}CMD}_{\text{Line 3}} + T\left(\underbrace{CM\mathfrak{m}}_{\text{Line 5}} + \underbrace{KM\mathfrak{d}^{\mathfrak{m}}}_{\text{Line 6}}\right)\right) = O(M\mathfrak{d}^2(e\mathfrak{d})^{\mathfrak{m}}(8^{\mathfrak{m}} + L)\operatorname{polylog}(\tfrac{1}{\beta\varepsilon})). \qquad (120)$$

Let us examine $\mathfrak{m}$ more carefully. Let $d := \mathfrak{d} + 1$. There are two asymptotically small parameters ε and β, and one large parameter d. The inverse temperature β is at most $\beta_c = (25e^6d^{10})^{-1}$ by Eq. (116). Pulling $\mathfrak{m}$ from Eq. (117) (and recalling that $\tau = (1 + e(d - 2))2ed$ from Eq. (98)), we have

$$\mathfrak{m} - 1 = \left\lceil \frac{e}{e-1}\left(\frac{\ln(\frac{12e^2d^2}{\beta\varepsilon}\ln(\frac{1}{\beta\tau}))}{\ln(\frac{1}{\beta\tau})}\right)\right\rceil - 1$$

$$\leq \frac{e}{e-1}\left(\frac{\ln(\frac{12e^2d^2}{\beta\varepsilon}\ln(\frac{1}{\beta\tau}))}{\ln(\frac{1}{\beta\tau})}\right)$$

$$= \frac{e}{e-1}\left(1 + \frac{\ln(d^2\tau) + \ln(\frac{1}{\varepsilon})}{\ln(\frac{1}{\beta\tau})} + \frac{\ln(12e^2\ln(\frac{1}{\beta\tau}))}{\ln(\frac{1}{\beta\tau})}\right)$$

$$\leq \frac{e}{e-1}\left(1 + \frac{\ln(d^2\tau) + \ln(\frac{1}{\varepsilon})}{\ln(\frac{1}{\beta_c\tau})} + \frac{\ln(12e^2\ln(\frac{1}{\beta_c\tau}))}{\ln(\frac{1}{\beta_c\tau})}\right)$$

$$= \frac{e}{e-1}\left(\frac{\ln(\frac{d^2}{\beta_c})}{\ln(\frac{1}{\beta_c\tau})} + \frac{\ln(\frac{1}{\varepsilon})}{\ln(\frac{1}{\beta_c\tau})}\right) + O\left(\frac{\ln\ln d}{\ln d}\right). \qquad (121)$$

So far, we have refrained from bounding the leading-order term (apart from taking $\beta \leq \beta_c$). We do

36

this now to bound the runtime. We use that $\tau \leq 2e^2 d^2$ by Eq. (98), so $\frac{1}{\beta\tau} \geq \frac{1}{\beta_c\tau} \geq 12e^4 d^8$.

$$
\begin{aligned}
(8ed)^{\mathfrak{m}} &= \exp\left(\ln(8ed)\left(1 + \frac{e}{e-1}\left(\frac{\ln(\frac{d^2}{\beta_c})}{\ln(\frac{1}{\beta_c\tau})} + \frac{\ln(\frac{1}{\varepsilon})}{\ln(\frac{1}{\beta_c\tau})}\right)\right) + O(\ln\ln d)\right) \\
&\leq \exp\left(\ln(8ed)\left(1 + \frac{e}{e-1}\left(\frac{\ln(25e^6 d^{12})}{\ln(12e^4 d^8)} + \frac{\ln(\frac{1}{\varepsilon})}{\ln(12e^4 d^8)}\right)\right) + O(\ln\ln d)\right) \\
&\leq \exp\left(\ln(8ed)\left(1 + \frac{e}{e-1}\frac{3}{2}\right) + \frac{e\ln(8ed)}{(e-1)\ln(12e^4 d^8)}\ln\frac{1}{\varepsilon} + O(\ln\ln d)\right) \quad (122)
\end{aligned}
$$

One can verify that $1 + \frac{e}{e-1}\frac{3}{2} < 3.5$ and $\frac{e}{e-1}\frac{\ln(8ed)}{\ln(12e^4 d^8)} \leq \frac{e\ln(16e)}{(e-1)\ln(3072e^5)} < 0.5$ for all $d \geq 2$. Hence,

$$
(8ed)^{\mathfrak{m}} = (8ed)^{1+\frac{3e}{2(e-1)}}e^{O(\ln\ln(d))}(1/\varepsilon)^{0.5} = O\left(d^{3.5}(1/\varepsilon)^{0.75}\right) = O\left(\frac{\mathrm{poly}(\mathfrak{d})}{\varepsilon}\right). \quad (123)
$$

This leads to an upper bound on the time complexity $O\left(\frac{ML}{\varepsilon}\mathrm{poly}(\mathfrak{d}, \log\frac{1}{\beta\varepsilon})\right)$ as promised in Theorem 4.6.

Correctness and error analysis. We begin by explaining the choice of $\mathfrak{m}$ that we stated above in Eq. (117). We want to choose a large enough $\mathfrak{m}$ so that the magnitude $|\mathcal{F}_a(\lambda)|$ will be small (say, at most $2\beta\varepsilon$). The convergence of the β-series by Theorem 3.1 implies that for all a,

$$
\begin{aligned}
|\mathcal{F}_a(\lambda)| &\leq |-\hat{E}_a + \langle E_a\rangle| + |-\langle E_a\rangle - \beta\lambda_a + \beta^2 p_2^{(a)}(\lambda) + \cdots + \beta^{\mathfrak{m}} p_{\mathfrak{m}}^{(a)}(\lambda)| \\
&\leq \beta\varepsilon + \sum_{m>\mathfrak{m}}\beta^m c_m \\
&= \beta\varepsilon + 2e^2\mathfrak{d}(\mathfrak{d}+1)\frac{(\beta\tau)^{\mathfrak{m}}}{(1-\beta\tau)^2}(\mathfrak{m}(1-\beta\tau) + 1) \\
&\leq \beta\varepsilon + 12e^2(\mathfrak{d}+1)^2(\beta\tau)^{\mathfrak{m}}\mathfrak{m} \qquad\qquad\qquad\qquad \text{if } \beta\tau \leq \frac{1}{2}. \quad (124)
\end{aligned}
$$

To obtain an $\mathfrak{m}$ such that $\|\mathcal{F}(\lambda)\|_\infty \leq 2\beta\varepsilon$, we require

$$
e^{-\mathfrak{m}'}\mathfrak{m}' \leq \frac{\beta\varepsilon}{12e^2(\mathfrak{d}+1)^2}\ln\frac{1}{\beta\tau} \text{ where } \mathfrak{m}' = \mathfrak{m}\log\frac{1}{\beta\tau}. \quad (125)
$$

Using the fact that, for $0 < b < 1$, $x = \frac{e}{e-1}\ln\frac{1}{b}$ is a solution to $xe^{-x} \leq b$, it is enough to have $\mathfrak{m}$ chosen as in Eq. (117).

Then recall that in our algorithm we wanted to apply J^{-1}, but settled for an approximation to make it more time efficient. There is a deviation incurred from this approximation of J^{-1} in each time step t:

$$
\begin{aligned}
e^{(t)} &:= \left(J(x^{(t)})^{-1} + \frac{1}{\beta}\sum_{k=0}^{K-1}(I + \beta^{-1}J(x^{(t)}))^k\right)\mathcal{F}(x^{(t)}) \\
&= -\frac{1}{\beta}\sum_{k=K}^{\infty}(I + \beta^{-1}J(x^{(t)}))^k\mathcal{F}(x^{(t)}) \qquad\qquad\qquad\qquad (126) \\
&= J^{-1}(x)(I + \beta^{-1}J(x^{(t)}))^K\mathcal{F}(x^{(t)})
\end{aligned}
$$

From Lemma 4.3, this error decays exponentially with K.

We can now begin analyzing the convergence of the Newton–Raphson method. Consider $\mathcal{F}_a(s) : [0, 1] \to \mathbb{R}$, where $\mathcal{F}_a(s) := \mathcal{F}_a(x + s(\lambda - x))$, which is coordinate a of $\mathcal{F}$ along the straight-line path between x and λ. Then using Taylor's theorem, which gives us a form for the remainder term in a Taylor series expansion, there is some $s' \in [0, 1]$ such that

$$\mathcal{F}_a(1) = \mathcal{F}_a(0) + (\partial_s \mathcal{F}_a)(0) + \frac{1}{2}(\partial_s^2 \mathcal{F}_a)(s'). \tag{127}$$

Now, we use that $\partial_s = \sum_b (\lambda_b - x_b)\partial_b$ and substitute our previous definition of $\mathcal{F}_a$ to get that, for $y^{(a)} := s'\lambda + (1 - s')x$,

$$\mathcal{F}_a(\lambda) = \mathcal{F}_a(x) + \sum_b (\lambda_b - x_b)\underbrace{(\partial_b \mathcal{F}_a)}_{J_{ab}}(x) + \frac{1}{2}\sum_{b,c}(\lambda_b - x_b)(\lambda_c - x_c)(\partial_b \partial_c \mathcal{F}_a)(y^{(a)}). \tag{128}$$

Using this, we will analyze how a Newton–Raphson method iteration decreases the distance to the solution λ. Let $x := x^{(t)}$, $x' := x^{(t+1)}$, $e := e^{(t)}$, $\Delta := x - \lambda$, and $\Delta' := x' - \lambda$.

$$|\Delta_d'| = |\mathrm{Proj}_{[-1,1]}[(x - (J^{-1}\mathcal{F})(x) + e)_d] - \lambda_d| \tag{129}$$

$$\leq \left| (x - (J^{-1}\mathcal{F})(x) + e)_d - \lambda_d \right|$$

$$= \left| e_d + \Delta_d - \sum_a (J(x)^{-1})_{da}(\mathcal{F}(x))_a \right|$$

$$= \left| e_d + \Delta_d - \sum_a J(x)_{da}^{-1}\left(\mathcal{F}_a(\lambda) - \sum_b (\lambda_b - x_b)J(x)_{ab} - \frac{1}{2}\sum_{b,c}(\lambda_b - x_b)(\lambda_c - x_c)[\partial_b \partial_c \mathcal{F}_a](y^{(a)})\right) \right|$$

$$= \left| \left[e + \Delta - J(x)^{-1}\mathcal{F}(\lambda) - J(x)^{-1}J(x)\Delta \right]_d + \frac{1}{2}\sum_{a,b,c} J(x)_{da}^{-1}\Delta_b \Delta_c [\partial_b \partial_c \mathcal{F}_a](y^{(a)}) \right|$$

$$= \left| \left[J(x)^{-1}\left((I + \beta^{-1}J(x))^K \mathcal{F}(x) - \mathcal{F}(\lambda)\right) \right]_d + \frac{1}{2}\sum_{a,b,c} J(x)_{da}^{-1}\Delta_b \Delta_c [\partial_b \partial_c \mathcal{F}_a](y^{(a)}) \right|$$

