Marginal thermal trajectory reconstruction as an initial value problem

Sriram Gopalakrishnan
(Perimeter Institute & University of Waterloo)

Abstract

We study simple few-qubit Hamiltonians in the context of reconstructing the trajectory of their quantum thermal state's local marginals. We find that a general set of ordinary differential equations accurately model the evolution of local marginals as an initial value problem. The first four sections cover original work, while the last two sections (Appendix A and B) for the most part are reviews of the topics of Hamiltonian Learning from Gibbs states, and Shadow Tomography respectively.

1 Introduction

The thermal state of a quantum many-body system, also known as the Gibbs state [1], plays a fundamental role in our understanding of equilibrium statistical mechanics. In an engineered quantum system (a quantum computer, equivalently), the ability to prepare Gibbs states of humanly chosen Hamiltonians plays a crucial role both in benchmarking these devices, as well as in simulating quantum statistical mechanics on them.

At a fixed temperature, we say two quantum Gibbs states are equivalent if they are generated by Hamiltonians with the same matrix form.

Proposition 4 of ref. [2] points at an additional remarkable property of equivalent Gibbs states. Besides fixing temperature, suppose we also fix the *structure* of local interactions, while allowing freedom in the *strengths* of the interactions. If so, the authors can prove that two Gibbs states are equivalent (hence have equal interaction strengths) if and only if they have the same set of "local marginals". Local marginals are the expectation values of canonical observables (that fix the structure of local interactions) with respect to the Gibbs state.

The proposition motivates a powerful complementary idea. If both the structure and strengths of local interactions are held fixed, then the trajectory of the vector of local marginals *never* intersects or becomes periodic as temperature is swept. This is remarkable because it means that at least in principle - it should be possible to reconstruct the *entire* trajectory of *all* local marginals across *all* temperatures *uniquely* given a minimal set of boundary conditions [1]

It is important to keep in mind that in the context of engineered quantum systems, local marginals are not directly measurable. Estimating local marginals even at a single temperature requires

[1] We originally had this idea during early discussions of ref. [2] with Beni Yoshida. We were inspired to pursue it recently when learning about differential equation solvers in a Computational Physics class

rounds of randomized circuit sampling and post-processing, where a scheme like shadow tomography comes in handy [3]. While shadow tomography is efficient in its own right, we cannot rely on it alone if we require access to the marginals continuously across temperatures.

Since the Gibbs state is itself the solution of a Schrodinger-like equation, it is possible to write down a simple set of first-order differential equations for the evolution of the local marginals across temperature. Let us see how this works. Recall from Statistical Mechanics that Gibbs states of a Hamiltonian H can be defined as :

$$\rho_\beta(H) = \frac{e^{-\beta H}}{\mathrm{tr}(e^{-\beta H})} \tag{1}$$

where $\beta = 1/T$ is the inverse-temperature ($k_B = 1$). By differentiating the above equation w.r.t. β, we can show that :

$$\frac{d\rho}{d\beta} = -H\rho + \mathrm{tr}(H\rho)\rho \tag{2}$$

The differential equation above satisfies one trivial boundary condition given by: $\rho_{\beta=0} = \mathbb{1}/d$. This points to fact that infinite temperature Gibbs states are always maximally mixed. The remaining evolution for $\beta > 0$ is dependent on the specific Hamiltonian H.

A local marginal of a Gibbs state at inverse temperature β is defined as :

$$\langle E_j \rangle_\beta = \mathrm{tr}(E_j \rho_\beta) \tag{3}$$

where $\{E_j : j \in [m]\}$ are the set of m canonical observables fixing the structure of local interactions.

Then, using (2) in (3), one can note :

$$\frac{d}{d\beta}\langle E_j \rangle_\beta = \mathrm{tr}\left(E_j \frac{d\rho}{d\beta}\right) \tag{4}$$

$$= -\langle E_j H \rangle_\beta + \langle H \rangle_\beta \langle E_j \rangle_\beta \tag{5}$$

Now, it is crucial in this study that the canonical observables $\{E_j\}$ fix the structure of local interactions. This means that the Hamiltonian can be expanded in this set as :

$$H = \sum_{j=1}^{m} \mu_j E_j \tag{6}$$

where μ_j are essentially the interaction strengths of the model. It follows that :

$$\frac{d}{d\beta}\langle E_j \rangle_\beta = -\sum_{k=1}^{m} \mu_k (\langle E_j E_k \rangle_\beta - \langle E_j \rangle_\beta \langle E_k \rangle_\beta) \tag{7}$$

$$= -\sum_{k=1}^{m} \mu_k \langle\!\langle E_j, E_k \rangle\!\rangle_\beta \tag{8}$$

In general, a term like $\langle\!\langle E_j, E_k \rangle\!\rangle_\beta$ represents a long-range thermal correlation function which is not directly expressible in terms of local marginals $\langle E_j \rangle_\beta$. The double angled bracket notation $\langle\!\langle \cdot, \cdot \rangle\!\rangle$ is commonly used for cumulants of a random variable (single or multivariate). In the previous

equation in particular, it points to a covariance, albeit an "asymmetric covariance" in general due to the non-commutativity of quantum observables.

For simplicity, and as a proof-of-concept, we'd like to work in a setting where H as a matrix is computable explicitly and given, so are the canonical observables $\{E_j : j \in [m]\}$. We then try to simulate the following system of coupled ODEs :

$$\boxed{\frac{d}{d\beta}\langle E_j\rangle_\beta = -\langle E_j H\rangle_\beta + \langle H\rangle_\beta \langle E_j\rangle_\beta \qquad j \in 1, 2 \dots m} \tag{9}$$

along with an extra set of ODEs for $\langle H\rangle$ and $\langle E_j H\rangle$ which we will get to later.