We will bound each expression above in turn. We can bound the first expression using Lemma 4.3 and Eq. (124), and that $K = \lceil \log_2(\frac{3}{\beta \varepsilon}) \rceil$:

$$\left| \left[J(x)^{-1}\left((I + \beta^{-1}J(x))^K \mathcal{F}(x) - \mathcal{F}(\lambda)\right) \right]_d \right|$$

$$\leq \|J(x)^{-1}\|_{\infty \to \infty}\left(\|I + \beta^{-1}J(x)\|_{\infty \to \infty}^K \|\mathcal{F}(x)\|_\infty + \|\mathcal{F}(\lambda)\|_\infty \right)$$

$$\leq 2\beta^{-1}\left(2^{-K}(2 + \beta\varepsilon) + 2\beta\varepsilon \right) \leq 6\varepsilon. \tag{130}$$

The second expression can be bounded through an argument similar to that of Lemma 4.3, in particular, that $\mathcal{F}_a(y)$ decomposes into degree-k polynomials $p_k^{(a)}(y)$ that depend only on y_b where b are within $\mathfrak{G}$-distance k from a and that have a bound on the magnitude of the coefficients (given

by c_k defined in Eq. (97)). We have that for all d,

$$\left|\frac{1}{2}\sum_{a,b,c}J(x)_{da}^{-1}\Delta_b\Delta_c[\partial_b\partial_c\mathcal{F}_a](y^{(a)})\right| \leq \frac{1}{2}\|J(x)^{-1}\|_{\infty\to\infty}\max_a|\sum_{b,c}\Delta_b\Delta_c[\partial_b\partial_c\mathcal{F}_a](y^{(a)})| \tag{131}$$

$$\leq \frac{1}{\beta}\max_a\sum_{k\geq 0}\sum_{b,c}|\Delta_b\Delta_c|\cdot\beta^k\cdot|\partial_b\partial_c p_k^{(a)}(y)|$$

$$\leq \frac{1}{\beta}\max_a\sum_{k\geq 0}\sum_{\substack{b,c:\\ \text{dist}(b,a)\leq k\\ \text{dist}(c,a)\leq k}}\|\Delta\|_\infty^2\cdot\beta^k\cdot k(k-1)c_k$$

$$\leq \frac{1}{\beta}\sum_{k\geq 0}(\mathfrak{d}+1)^{2k}\|\Delta\|_\infty^2\cdot\beta^k\cdot k(k-1)c_k$$

$$= \frac{12e^2}{\beta}\|\Delta\|_\infty^2(\mathfrak{d}+1)^2\frac{(\beta(\mathfrak{d}+1)^2\tau)^2}{(1-\beta(\mathfrak{d}+1)^2\tau)^4} \quad\text{if } \beta(\mathfrak{d}+1)^2\tau < 1$$

$$\leq 12.5e^2\beta(\mathfrak{d}+1)^6\tau^2\|\Delta\|_\infty^2 \quad\text{if } \beta\mathfrak{d}^2\tau \leq 1-\sqrt[4]{\frac{12}{12.5}}.$$

These two computations, together with Eq. (129), gives us our bound on $\|\Delta'\|_\infty$.

$$\|\Delta'\|_\infty \leq 6\varepsilon + 12.5e^2\beta(\mathfrak{d}+1)^6\tau^2\|\Delta\|_\infty^2 \tag{132}$$

To summarize, we have just shown that for the Newton–Raphson method iteration shown in Eq. (118), the error decays as

$$\|\Delta_{t+1}\|_\infty \leq 6\varepsilon + 12.5e^2\beta(\mathfrak{d}+1)^6\tau^2\|\Delta_t\|_\infty^2. \tag{133}$$

We can solve this recursion: By Lemma 4.7 below, provided that $75\varepsilon e^2\beta(\mathfrak{d}+1)^6\tau^2 \leq \frac{1}{4}$ and $\|\Delta_0\|_\infty \leq \frac{1}{25e^2\beta(\mathfrak{d}+1)^6\tau^2}$, we have that $\|x_t-\lambda\|_\infty \leq 18\varepsilon$ after T iterations where

$$T = \lceil -\log_2(75e^2(\mathfrak{d}+1)^6\tau^2\beta\varepsilon)\rceil$$
$$\leq \lceil -\log_2(300e^6(\mathfrak{d}+1)^{10}\beta\varepsilon)\rceil. \tag{134}$$

Since $\|\Delta_0\|_\infty \leq 1$, the condition is satisfied when $\beta \leq (25e^2(\mathfrak{d}+1)^6\tau^2)^{-1}$ and $\varepsilon \leq \frac{1}{12}$. This completes the proof of Theorem 4.6.

Lemma 4.7. *Let $c,d \in \mathbb{R}_{>0}$ be such that $cd \leq \frac{1}{4}$. Consider a sequence of positive real numbers $z_0,z_1,z_2,\ldots$ that satisfy for all $n \geq 0$,*

$$z_0 \leq \frac{1}{2d} \quad\text{and}\quad z_{n+1} \leq c+dz_n^2. \tag{135}$$

Then, for all $n \geq \log_2\frac{1}{cd}-1$ it holds that $z_n \leq 3c$.

Proof. With $y_n := dz_n$, the recursion is

$$y_0 \leq \frac{1}{2} \quad\text{and}\quad y_{n+1} \leq cd+y_n^2. \tag{136}$$

Note that by induction, $y_n \leq \frac{1}{2}$ for all $n \geq 0$. So, $\{y_n\}$ also satisfies the inequality $y_{n+1} \leq cd+\frac{1}{2}y_n$. Unrolling the iteration, we get that

$$dz_n = y_n \leq cd\Big(1+\frac{1}{2}+\cdots+\frac{1}{2^{n-1}}\Big)+\frac{1}{2^n}y_0 \leq 2cd+\frac{1}{2^{n+1}}. \tag{137}$$

So, $z_n \leq 3c$ when $n \geq \log_2\frac{1}{cd}-1$. $\qquad\square$

5 Lower bounds

In this section we establish the lower bounds claimed in Theorem 1.2, starting with the lower bound for Hamiltonian learning with ℓ_∞ error ε. We then build on that argument to obtain the lower bound with ℓ_2 error ε.

5.1 Warmup for constant N

As a warmup, let's establish a lower bound for ℓ_∞ error for Hamiltonians on a constant number of qubits. In this case we want to show a lower bound of $\Omega(\exp(2\beta)/\beta^2\varepsilon^2)$ samples for any β and $\varepsilon \in (0, 1/2]$.

Consider two diagonal Hamiltonians H_0 and H_1 for a 2-qubit system (or a single qudit with local dimension 4) expressed in terms of the Pauli matrices

$$Z \otimes I = \begin{pmatrix} +1 \\ & +1 \\ & & -1 \\ & & & -1 \end{pmatrix}, \quad I \otimes Z = \begin{pmatrix} +1 \\ & -1 \\ & & +1 \\ & & & -1 \end{pmatrix}, \quad \text{and} \quad Z \otimes Z = \begin{pmatrix} +1 \\ & -1 \\ & & -1 \\ & & & +1 \end{pmatrix}. \tag{138}$$

For any $\varepsilon \in (0, 1/2]$, we define

$$H_0 = (-1)Z \otimes I + \left(-\frac{1}{2}\right)I \otimes Z + \left(-\frac{1}{2}\right)Z \otimes Z = \begin{pmatrix} -2 \\ & 0 \\ & & 1 \\ & & & 1 \end{pmatrix}, \quad \text{and} \tag{139}$$

$$H_1 = (-1)Z \otimes I + \left(-\frac{1}{2}+\varepsilon\right)I \otimes Z + \left(-\frac{1}{2}-\varepsilon\right)Z \otimes Z = \begin{pmatrix} -2 \\ & 0 \\ & & 1+2\varepsilon \\ & & & 1-2\varepsilon \end{pmatrix}. \tag{140}$$

The coefficients of these Hamiltonians lie in $[-1, 1]$, and if we learn an unknown Hamiltonian to error $< \varepsilon/2$, then we can distinguish these two Hamiltonians. We now show that distinguishing the Gibbs states of these Hamiltonians needs $\Omega(\exp(2\beta)/\beta^2\varepsilon^2)$ samples.

Since the Hamiltonians are diagonal, their Gibbs states are also diagonal and are simply probability distributions. The two Gibbs states ρ_0 and ρ_1 are

$$\rho_0 = \frac{1}{Z_0}\begin{pmatrix} e^{2\beta} \\ & 1 \\ & & e^{-\beta} \\ & & & e^{-\beta} \end{pmatrix} \quad \text{and} \quad \rho_1 = \frac{1}{Z_1}\begin{pmatrix} e^{2\beta} \\ & 1 \\ & & e^{-\beta+2\beta\varepsilon} \\ & & & e^{-\beta-2\beta\varepsilon} \end{pmatrix}, \tag{141}$$

where Z_0 and Z_1 are the respective partition functions (or normalization constants).

We want to lower bound the number of samples needed to distinguish the two probability distributions corresponding to ρ_0 and ρ_1, which we can call q_0 and q_1. The problem of distinguishing probability distributions given samples is called hypothesis testing, and its complexity is well understood.

One way to lower bound the number of samples needed is via the KL divergence between these distributions, which is defined as follows for two distributions p and q:

$$D_{\mathrm{KL}}(p \parallel q) = \sum_j p_j \log\left(\frac{p_j}{q_j}\right). \tag{142}$$

Lemma 5.1. *For the probability distributions q_0 and q_1 corresponding to ρ_0 and ρ_1 in Eq. (141), we have $\mathrm{D}_{\mathrm{KL}}(q_1 \parallel q_0) \leq 8\beta^2\varepsilon^2 e^{-3\beta+2\beta\varepsilon}$. When $\varepsilon \leq 1/2$, we have $\mathrm{D}_{\mathrm{KL}}(q_1 \parallel q_0) \leq 8\beta^2\varepsilon^2 e^{-2\beta}$.*

Proof. This follows from a straightforward calculation.

$$
\begin{aligned}
\mathrm{D}_{\mathrm{KL}}(q_1 \parallel q_0) &= \frac{e^{2\beta}\log\frac{e^{2\beta}}{e^{2\beta}} + 1\log\frac{1}{1} + e^{-\beta+2\beta\varepsilon}\log\frac{e^{-\beta+2\beta\varepsilon}}{e^{-\beta}} + e^{-\beta-2\beta\varepsilon}\log\frac{e^{-\beta-2\beta\varepsilon}}{e^{-\beta}}}{Z_1} + \log\frac{Z_0}{Z_1} \\
&= \frac{e^{-\beta+2\beta\varepsilon}2\beta\varepsilon - e^{-\beta-2\beta\varepsilon}2\beta\varepsilon}{Z_1} + \log\frac{Z_0}{Z_1} \\
&= \frac{2\beta\varepsilon e^{-\beta+2\beta\varepsilon}(1 - e^{-4\beta\varepsilon})}{e^{2\beta} + 1 + e^{-\beta+2\beta\varepsilon} + e^{-\beta-2\beta\varepsilon}} - \log\frac{Z_1}{Z_0} \\
&= \frac{2\beta\varepsilon e^{-\beta+2\beta\varepsilon}(1 - e^{-4\beta\varepsilon})}{e^{2\beta} + 1 + e^{-\beta+2\beta\varepsilon} + e^{-\beta-2\beta\varepsilon}} - \log\frac{e^{2\beta} + 1 + e^{-\beta-2\beta\varepsilon} + e^{-\beta+2\beta\varepsilon}}{e^{2\beta} + 1 + 2e^{-\beta}}. \quad (143)
\end{aligned}
$$

Now using the inequality $\log(x) \geq 1 - 1/x = (x - 1)/x$, which holds for $x > 0$, we get

$$
\leq \frac{2\beta\varepsilon e^{-\beta+2\beta\varepsilon}(1 - e^{-4\beta\varepsilon})}{e^{2\beta} + 1 + e^{-\beta+2\beta\varepsilon} + e^{-\beta-2\beta\varepsilon}} - \frac{e^{-\beta-2\beta\varepsilon} + e^{-\beta+2\beta\varepsilon} - 2e^{-\beta}}{e^{2\beta} + 1 + e^{-\beta-2\beta\varepsilon} + e^{-\beta+2\beta\varepsilon}}. \quad (144)
$$

The denominators in this expression are $\geq e^{2\beta}$, so we can continue

$$
\begin{aligned}
&\leq e^{-2\beta}\Big(2\beta\varepsilon e^{-\beta+2\beta\varepsilon}(1 - e^{-4\beta\varepsilon}) - (e^{-\beta-2\beta\varepsilon} + e^{-\beta+2\beta\varepsilon} - 2e^{-\beta})\Big) \\
&= e^{-3\beta+2\beta\varepsilon}\Big(2\beta\varepsilon(1 - e^{-4\beta\varepsilon}) - (1 - e^{-2\beta\varepsilon})^2\Big) \\
&\leq e^{-3\beta+2\beta\varepsilon}2\beta\varepsilon(1 - e^{-4\beta\varepsilon}) \\
&= e^{-3\beta+2\beta\varepsilon}2\beta\varepsilon(1 - e^{-2\beta\varepsilon})(1 + e^{-2\beta\varepsilon}) \\
&\leq e^{-3\beta+2\beta\varepsilon}4\beta\varepsilon(1 - e^{-2\beta\varepsilon}) \leq e^{-3\beta+2\beta\varepsilon}(4\beta\varepsilon)(2\beta\varepsilon) \\
&= 8\beta^2\varepsilon^2 e^{-3\beta+2\beta\varepsilon}, \quad (145)
\end{aligned}
$$

where the last inequality used $1 - e^{-x} \leq x$, which holds for $x > 0$. $\qquad\square$

The number of samples needed to distinguish the two probability distributions is lower bounded by the inverse of the KL divergence between the two, as we make precise in the next section, which gives us the desired lower bound for constant N.

5.2 Lower bound for ℓ_∞ error

To prove the general lower bound for non-constant N, we will need Fano's lemma, and specifically we use the version in [Tsy09, Cor. 2.6]:

Lemma 5.2 (Fano's lemma). *For any $N \geq 2$, let $P_0, P_1, \ldots, P_N$ be probability distributions that satisfy*

$$
\frac{1}{N+1}\sum_{j=1}^{N}\mathrm{D}_{\mathrm{KL}}(P_j \parallel P_0) \leq \alpha \quad (146)
$$

for some $\alpha \in (0, \log N)$. Then if p_{error} denotes the minimax error of the hypothesis testing problem, or the worst-case error of distinguishing the different distributions by the best strategy, we have

$$
p_{\mathrm{error}} \geq \frac{\log(N+1) - \log(2) - \alpha}{\log(N)} \geq 1 - \frac{\log(2) + \alpha}{\log(N)}. \quad (147)
$$

We're now ready to establish a more precise version of Theorem 1.2 for ℓ_∞ error ε.

Theorem 5.3. *For any $\varepsilon \in (0, 1/2]$, $\beta > 0$, $\delta > 0$, and N, there exists a 2-local Hamiltonian on $2N$ qubits such that the sample complexity of learning its coefficients to ℓ_∞ error ε with probability at least $1 - \delta$ is $\Omega\left(\frac{\exp(2\beta)}{\beta^2 \varepsilon^2} \log\left(\frac{N}{\delta}\right)\right)$.*

Proof. We divide our $2N$ qubits into N pairs and consider Hamiltonians that are either H_0 or H_1, as defined in Section 5.1, on each pair. We consider $N + 1$ possible Hamiltonians, corresponding to all pairs having Hamiltonian H_0, or all but one pair having Hamiltonian H_0 and one pair of qubits having Hamiltonian H_1. So the potential Gibbs states produced are of the form $\rho_0 \otimes \cdots \otimes \rho_0$ or $\rho_0 \otimes \ldots \otimes \rho_0 \otimes \rho_1 \otimes \rho_0 \otimes \ldots \otimes \rho_0$ where ρ_1 is the ith copy for $i \in [N]$. As noted, learning the Hamiltonian to ℓ_∞ error $\varepsilon/2$ allows us to distinguish all these distributions. We will show that these distributions are hard to distinguish unless we have enough samples.

Consider the problem of distinguishing between $N + 1$ distributions $P_0, P_1, \ldots, P_N$, where each P_i is S independent copies of a distribution p_i over N qubits. These distributions p_i correspond to the $2N$-bit probability distributions that we get from the density matrix that has ρ_0 on all qubits and ρ_1 on the ith qubit. Since these are diagonal density matrices, we'll just think of them as probability distributions over $2SN$ bits.

To employ Fano's lemma, we need to bound $\mathrm{D}_{\mathrm{KL}}(P_j \parallel P_0)$. Since each P_i is simply S copies of a distribution p_i, we have $\mathrm{D}_{\mathrm{KL}}(P_j \parallel P_0) = S \, \mathrm{D}_{\mathrm{KL}}(p_j \parallel p_0)$ due to the chain rule for KL divergence. Any p_j with $j \neq 0$ and p_0 only differ at one site, where one distribution is q_1 and the other is q_0, so by the chain rule again we have $\mathrm{D}_{\mathrm{KL}}(p_j \parallel p_0) = \mathrm{D}_{\mathrm{KL}}(q_1 \parallel q_0)$, which we have already computed in Lemma 5.1. Thus we have that

$$\mathrm{D}_{\mathrm{KL}}(P_j \parallel P_0) = S \, \mathrm{D}_{\mathrm{KL}}(p_j \parallel p_0) = S \, \mathrm{D}_{\mathrm{KL}}(q_1 \parallel q_0) \leq 8 S \beta^2 \varepsilon^2 e^{-2\beta}. \tag{148}$$

Then we can take α to be this value and apply Fano's lemma to get

$$\begin{aligned} p_{\mathrm{error}} &\geq 1 - \frac{\log(2) + \alpha}{\log(N)} \\ &= 1 - \frac{\log(2) + 8S\beta^2\varepsilon^2 e^{-2\beta}}{\log(N)}. \end{aligned} \tag{149}$$

This error can be a small constant only if $S = \Omega\left(\frac{e^{2\beta} \log(N)}{\beta^2 \varepsilon^2}\right)$. This gives us the lower bound for constant δ.

To get the δ dependence, we show a reduction to this case. Assume there is an algorithm that can solve the Hamiltonian learning problem with error δ on the above instance on N qubits using T samples. Let's use the same algorithm to solve the hard instance we constructed above on $N/(3\delta)$ qubits with probability $2/3$. For this problem we already have a lower bound of $\Omega\left(\frac{\exp(2\beta)}{\beta^2 \varepsilon^2} \log\left(\frac{N}{\delta}\right)\right)$, which we get by replacing N by $N/(3\delta)$ in our previous lower bound.

We can split this problem up into $1/(3\delta)$ instances of size N, and apply the assumed algorithm that solves N-size instances with error δ. This algorithm needs T samples of the each of the $1/(3\delta)$ N-qubit Hamiltonians, but each sample of the $N/(3\delta)$-qubit Hamiltonian provides one sample each for the N-qubit Hamiltonians. So the sample complexity of our new algorithm remains T. Finally, this algorithm learns all $1/(3\delta)$ N-qubit Hamiltonians with error probability at most δ per instance. So by the union bound, it correctly learns all $1/(3\delta)$ instances with error at most $1/3$. Thus our assumed algorithm solves the Hamiltonian learning problem on $N/(3\delta)$ qubits and hence must use $\Omega\left(\frac{\exp(2\beta)}{\beta^2 \varepsilon^2} \log\left(\frac{N}{\delta}\right)\right)$ samples. $\square$

Remark 5.4. The lower bound above even applies to a slightly more general learning setting where we can choose different β (inverse temperature) for different samples. This scenario may arise in a physical situation where one wishes to examine the temperature dependence of some observable's expectation value to learn the Hamiltonian. In this case, the probability distributions P_j that we will distinguish is a product of p_j at possibly different temperatures. The KL divergence can be upper bounded similarly, and in the application of Fano's lemma we can take the maximum of the upper bounds on the KL divergence. If the temperatures are chosen nonadaptively, i.e., β are chosen beforehand and the samples are prepared for us accordingly, then the sample complexity is $\Omega(\min_\beta \frac{e^{2\beta}}{\beta^2 \varepsilon^2} \log N)$ for a constant probability of success, where the minimum is taken over β that are used in the samples.

5.3 Lower bound for ℓ_2 error

Our lower bound for ℓ_2 error ε' builds on the previous construction. Let us use ε' to denote the ℓ_2 error and reserve ε to be the parameter that appears in the definition of H_1 in Eq. (140).

Theorem 5.5. *For any $\beta > 0$, N, and $\varepsilon' \in (0, 0.01\sqrt{N}]$, there exists a 2-local Hamiltonian on $2N$ qubits such that the sample complexity of learning its coefficients to ℓ_2 error ε' with probability $\geq \frac{2}{3}$ is $\Omega\left(\frac{\exp(2\beta)}{\beta^2 \varepsilon'^2} N\right)$.*

Proof. To get this result for ℓ_2 error ε', we consider a different collection of Hamiltonians. Consider an error correcting code C over N bits that encodes $\Omega(N)$ logical bits and has code distance at least $N/10$.[10] We know such codes exist that achieve the Hamming bound and have size $|C| = \Theta\left(\frac{2^N}{\sum_{t=0}^{0.05N-1} \binom{N}{k}}\right) \gtrsim \frac{2^{cN}}{\sqrt{N}} = 2^{\Omega(N)}$.

Consider a Hamiltonian on $2N$ qubits that is specified by a codeword $x \in \{0,1\}^N$. We divide the $2N$ qubits into pairs again and each pair will have Hamiltonian either H_0 or H_1 as before. Recall that H_1 depends on a parameter ε, which will be different from ε' and will be chosen later.

The Hamiltonian for the first pair of qubits is H_{x_1}, for the next pair is H_{x_2} and so on. So just as before, the Gibbs state will be $\rho_{x_1} \otimes \rho_{x_2} \otimes \cdots \otimes \rho_{x_N}$, and as before, these are diagonal states, so the resulting probability distributions will be $q_{x_1} \otimes q_{x_2} \otimes \cdots \otimes q_{x_N}$.

Just like before, we want to show that identifying the Hamiltonian (with probability $\geq 2/3$), which is equivalent to identifying the codeword x from which the Hamiltonian was constructed, requires many samples. We claim that if we learn the Hamiltonian to ℓ_2 error $\varepsilon' = 0.01\sqrt{N}\varepsilon$, then we can exactly identify the string x (with probability $\geq 2/3$). This step converts learning with ℓ_2 error to exact identification and this conversion dictates the value of ε in our definition of H_1.

Consider the unknown Hamiltonian on the first pair of qubits, H_{x_1}. This has two unknown coefficients, which are $-1/2 + \varepsilon x_0$ and $-1/2 - \varepsilon x_0$. Let's only consider the problem of learning the first of these coefficients for all our Hamiltonians H_{x_i}. Now if we have learned the coefficients to ℓ_2 error ε', it means we have a string λ_i that satisfies $\sqrt{\sum_i (\lambda_i + 1/2 - \varepsilon x_0)^2} \leq \varepsilon'$. By setting $y_i = \varepsilon(\lambda_i + 1/2)$, this means we have a string $y \in \mathbb{R}^N$ that satisfies $\sqrt{\sum_i (y_i - x_i)^2} \leq \varepsilon'/\varepsilon$.

We now use the property that x is a codeword of an error correcting code with large distance, so we can identify x given a close enough y. We know that any two codewords are at least $N/10$ apart in Hamming distance. This means any two codewords are at least $\sqrt{N/10}$ apart in ℓ_2 distance. Hence if we have a point in $\mathbb{R}^N$ (not just on the Boolean hypercube) that is ℓ_2 distance strictly less than $\sqrt{N/10}/2$ from a codeword x, it can be uniquely decoded to x. So, if we have a string

_____

[10]In other words, let C be a set of $2^{\Omega(N)}$ length-N bitstrings such that any two elements of C differ in $N/10$ bits.

$y \in \mathbb{R}^N$ such that $\sqrt{\sum_i (y_i - x_i)^2} \leq 0.01\sqrt{N}$, that will suffice. Thus we can choose ε to satisfy $\varepsilon' = 0.01\sqrt{N}\varepsilon$.

Now that we know that solving the Hamiltonian learning task allows us to exactly distinguish this set of Hamiltonians, let's show that distinguishing the Gibbs states of this set of Hamiltonians requires many samples using Fano's lemma.

To employ Fano's lemma, we need to bound the pairwise KL divergences again. We now consider $2^{\Omega(N)}$ probability distributions p_x, each corresponding to the Gibbs state of the Hamiltonian constructed from a codeword $x \in C$. Without loss of generality let us assume that $x = 0^N$ is part of the code, and let p_0 refer to the distribution corresponding to this Hamiltonian. As before, we let P_x be S copies of p_x. For any codeword $x \neq 0^N$, $D_{KL}(P_x \parallel P_0) = S\, D_{KL}(p_x \parallel p_0) \leq SN\, D_{KL}(q_1 \parallel q_0) \leq 8SN\beta^2\varepsilon^2 e^{-2\beta}$ using the chain rule and Lemma 5.1. So we can choose the parameter α in Fano's lemma to be $8SN\beta^2\varepsilon^2 e^{-2\beta} = O(S\beta^2\varepsilon'^2 e^{-2\beta})$. Applying Fano's inequality, we get

$$p_{\text{error}} = 1 - \frac{\log(2) + O(S(e^{-2\beta}\beta^2\varepsilon'^2))}{\log(2^{\Omega(N)})}, \tag{150}$$

which can be a small constant only if $S = \Omega\left(\frac{\exp(2\beta)N}{\beta^2\varepsilon'^2}\right)$. $\qquad\square$

6 Discussion

In this paper, we have addressed the Hamiltonian learning problem in a high-temperature regime. We have analyzed an algorithm to show that it has optimal sample complexity and time complexity. We were able to claim time optimality because our time complexity is simply linear in the sample size, the number of qubits in the total of all samples used in the algorithm. The critical temperature above which our algorithm is guaranteed to work depends only on the degree of the dual interaction graph, which we have treated as a constant in the optimality claims for sample and time complexity.

Although our algorithm is optimal for any fixed $\mathfrak{d}$, it might be possible to enlarge the temperature domain where our method works. The critical temperature to guarantee the convergence of the Newton–Raphson method is higher than that to ensure the convergence of the β-series expansion of $\langle E_a \rangle$; the former is $O(\mathfrak{d}^{10})$ (Theorem 4.6) while the latter is $O(\mathfrak{d}^2)$ (Theorem 3.1). This rather large discrepancy occurred when we used the "band-diagonal" property of the Jacobian of $\mathcal{F}$, and it will require finer understanding of these correlations to improve our bounds in terms of $\mathfrak{d}$. It is also feasible to extend our algorithm beyond low-intersection Hamiltonians to local Hamiltonians, where $\mathfrak{d}$ need not be constant, since cluster expansion works in the more general setting where one-spin energy is bounded. However, the number of monomials still scales exponentially in $\mathfrak{d}$, so even writing down the truncated Taylor series expansion could be computationally expensive.

The problem of finding an efficient learning algorithm in the low-temperature regime remains completely open. Our high-temperature expansion does not converge in general for large β since there are systems that undergo phase transitions as we lower the temperature, where the partition function is not analytically continued from the high-temperature domain. In fact, an efficient algorithm for all temperatures, if it exists, should not attempt to evaluate partition functions since low-temperature partition functions are generally (at least) NP-hard to compute. The classical polynomial-time algorithms avoid evaluating partition functions using conditional independence (the Markov property), but this does not hold in general for quantum noncommuting Hamiltonians.

Acknowledgements

A   Learning Hamiltonians from real time dynamics

Suppose we are given a blackbox that implements unitary time evolution

$$U = e^{-itH} \tag{151}$$

governed by a fixed, time-independent, unknown Hamiltonian $H = \{(a, E_a, \lambda_a) : a \in [M]\}$. We assume that the evolution time t is known to us, and the Hamiltonian follows the same normalization as in the main text: E_a are distinct Pauli matrices and $\lambda_a \in [-1, 1]$ for all a. The blackbox converts any input state represented by a density matrix ρ to $U\rho U^\dagger$. Now, the learning problem is to estimate λ_a to additive accuracy ε with as few uses of blackbox U as possible.

We consider a scenario where t is smaller than some constant t_c that only depends on the structure of Hamiltonian terms E_a but not on the coefficients λ_a. The learning algorithm and its analysis will be very similar to that in the main text, so we will be brief. We restate the theorem we will prove.

Theorem 1.3 (Real-time dynamics). *Let H be a low-intersection Hamiltonian on N qubits and let $U = e^{-itH}$ be a blackbox unitary with $t < t_c$. Then we can learn the coefficients of H to ℓ_∞ error ε with success probability $1 - \delta$, using U $O\left(\frac{1}{t\varepsilon^2} \log \frac{N}{\delta}\right)$ times, with time complexity $O\left(\frac{N}{t\varepsilon^2} \log \frac{N}{\delta}\right)$.*

Remark A.1. In what follows below, we prove Theorem 1.3 with t replaced by t^2. We can improve this dependence from $1/t^2$ to $1/t$ by reducing to a setting where t is constant: by applying U $n = \lfloor t_c/t \rfloor$ times, we can produce a black box for the unitary $V = U^n = \exp(-iHt\lfloor t_c/t \rfloor)$. Learning parameters from V is the same problem as learning parameters from U, except the parameter t becomes $t\lfloor t_c/t \rfloor = \Theta(t_c)$, which is constant (determined only by $\mathfrak{d}$). So, we can use the algorithm below with the unitary V, requiring $O\left(\frac{1}{\varepsilon^2} \log \frac{N}{\delta}\right)$ applications of V, and therefore $O\left(\frac{1}{t\varepsilon^2} \log \frac{N}{\delta}\right)$ applications of U. The time complexity is also inflated by $n = \Theta(1/t)$ in a similar fashion.

Note that the complexity in the time parameter is optimal; two time-evolution operators I and e^{-itZ} on one qubit differ by $O(t)$ in operator norm and hence also in completely bounded (diamond) norm as quantum channels.

A.1   Series expansion of time-evolved operators

Similarly to Theorem 3.1 for the β-series expansion of $\text{Tr}(E_a e^{-\beta H})/\text{Tr}\, e^{-\beta H}$, in this section we prove properties about the t-series expansion of $UPU^\dagger$ where P is a single-qubit Pauli operator. The relevance of this quantity to the learning problem will be evident in the next subsection. In the following theorem, the pink text indicates where it differs from Theorem 3.1; morally, the same properties are proven, just for a different series. All of the quantitative bounds are at least as strong as those in Theorem 3.1, and though we prove results for a matrix-valued polynomial, they are indeed comparable when considering their trace against the operator Q, as defined in $\mathcal{F}(Q, P)$.

Theorem A.2. *Consider a Hamiltonian $\{(a, E_a, \lambda_a) : a \in [M]\}$. Then, for every single-qubit Pauli operator P and L-qubit Pauli operator Q, we have a Taylor series expansion*

$$UPU^\dagger = \sum_{m=1}^{\infty} t^m q_m(\lambda_1, \ldots, \lambda_M) \tag{152}$$

$$\mathcal{F}(Q, P) := \frac{1}{\mathsf{D}} \mathrm{Tr}(QUPU^\dagger) = \frac{1}{\mathsf{D}} \sum_{m=1}^{\infty} \mathrm{Tr}(Qq_m(\lambda)) \tag{153}$$

where equality holds whenever the series converges absolutely. For any $m \in \mathbb{Z}_{>0}$, the following hold:

1. *$q_m \in \mathbb{C}^{\mathsf{D} \times \mathsf{D}}[\lambda_1, \ldots, \lambda_M]$ is a degree m homogeneous matrix-valued polynomial in the Hamiltonian term coefficients.*

2. *Let $\mathfrak{G}(P)$ denote the dual interaction graph among operators $\{E_1, \ldots, E_M, P\}$, i.e., $\mathfrak{G}(P)$ is $\mathfrak{G}$ with an extra node p and an extra edge (p, a) if and only if $\mathrm{Supp}(E_a) \cap \mathrm{Supp}(P) \neq \varnothing$. Then q_m involves λ_a only if the distance between p and a on $\mathfrak{G}(P)$, $\mathrm{dist}_{\mathfrak{G}(P)}(p, a)$, is at most m.*

3. *q_m consists of at most $\max(L, \mathfrak{d})e\mathfrak{d}(1 + e(\mathfrak{d} - 1))^{m-1} \leq e\mathfrak{d}(1 + e(\mathfrak{d} - 1))^m$ monomials.*

4. *The coefficient matrix in front of any monomial of q_m has spectral norm at most 2^m in magnitude.*

Suppose further that every E_a is a tensor product of Pauli matrices, supported on at most L qubits. Then, after $O(LM\mathfrak{d} \log \mathfrak{d})$ pre-processing time (see Remark 2.4), the following are true for every $m \in \mathbb{Z}_{>0}$.

A. *The list of monomials that appear in q_m can be enumerated in time $O(m\mathfrak{d}C)$, where C is the number of monomials (so, in particular, in time $O(m\mathfrak{d}^2(1 + e(\mathfrak{d} - 1))^m)$).*

B. *The truncated series of $\mathcal{F}(Q, P)$, $\frac{1}{\mathsf{D}} \sum_{\ell=1}^{m} \mathrm{Tr}(Qq_\ell(\lambda))$, can be computed exactly as a rational polynomial in $C(4^m + L) \mathrm{poly}(m)$ time.*

To understand $UPU^\dagger$, we recall the well-known formula for square matrices A and B,

$$e^A B e^{-A} = \sum_{n=0}^{\infty} \frac{[A, B]_n}{n!} \tag{154}$$

$$\text{where } [A, B]_k = \begin{cases} B & (k = 0) \\ A[A, B]_{k-1} - [A, B]_{k-1}A & (k \geq 1). \end{cases}$$

Since the nested commutator $[A, B]_n$ has norm upper bounded by $2^n \|A\|^n \|B\|$, which grows only exponentially with n, this series always converges absolutely for any finite dimensional matrices over complex numbers. Applying it to our case, we have

$$UPU^\dagger = \sum_{n=0}^{\infty} \frac{(-it)^n}{n!} [H, P]_n = P - it[H, P] - \frac{t^2}{2}[H, [H, P]] + \cdots \tag{155}$$

$$= \sum_{\mathbf{V}} \frac{\lambda^{\mathbf{V}}}{\mathbf{V}!} \mathcal{D}_{\mathbf{V}}(UPU^\dagger) \tag{156}$$

where in the second line we re-express the t-series as a multivariate Taylor series in λ_a. This is our series: Item A.2(1) follows from Eq. (155) upon taking $q_m(\lambda) := \frac{(-i)^m}{m!}[H(\lambda), P]_m$.

We now examine a cluster derivative with respect to $\mathbf{V}$, which has total weight $n = |\mathbf{V}|$. Let us enumerate all elements of $\mathbf{V}$ as $a_1, a_2, \ldots, a_n$; this is a list of nodes of $\mathfrak{G}$ with no particular order and the elements a_j are repeated as many times as their multiplicities. In this context, the cluster derivative $\mathcal{D}_{\mathbf{V}}(UPU^\dagger) = [\partial_{a_1} \cdots \partial_{a_n} UPU^\dagger]|_{\lambda=(0,\ldots,0)}$ is a constant matrix which comes from evaluating a derivative at the origin of the λ-space $[-1, 1]^M$.

$$\mathcal{D}_{\mathbf{V}}(UPU^\dagger) = \mathcal{D}_{\mathbf{V}} \frac{(-it)^n}{n!} [\underbrace{H, [H, \cdots [H}_{n}, P] \cdots]] \tag{157}$$

$$= \frac{(-it)^n}{n!} \sum_{\sigma \in S_n} [\partial_{a_{\sigma(1)}} H, [\partial_{a_{\sigma(1)}} H, \cdots [\partial_{a_{\sigma(n)}} H, P] \cdots]] \qquad \text{by the Leibniz rule}$$

$$= \frac{(-it)^n}{n!} \sum_{\sigma \in S_n} [E_{a_{\sigma(1)}}, [E_{a_{\sigma(1)}}, \cdots [E_{a_{\sigma(n)}}, P] \cdots]]$$

where S_n is the permutation group on $\{1, 2, \ldots, n\}$. From Equation (157), the rest of Theorem A.2 will follow.

Lemma A.3. *For any cluster $\mathbf{V}$ on $\mathfrak{G}$, if $\mathbf{V} \sqcup \{(p, 1)\}$ is disconnected on $\mathfrak{G}(P)$, then $\mathcal{D}_{\mathbf{V}}(UPU^\dagger) = 0$.*

Proof. Consider a term in Eq. (157), which we can label as $[E_{a_1}, [E_{a_2}, \cdots [E_{a_n}, P] \cdots]]$ without loss of generality. If $\mathbf{V} \sqcup \{(p, 1)\}$ is disconnected, then there exists an $k \in [n]$ such that a_k is disconnected from all of $a_{k+1}, \ldots, a_n, p$ (otherwise, every a_k would have a path to p by strong induction, making the cluster connected). Consequently, E_{a_k} commutes with the intermediate commutator $C_{k+1} = [E_{a_{k+1}}, \cdots [E_{a_n}, P] \cdots]$, which is supported on $R = \mathrm{Supp}(P) \cup \bigcup_{j=k+1}^n \mathrm{Supp}(E_{a_j})$. This means that the next $C_k = [E_{a_k}, C_k]$ is zero and so the whole term $[E_{a_1}, [E_{a_2}, \cdots [E_{a_n}, P] \cdots]]$ is zero. This argument applies to every term, so the whole sum, and the cluster derivative, must also be zero. $\quad\square$

By this lemma, Item A.2(2) follows immediately, since λ_a is present in q_m, then there must be a cluster $\mathbf{V}$ of size m such that $a \in \mathbf{V}$ and $\mathbf{V} \sqcup \{(p, 1)\}$ is connected. This implies that the distance between a and p is at most m. Similarly, the number of monomials of q_m can be bounded by the number of weight-m connected clusters in $\mathfrak{G}$ neighboring p in $\mathfrak{G}(P)$. By Proposition 3.6, this can be bounded by $\max(L, \mathfrak{d} + 1)e\mathfrak{d}(1 + e(\mathfrak{d} - 1))^{m-1}$, where the additional factor of $\max(L, \mathfrak{d} + 1)$ comes from needing to count clusters that start at any of the terms adjacent to p. This gives Item A.2(3). The lemma below gives Item A.2(4).

Lemma A.4. *For any cluster $\mathbf{V}$ with $|\mathbf{V}| = n$, we have $\|\mathcal{D}_{\mathbf{V}}(UPU^\dagger)\| \le 2^n |t|^n$.*

Proof. The norm of a nested commutator in the last line of Eq. (157) is at most 2^n. $\quad\square$

Finally, for the time complexity results, note the same algorithm for computing clusters works in this setting, giving Item A.2(A). To compute the series q_m, one could use the same approach as Section 3.6, but we take a simpler and faster approach: we have an explicit form for the series, Eq. (155), so all we need to do is compute the commutators $[H, P]_n$ iteratively, for n from 1 to m. We can maintain $[H, P]_n$ as a sum over clusters of monomials $\lambda^{\mathbf{V}}$ with corresponding integer matrices (where each integer is bounded by $2^n n!$ by Item A.2(4)), of which faithful representations can be maintained as done in Section 3.6. For each of these integer matrices X, one can compute the corresponding commutator $[H, X]$ in $O(4^n \mathrm{poly}(n))$ time, giving the matrices for the next commutator $[H, P]_{n+1}$. This gives the specified runtime.

A.2 Learning algorithm

Our learning algorithm in the "real-time dynamics" setting will be essentially the same as that of learning from the "Gibbs state" setting. For each node a of $\mathfrak{G}$, choose P_a to be any single-qubit Pauli that anticommutes with E_a, and let $Q_a = i[P_a, E_a] = 2iP_aE_a$. Define $\mathcal{F}_a = \mathcal{F}(Q_a, P_a)$. Our learning algorithm consists of two parts.

1. Find estimates $\tilde{\mathcal{F}}_a$ such that $\left|\tilde{\mathcal{F}}_a - \mathcal{F}_a\right| \leq t\varepsilon$ for all $a \in [M]$.

2. Approximately invert the function $\{\lambda_a\} \mapsto \{\mathcal{F}_a\}$ given $\tilde{\mathcal{F}}_a$.

Lemma A.5. *Consider a Hamiltonian $\{(a, E_a, \lambda_a) : a \in [M]\}$ on N qubits. We can find estimates $\tilde{\mathcal{F}}_a$ such that $\left|\tilde{\mathcal{F}}_a - \mathcal{F}_a\right| \leq t\varepsilon$ for all $a \in [M]$, with probability at least $1 - \delta$, using only $O(\frac{\mathfrak{d}}{t^2\varepsilon^2}\log(\frac{M}{\delta}))$ applications of U and with time complexity $O(\frac{N\mathfrak{d}}{t^2\varepsilon^2}\log(\frac{M}{\delta}))$.*

Proof. Recall that in Lemma 4.1, we argued that, from one copy of ρ, it's possible to generate a bounded random variable $Y_a \in [-1, 1]$ that is an unbiased esimator of $\text{Tr}(E_a\rho)$. Moreover, for a set of terms $S \subset [M]$, it is possible to generate Z_a's for all such $a \in S$ from one copy of ρ, provided the E_a's are non-overlapping, or in other words, provided S is an independent set in the dual interaction graph $\mathfrak{G}$. We will use this again here; the main challenge is that for each E_a we wish to measure against a different P_a, so it's not immediately clear how to use one application of U to produce estimators for multiple different terms. We resolve this by thinking of ρ as a distribution over states, and then conditioning on this distribution to measure expectations over what are effectively different mixed states.

Consider the procedure of sampling a string $s \sim \{0, 1, +, -, i, -i\}^N$ uniformly at random, and then preparing the state $\rho = |s_1\rangle\langle s_1| \otimes \cdots \otimes |s_N\rangle\langle s_N|$. Notice that if we discard our initial string s, thereby averaging over s, then ρ is the maximally mixed state; further, if we discard the entire initial string apart from one qubit i, then $\rho = \frac{I}{2} \otimes \cdots \otimes \frac{I}{2} \otimes |s_i\rangle\langle s_i| \otimes \frac{I}{2} \otimes \cdots \otimes \frac{I}{2}$. Note that, for $s = 0, 1, +, -, i, -i$, $2|s\rangle\langle s| - I$ is a Pauli matrix $Z, -Z, X, -X, Y, -Y$, respectively.

Suppose we apply U to ρ, measure it on the support of E_a to get the unbiased estimator Y_a of $\text{Tr}(E_aU\rho U^\dagger)$, and define the following random variable.

$$Z_a = \begin{cases} Y_a & \text{if } s_{\text{Supp}(P_a)} \text{ satisfies } 2|s_{\text{Supp}(P_a)}\rangle\langle s_{\text{Supp}(P_a)}| - I = P_a \\ 0 & \text{otherwise} \end{cases}$$

Here, we abuse notation by using P_a to refer both to the 1-qubit Pauli and the n-qubit tensor of that Pauli with the identity matrix. The random variable Z_a is bounded in $[-1, 1]$ because Y_a is, and furthermore,

$$\begin{aligned}
\text{E}[Z_a] &= \Pr_s\left[2|s_{\text{Supp}(P_a)}\rangle\langle s_{\text{Supp}(P_a)}| - I = P_a\right]\text{E}_s\left[Y_a \mid 2|s_{\text{Supp}(P_a)}\rangle\langle s_{\text{Supp}(P_a)}| - I = P_a\right] \\
&= \frac{1}{6}\text{Tr}\left(E_aU\left(\text{E}[\rho \mid 2|s_{\text{Supp}(P_a)}\rangle\langle s_{\text{Supp}(P_a)}| - I = P_a]\right)U^\dagger\right) \\
&= \frac{1}{6}\text{Tr}\left(E_aU\left(\frac{I}{2} \otimes \cdots \otimes \frac{I}{2} \otimes \frac{I + P_a}{2} \otimes \frac{I}{2} \otimes \cdots \otimes \frac{I}{2}\right)U^\dagger\right) \\
&= \frac{1}{12}\text{Tr}\left(E_aUP_aU^\dagger\right)
\end{aligned}$$

So, $12Z_a$ is an unbiased estimator for $\mathcal{F}_a$, and the rest of the result follows exactly like it did in Lemma 4.1: one can use one copy of ρ to get multiple estimators Z_a, provided their corresponding terms do not overlap. Thus, in $\mathfrak{d} + 1$ rounds of $O(\frac{1}{t^2\varepsilon^2}\log\frac{M}{\delta})$ applications of U each, one can get

$O(\frac{1}{t^2\varepsilon^2}\log\frac{M}{\delta})$ copies of Z_a for every $a \in [M]$. By Chernoff bound, each rescaled average $12\bar{Z}_a$ will then satisfy $\left|12\bar{Z}_a - \mathcal{F}_a\right| \leq \varepsilon$ with probability $\geq 1 - \delta/M$, and so they all satisfy $\left|12\bar{Z}_a - \mathcal{F}_a\right| \leq \varepsilon$ with probability $\geq 1 - \delta$. The time complexity is the same, since the only change to the procedure is doing one additional $O(1)$-time check per sample Z_a. $\qquad\square$

Finally, for the second, classical part of the algorithm, note that operators P_a and Q_a are chosen so that the leading term of $\mathcal{F}_a$ is a known constant multiple of λ_a:

$$\mathcal{F}(Q, P) = \frac{\text{Tr}}{\text{D}}(QP - itQ[H, P] + \cdots) \tag{158}$$
$$= t\frac{\text{Tr}}{\text{D}}\left([P, E_a]\sum_b \lambda_b[E_b, P] + \cdots\right)$$
$$= 4t\lambda_a + \cdots$$

where the last line uses the orthonormality of the Hamiltonian terms. This observation implies that the Jacobian is "band-diagonal", and suffices, along with Theorem A.2, for the full analysis of the Newton–Raphson method in Section 4.2 to go through identically. The only difference here is that β is replaced with $-4t$, so this part takes time $O(ML/\varepsilon)\,\text{poly}(\mathfrak{d}\log(1/t\varepsilon))$. The time complexity of the quantum part dominates.

Remark A.6. Since the series expansion in Eq. (158) is only shown to converge for $t < t_c$ where $t_c = 1/\text{poly}(\mathfrak{d})$, we can only claim that our algorithm works for small enough t. In the learning problem from Gibbs states, the analogous condition $\beta \leq 1/\text{poly}(\mathfrak{d})$ is due to the fact that our approach cannot handle arbitrarily low temperature; the sample complexity result [AAKS21] shows that learning is feasible for all temperatures, at least in an information-theoretic sense. In contrast, in the learning problem from real-time evolution, it is fundamental that we have to restrict the evolution time to be smaller than some constant set by $\mathfrak{d}$; for a certain long time, the learning is simply impossible. Consider a Hamiltonian $H = -\lambda I + \lambda(I + Z_1)(I + Z_2)\cdots(I + Z_n)$ on n qubits where Z_j is the Pauli Z on qubit j. This Hamiltonian is the sum of all nonidentity products of Z's with a uniform coefficient $\lambda \in [-1, 1]$, and obeys our normalization conditions for Hamiltonians. The intersection degree $\mathfrak{d}$ is exponentially large in n. The eigenspectrum of H consists of just two values, $(2^n - 1)\lambda$ and $-\lambda$. Hence, $e^{-itH} \propto I$ if $t = 2\pi/2^n\lambda = \Theta((\mathfrak{d}\lambda)^{-1})$. Since λ is unknown, we conclude that no general algorithm can determine λ unless we restrict t to be smaller than $1/\text{poly}(\mathfrak{d})$.

B   Algorithm for parameter learning of Markov random fields

In this section, we will prove a folklore result by giving a simple algorithm for parameter learning of Markov random fields. By the Hammersley–Clifford theorem, a Markov random field $\mathcal{D}$ over $\{-1, +1\}^N$ can be written as

$$\Pr_{Z\sim\mathcal{D}}[Z = z] \propto \exp(\sum_{S\subset[N]} \psi_S(z_S))$$

for some functions $\psi_S : \mathbb{R}^{|S|} \to \mathbb{R}$, where $z_S = (z_i)_{i\in S}$. Typically, the sum is restricted to be over S with size at most some constant. By writing every ψ_S as a sum of products of variables, this expression becomes

$$\Pr_{Z\sim\mathcal{D}}[Z = z] \propto \exp(\sum_{S\subset[N]} \lambda_S z^S),$$

49

where $z^S := \prod_{i \in S} z_i$. For the parameter learning problem, we assume we already know the structure of the MRF, so suppose we are given a hypergraph $G = (V = [N], E)$ on N vertices such that

$$\Pr_{Z \sim \mathcal{D}}[Z = z] \propto \exp(-\beta \sum_{S \in E} \lambda_S z^S).$$

Here, $-\beta$ is a rescaling factor so that we can assume without loss of generality that $\lambda_S \in [-1, 1]$ for all $S \in E$. We will interpret the λ parameters as a vector in $[-1, 1]^{|E|}$ and $\beta \in (0, \infty)$. This is precisely a Gibbs state of a classical Hamiltonian, following the definitions given in Section 2.1. Further, this is the setting where each term is is a product of Paulis, since each z^S is a product of Pauli Z operators.

For a vertex $i \in [N]$, let $E_i = \{S \in E \mid i \in S\}$ be the set of hyperedges containing i and let $N_i = (\cup_{S \in E_i} S) \setminus \{i\}$ be the neighborhood of i. Our algorithm will depend on two parameters: maximum degree $d := \max_{i \in [N]} |E_i|$ and an "average order" parameter $L := \max_{i \in [N]} \frac{1}{d} |N_i \cup \{i\}|$. We do not consider the empty graph, so that $d, Ld \geq 1$.

We will need the following lemma.

Lemma B.1 (Version of Lemma 2.1, [Bre15])**.** *For any node $u \in V$, subset $S \subset V$, and configuration $x_S \in \{\pm 1\}^{|S|}$,*

$$\min_{b \in \{-1, +1\}} \Pr[X_u = b \mid X_S = x_S] \geq \frac{1}{\exp(2\beta d) + 1} \geq \frac{1}{2} \exp(-2\beta d).$$

Proof. First, using the Markov property

$$
\begin{aligned}
|\mathrm{E}[X_u \mid X_{V \setminus \{u\}} = x_{V \setminus \{u\}}]| &= \left| \frac{\sum_{x_u} x_u \exp(-\beta(\sum_{S \in E} \lambda_S x^S))}{\sum_{x_u} \exp(-\beta(\sum_{S \in E} \lambda_S x^S))} \right| \\
&= \left| \frac{\sum_{x_u} x_u \exp(-\beta(\sum_{S \in E_u} \lambda_S x^{S \setminus \{u\}} x_u))}{\sum_{x_u} \exp(-\beta(\sum_{S \in E_u} \lambda_S x^{S \setminus \{u\}} x_u))} \right| \\
&= \left| \frac{\exp(-\beta(\sum_{S \in E_u} \lambda_S x^{S \setminus \{u\}})) - \exp(\beta(\sum_{S \in E_u} \lambda_S x^{S \setminus \{u\}}))}{\exp(-\beta(\sum_{S \in E_u} \lambda_S x^{S \setminus \{u\}})) + \exp(\beta(\sum_{S \in E_u} \lambda_S x^{S \setminus \{u\}}))} \right| \\
&= \frac{\exp(\beta |\sum_{S \in E_u} \lambda_S x^{S \setminus \{u\}}|) - \exp(-\beta |\sum_{S \in E_u} \lambda_S x^{S \setminus \{u\}}|)}{\exp(\beta |\sum_{S \in E_u} \lambda_S x^{S \setminus \{u\}}|) + \exp(-\beta |\sum_{S \in E_u} \lambda_S x^{S \setminus \{u\}}|)} \\
&= 1 - \frac{2}{\exp(2\beta |\sum_{S \in E_u} \lambda_S x^{S \setminus \{u\}}|) + 1}
\end{aligned}
$$

Further, by the tower property of conditional expectation and Jensen's inequality,

$$
\begin{aligned}
&|\mathrm{E}[X_u \mid X_S = x_S]| \\
&= |\mathrm{E}[\mathrm{E}[X_u \mid X_{V \setminus \{u\}} = x_{V \setminus \{u\}}] \mid X_S = x_S]| \\
&\leq \mathrm{E}[|\mathrm{E}[X_u \mid X_{V \setminus \{u\}} = x_{V \setminus \{u\}}]| \mid X_S = x_S] \\
&= \mathrm{E}\left[1 - \frac{2}{\exp(2\beta |\sum_{S \in E_u} \lambda_S x^{S \setminus \{u\}}|) + 1} \mid X_S = x_S\right] \\
&= 1 - \frac{2}{\exp(2\beta |\sum_{S \in E_u} \lambda_S x^{S \setminus \{u\}}|) + 1} \\
&\leq 1 - \frac{2}{\exp(2\beta d) + 1}
\end{aligned}
$$

50

For a $\{\pm 1\}$-valued random variable X, $\min\{\Pr[X = 1], \Pr[X = -1]\} = \frac{1}{2}(1 - |\mathrm{E}[X]|)$, so

$$\min_{b \in \{-1,+1\}} \{\Pr[X_u = b \mid X_S = x_S]\} = \frac{1}{2}(1 - |\mathrm{E}[X_u \mid X_S = x_S]|) \geq \frac{1}{\exp(2\beta d) + 1}. \qquad \square$$

Define the sigmoid function $\sigma(x) := \frac{e^x}{1+e^x}$. We will need the following fact about the sigmoid:

Lemma B.2 (Claim 4.2, [KM17]). *For all x, y, $|\sigma(x) - \sigma(y)| \geq \exp(-|x| - 3)\min(1, |x - y|)$.*

Theorem B.3. *Fix $\varepsilon \in (0, 1)$. Given samples from the MRF $X^{(1)}, \ldots, X^{(T)}$ with $T = \Theta(\exp(8\beta L d^2 + 2Ld)\frac{1}{\beta^2 \varepsilon^2} \log \frac{N}{\delta})$, we can compute an estimate $\hat{\lambda}$ such that $\|\hat{\lambda} - \lambda\|_\infty \leq \varepsilon$ with probability $\geq 1 - \delta$. The algorithm takes $O(TN(Ld^2 + d2^d))$ time.*

For low-intersection Hamiltonians (as defined in the introduction), $L = O(1)$ and $d \leq \mathfrak{d}+1 = O(1)$. For constant L and d, the sample complexity and time complexity of learning a classical Hamiltonian to ℓ_∞ error ε become

$$\frac{\exp(O(\beta))}{\beta^2 \varepsilon^2} \log \frac{N}{\delta} \quad \text{and} \quad \frac{\exp(O(\beta))}{\beta^2 \varepsilon^2} N \log \frac{N}{\delta},$$

respectively.

Proof. First, fix a particular $v \in [n]$, and consider conditioning on its neighbors N_v. The distribution on x_v after conditioning is

$$\Pr[X_v = x_v \mid X_{N_v} = x_{N_v}] = \sigma(2\beta \sum_{S \in E_v} \lambda_S x^S).$$

We now show that it suffices to be able to estimate such conditional probabilities, for a particular setting of $X_{N(v)}$. Let $q_x^{(v)}$ be the argument inside the σ above, so

$$q_x^{(v)} := \sigma^{-1}(\Pr[X_v = x_v \mid X_{N_v} = x_{N_v}]) = 2\beta \sum_{S \in E_v} \lambda_S x^S.$$

Note that $q_x^{(v)}$ only depends on those x_u where $u \in N_v$. We will argue below that we can get an estimate of our conditional probability $\sigma(q_x^{(v)})$ to $\exp(-|q_x^{(v)}| - 3)\min(0.5, 2\beta\varepsilon)$ error. We will invert σ on this estimate to get an estimate $\hat{q}_x^{(v)}$ for $q_x^{(v)}$, so we denote the original estimate to be $\sigma(\hat{q}_x^{(v)})$. By Lemma B.2,

$$\min(1, |q_x^{(v)} - \hat{q}_x^{(v)}|) \leq \exp(|q_x^{(v)}| + 3)|\sigma(q_x^{(v)}) - \sigma(\hat{q}_x^{(v)})| \leq \min(0.5, 2\beta\varepsilon)$$
$$\text{so } |q_x^{(v)} - \hat{q}_x^{(v)}| \leq \min(0.5, 2\beta\varepsilon) \leq 2\beta\varepsilon$$

So $\hat{q}_x^{(v)}$ is an estimate for $q_x^{(v)}$ up to additive $2\beta\varepsilon$ error. Now, we show that we can use these $\hat{q}_x^{(v)}$'s to get a good estimate of the parameters $\{\lambda_S\}_{S \in E}$.

Suppose we want to know λ_S. Pick $v \in S$ and consider all $T \in E_v$ that are not S. If $T \nsubseteq S$, then choose some vertex $u \in T \setminus S$ and place it in a set N_{out}. Otherwise, $T \subsetneq S$, so we choose some vertex $u \in S \setminus T$ and place it in a set N_{in}. Note that $N_{\text{in}} \subset S$ and $N_{\text{out}} \subset [N] \setminus S$, so they are disjoint.[11] We can get an estimate for λ_S by averaging $q_z^{(v)}$ over the coordinates of N_{in} and N_{out} in

---

[11]For some intuition, in the conventional setting where vertices are on a lattice and a term S is a connected piece of the lattice, one can think of taking N_{in} to be S and N_{out} to be the neighborhood of S (or the boundary of S^c).

a particular way; for a set of indices $I \subset [n]$, let $P_I = \{z \in \{\pm 1\}^N \mid z_i = 1 \text{ for all } i \notin I\}$ be a slice of the Hamming cube.

$$
\begin{aligned}
\frac{1}{2\beta} \underset{z \sim P_{N_{\mathrm{in}} \cup N_{\mathrm{out}}}}{\mathrm{E}} [z^{N_{\mathrm{in}}} q_z^{(v)}] &= \Big[\underset{z_{N_{\mathrm{in}}}}{\mathrm{E}} \underset{z_{N_{\mathrm{out}}}}{\mathrm{E}} \Big[z^{N_{\mathrm{in}}} \sum_{T \in E_v} \lambda_T z^T \Big] \Big]_{z=\vec{1}} \\
&= \Big[\underset{z_{N_{\mathrm{in}}}}{\mathrm{E}} \Big[\sum_{\substack{T \in E_v \\ T \cap N_{\mathrm{out}} = \varnothing}} \lambda_T z^{N_{\mathrm{in}}} z^T \Big] \Big]_{z=\vec{1}} = \Big[\sum_{\substack{T \in E_v \\ T \cap N_{\mathrm{out}} = \varnothing \\ T^c \cap N_{\mathrm{in}} = \varnothing}} \lambda_T z^{T \setminus N_{\mathrm{in}}} \Big]_{z=\vec{1}} = \Big[\lambda_S z^{S \setminus N_{\mathrm{in}}} \Big]_{z=\vec{1}} = \lambda_S
\end{aligned}
$$

Suppose we have an estimate of $q_z^{(v)}$, $\hat{q}_z^{(v)}$, to $2\beta\varepsilon$ error. Then

$$
\begin{aligned}
\Big| \frac{1}{2\beta} \underset{z \sim P_{N_{\mathrm{in}} \cup N_{\mathrm{out}}}}{\mathrm{E}} z^{N_{\mathrm{in}}} \hat{q}_z^{(v)} - \lambda_S \Big| &= \frac{1}{2\beta} \Big| \underset{z \sim P_{N_{\mathrm{in}} \cup N_{\mathrm{out}}}}{\mathrm{E}} z^{N_{\mathrm{in}}} (\hat{q}_z^{(v)} - q_z^{(v)}) \Big| \\
&\leq \frac{1}{2\beta} \underset{z \sim P_{N_{\mathrm{in}} \cup N_{\mathrm{out}}}}{\mathrm{E}} \Big| z^{N_{\mathrm{in}}} \Big| \Big| \hat{q}_z^{(v)} - q_z^{(v)} \Big| \\
&\leq \frac{1}{2\beta} \underset{z \sim P_{N_{\mathrm{in}} \cup N_{\mathrm{out}}}}{\mathrm{E}} \Big| z^{N_{\mathrm{in}}} \Big| 2\beta\varepsilon \\
&= \varepsilon
\end{aligned}
$$

Note that $|P_{N_{\mathrm{in}} \cup N_{\mathrm{out}}}| = 2^{|N_{\mathrm{in}}| + |N_{\mathrm{out}}|} \leq 2^d$. So, now we just need to show how to estimate $\sigma(q_z^{(v)}) = \Pr[X_v = z_v \mid X_{N_v} = z_{N_v}]$ to $\exp(-\left| q_x^{(v)} \right| - 3) \min(0.5, 2\beta\varepsilon)$ error for $\leq 2^{d+1}$ choices of z_v, z_{N_v}, over all M choices of S. Since $\left| q_x^{(v)} \right| \leq 2\beta d$ always, it suffices to estimate to $\Theta(\exp(-2\beta d) \min(0.5, \beta\varepsilon))$ error.

Recall that estimating the probability p of an event occurring to α relative error with probability $\geq 1 - \delta$ requires $\Theta(\frac{1}{p\alpha^2} \log \frac{1}{\delta})$ samples. (The estimator we use is simply the empirical probability of the event over the samples, and the proof follows from a Chernoff bound). So, if we pull $\Theta(\frac{1}{p_{\min} \alpha^2} \log \frac{M 2^{d+2}}{\delta})$ samples of our Markov random field, then we can get an estimate to any particular choice of $\Pr[X_v = z_v \text{ and } X_{N_v} = z_{N_v}]$ and $\Pr[X_{N_v} = z_{N_v}]$ to α relative error that is correct with probability $\geq 1 - \frac{\delta}{M 2^{d+2}}$, provided that $p_{\min}$ is chosen to be smaller than this probability. By union bound, we can get an estimate to all the probabilities we would need to compute the C conditional probabilities with probability $\geq 1 - \delta$ using this number of samples, provided $p_{\min}$ is smaller than *all* the probabilities we wish to estimate. Also recall that by Lemma B.1,

$$
\Pr[X_S = x_S] = \prod_{i=1}^{|S|} \Pr[X_{u_i} = x_{u_i} \mid X_{u_j} = x_{u_j} \text{ for } j < i] \geq 2^{-|S|} \exp(-2\beta |S| d).
$$

Since we want to compute these probabilities for $S = N_v$ or $S = N_v \cup \{v\}$, we can take $p_{\min} = \exp(-2\beta(Ld)d - Ld)$.

With these estimates, we can get good estimates to the conditional probabilities, assuming α is sufficiently small:

$$
\begin{aligned}
\frac{\Pr[X_v = z_v \text{ and } X_{N_v} = z_{N_v}](1 \pm \alpha)}{\Pr[X_{N_v} = z_{N_v}](1 \pm \alpha)} &\in \frac{\Pr[X_v = z_v \text{ and } X_{N_v} = z_{N_v}]}{\Pr[X_{N_v} = z_{N_v}]} (1 \pm O(\alpha)) \\
&= \Pr[X_v = z_v \mid X_{N_v} = z_{N_v}](1 \pm O(\alpha))
\end{aligned}
$$

We set $\alpha = \Theta(\exp(-2\beta d) \min(0.5, \beta\varepsilon)) < 0.5$ to conclude. The number of samples we need is

$$
T = \Theta\Big(\frac{\exp(2\beta(Ld)d + Ld)}{(\exp(-2\beta d) \min(1, \beta\varepsilon))^2} \log \frac{M 2^{d+2}}{\delta} \Big)
$$

Using that $\frac{1}{\min(1,\beta^2\varepsilon^2)} = \frac{1}{\beta^2\varepsilon^2}\max(1,\beta\varepsilon)^2 \le \frac{1}{\beta^2\varepsilon^2}\exp(2\beta\varepsilon)$, this is

$$= \Theta\Big(d\exp(2\beta(Ld)d + Ld + 4\beta d + 2\beta\varepsilon)\frac{1}{\beta^2\varepsilon^2}\log\frac{M}{\delta}\Big).$$

The bound in the theorem statement comes from simplifying and performing rough bounds on the above expression. We can run this algorithm in time $O(TM(Ld + 2^d))$, since for each term, we can compute the 2^d empirical conditional probabilities for it by taking Ld time per sample to sort them into the various (disjoint) events. Classifying each sample only requires looking at the bits corresponding to v and its neighbors, so checking all takes $O(TMLd)$ time, and we need to do this for all terms. The result in the statement comes from taking $M \le Nd$. $\qquad\square$

References

[AAKS21] Anurag Anshu, Srinivasan Arunachalam, Tomotaka Kuwahara, and Mehdi Soleimanifar. Sample-efficient learning of interacting quantum systems. *Nature Physics*, May 2021. Preliminary version in FOCS 2020. arXiv:2004.07266, doi:10.1038/s41567-021-01232-0. [pp. 1, 3, 4, 5, 7, 33, 34, 49]

[AG04] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Phys. Rev. A*, 70:052328, Nov 2004. doi:10.1103/PhysRevA.70.052328. [p. 26]

[AKN06] Pieter Abbeel, Daphne Koller, and Andrew Y. Ng. Learning factor graphs in polynomial time and sample complexity. *Journal of Machine Learning Research*, 7(64):1743–1788, 2006. URL: http://jmlr.org/papers/v7/abbeel06a.html. [p. 2]

[BAL19] Eyal Bairey, Itai Arad, and Netanel H. Lindner. Learning a local Hamiltonian from local measurements. *Phys. Rev. Lett.*, 122:020504, Jan 2019. arXiv:1807.04564, doi:10.1103/PhysRevLett.122.020504. [pp. 2, 4, 6]

[BGP+20] Eyal Bairey, Chu Guo, Dario Poletti, Netanel H Lindner, and Itai Arad. Learning the dynamics of open quantum systems from their steady states. *New Journal of Physics*, 22(3):032001, mar 2020. doi:10.1088/1367-2630/ab73cd. [p. 2]

[BMS13] Guy Bresler, Elchanan Mossel, and Allan Sly. Reconstruction of Markov random fields from samples: Some observations and algorithms. *SIAM Journal on Computing*, 42(2):563–578, January 2013. doi:10.1137/100796029. [pp. 2, 3]

[BP12] Winton Brown and David Poulin. Quantum Markov networks and commuting Hamiltonians, 2012. arXiv:1206.0755v1. [p. 8]

[Bre15] Guy Bresler. Efficiently learning Ising models on arbitrary graphs. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing - STOC '15*. ACM Press, 2015. doi:10.1145/2746539.2746631. [pp. 2, 3, 50]

[CL68] C. Chow and C. Liu. Approximating discrete probability distributions with dependence trees. *IEEE Transactions on Information Theory*, 14(3):462–467, May 1968. doi:10.1109/tit.1968.1054142. [p. 3]

[Cli90] Peter Clifford. Markov random fields in statistics. In *Disorder in Physical Systems. A Volume in Honour of John M. Hammersley*. Clarendon Press, 1990. [p. 3]

[Dom96]   Cyril Domb. *The critical point.* Taylor & Francis, third edition, 1996. [p. 12]

[EHF19]   Tim J. Evans, Robin Harper, and Steven T. Flammia. Scalable Bayesian Hamiltonian learning, 2019. `arXiv:1912.07636v1`. [p. 2]

[FS09]   Philippe Flajolet and Robert Sedgewick. *Analytic Combinatorics.* Cambridge University Press, 2009. `doi:10.1017/cbo9780511801655`. [p. 17]

[GKP94]   Ronald Graham, Donald Knuth, and Oren Patashnik. *Concrete Mathematics: A Foundation for Computer Science.* Addison Wesley Pub Co Inc, 1994. [p. 16]

[HC71]   J.M. Hammersley and P. Clifford. Markov field on finite graphs and lattices. Available at `http://www.statslab.cam.ac.uk/~grg/books/hammfest/hamm-cliff.pdf`, 1971. [p. 8]

[HvdH21]   David Harvey and Joris van der Hoeven. Integer multiplication in time $O(n \log n)$. *Annals of Mathematics*, 193(2):563, 2021. `doi:10.4007/annals.2021.193.2.4`. [p. 27]

[JEMF06]   Ariel Jaimovich, Gal Elidan, Hanah Margalit, and Nir Friedman. Towards an integrated protein–protein interaction network: A relational markov network approach. *Journal of Computational Biology*, 13(2):145–164, mar 2006. `doi:10.1089/cmb.2006.13.145`. [p. 3]

[KF09]   Daphne Koller and Nir Friedman. *Probabilistic graphical models :Principles and techniques.* The MIT Press, Cambridge, Massachusetts, 2009. [p. 3]

[KKB20]   Tomotaka Kuwahara, Kohtaro Kato, and Fernando G. S. L. Brandão. Clustering of conditional mutual information for quantum Gibbs states above a threshold temperature. *Physical Review Letters*, 124(22), June 2020. `doi:10.1103/physrevlett.124.220601`. [pp. 4, 6, 7, 12]

[KM17]   Adam Klivans and Raghu Meka. Learning graphical models using multiplicative weights. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, October 2017. `doi:10.1109/focs.2017.39`. [pp. 2, 3, 4, 8, 51]

[KS80]   Ross Kindermann and J. Laurie Snell. *Markov random fields and their applications.* American Mathematical Society, Providence, R.I, 1980. `doi:10.1090/conm/001`. [p. 3]

[KS01]   David Karger and Nathan Srebro. Learning markov networks: Maximum bounded tree-width graphs. In *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '01, page 392–401, USA, 2001. Society for Industrial and Applied Mathematics. [p. 2]

[KS20]   Tomotaka Kuwahara and Keiji Saito. Gaussian concentration bound and ensemble equivalence in generic quantum many-body systems including long-range interaction. *Annals of Physics*, 421:168278, 2020. `arXiv:1906.10872`, `doi:10.1016/j.aop.2020.168278`. [pp. 6, 12, 20]

[Lau96]   Steffen Lauritzen. *Graphical models.* Clarendon Press Oxford University Press, Oxford New York, 1996. [p. 3]

[Li09]   Stan Z. Li. *Markov Random Field Modeling in Image Analysis.* Advances in Computer Vision and Pattern Recognition. Springer London, 2009. [p. 3]

[QR19]   Xiao-Liang Qi and Daniel Ranard. Determining a local Hamiltonian from a single eigenstate. *Quantum*, 3:159, July 2019. `doi:10.22331/q-2019-07-08-159`. [p. 2]

[Rud87]   Walter Rudin. *Real and Complex Analysis*. McGraw-Hill, third edition, 1987. [p. 10]

[SW12]   Narayana P. Santhanam and Martin J. Wainwright. Information-theoretic limits of selecting binary graphical models in high dimensions. *IEEE Transactions on Information Theory*, 58(7):4117–4134, July 2012. `doi:10.1109/tit.2012.2191659`. [pp. 2, 3, 4, 6]

[Tsy09]   Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer New York, 2009. `doi:10.1007/b13794`. [p. 41]

[VMLC16] Marc Vuffray, Sidhant Misra, Andrey Lokhov, and Michael Chertkov. Interaction screening: Efficient and sample-optimal learning of Ising models. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL: `https://proceedings.neurips.cc/paper/2016/file/861dc9bd7f4e7dd3cccd534d0ae2a2e9-Paper.pdf`. [pp. 2, 3, 4]

[WA22]   Dominik S. Wild and Álvaro M. Alhambra. Classical simulation of short-time quantum dynamics, 2022. `arXiv:2210.11490v1`. [pp. 12, 19, 20, 21, 23, 24]

[ZYLB21] Assaf Zubida, Elad Yitzhaki, Netanel H. Lindner, and Eyal Bairey. Optimal short-time measurements for hamiltonian learning, 2021. `arXiv:2108.08824`, `doi:10.48550/ARXIV.2108.08824`. [p. 6]