The RHS in (9) is certainly not an efficient computation as the system-size scales up. Rather, what we are really interested in is the following question: Is it possible to reconstruct the entire thermal trajectory of $[\langle E_1\rangle_\beta \ \langle E_2\rangle_\beta \ \dots \ \langle E_m\rangle_\beta]$ from very few boundary conditions?

When we say "very few boundary conditions", we mean one of two situations: i)We have access to *all* local marginals at *few* temperatures, or ii)We have access to *few* [perhaps even just a single] local marginal(s) at *many* [at most m] temperatures. While the former resembles a typical initial value problem (IVP), the latter is reminiscent of the more interesting boundary value problems (BVP) in the study of ordinary differential equations. In this work, we are focused on the IVP setting, while remarking about BVPs in the discussion section.

2 Model Hamiltonians

Let us start by considering a very simple two-qubit Hamiltonian

$$H = JX_1X_2 \tag{10}$$

where X_1 and X_2 are Pauli X operators. It easily diagonalized in the $\{|+\rangle, |-\rangle\}$ basis corresponding to the eigenstates of X as :

$$H = J\Big(|{+}{+}\rangle\langle{+}{+}| + |{-}{-}\rangle\langle{-}{-}| \quad - |{+}{-}\rangle\langle{+}{-}| - |{-}{+}\rangle\langle{-}{+}| \Big) \tag{11}$$

It follows that the partition function of the model is given by :

$$Z(\beta) = \mathrm{tr}(e^{-\beta J X_1X_2}) = 2(e^{-\beta J} + e^{\beta J}) = 4\cosh(\beta J) \tag{12}$$

In this simple case, the Gibbs state can be easily expressed as a diagonal matrix in the X basis as follows :

$$\rho_\beta(J) = \frac{e^{-\beta J X_1X_2}}{\mathrm{tr}(e^{-\beta J X_1X_2})} \tag{13}$$

$$= \frac{e^{-\beta J}}{4\cosh(\beta J)}(|{+}{+}\rangle\langle{+}{+}| + |{-}{-}\rangle\langle{-}{-}|) + \frac{e^{\beta J}}{4\cosh(\beta J)}(|{+}{-}\rangle\langle{+}{-}| + |{-}{+}\rangle\langle{-}{+}|) \tag{14}$$

The only local marginal in this case is simply $\langle X_1X_2\rangle_\beta$, and is given by :

$$\langle X_1X_2\rangle_\beta = \mathrm{tr}(\rho_\beta(J)X_1X_2) \tag{15}$$

Since $\langle +|X|+\rangle = 1$ and $\langle -|X|-\rangle = -1$, we can easily note :

$$\langle X_1 X_2 \rangle_\beta = \frac{2(e^{-\beta J} - e^{\beta J})}{4\cosh(\beta J)} = -\tanh(\beta J) \tag{16}$$

Indeed we have

$$\frac{d}{d\beta}\langle X_1 X_2 \rangle = -\langle X_1 X_2 H \rangle + \langle H \rangle \langle X_1 X_2 \rangle \tag{17}$$

$$= -J + J\langle X_1 X_2 \rangle^2 \tag{18}$$

$$= -J + J\tanh^2(\beta J) = -J\mathrm{sech}^2(\beta J) \tag{19}$$

which is consistent with differentiating (16) directly.

So in this case, we would like to compare a numerical simulation of

$$\frac{d}{d\beta}\langle X_1 X_2 \rangle = -J + J\langle X_1 X_2 \rangle^2 \tag{20}$$

with the analytical solution in (16).

Let us now consider a three qubit Hamiltonian given by

$$H = J_x X_1 X_2 + J_z Z_1 Z_2 \tag{21}$$

The two terms are non-commuting, and even if there exists an analytical partition function, it is not likely to be simple. Nevertheless, we can say that any diagonalization of H should ultimately effect the following computations :

$$\langle X_1 X_2 \rangle = \langle X_1 X_2 e^{-\beta H} \rangle / \operatorname{tr}(e^{-\beta H}) \tag{22}$$

$$\langle Z_2 Z_3 \rangle = \langle Z_2 Z_3 e^{-\beta H} \rangle / \operatorname{tr}(e^{-\beta H}) \tag{23}$$

On the other hand, the differential equations representing the marginal evolution as in (9), are given by :

$$\frac{d}{d\beta}\langle X_1 X_2 \rangle = -\langle X_1 X_2 H \rangle + \langle H \rangle \langle X_1 X_2 \rangle \tag{24}$$

$$\frac{d}{d\beta}\langle Z_2 Z_3 \rangle = -\langle Z_2 Z_3 H \rangle + \langle H \rangle \langle Z_2 Z_3 \rangle \tag{25}$$

We will need additional equations for the evolution of $\langle H \rangle$, $\langle X_1 X_2 H \rangle$, $\langle Z_2 Z_3 H \rangle$ in order to bring the system of ODEs to a form amenable to an algorithm like RK4.

Now, here is where a crucial property of Paulis turns out to be useful : while they certainly don't commute, they do *anti-commute*. In our case, it can be seen that $\{X_1 X_2, Z_2 Z_3\} = 0$ [ERROR FIX]. A consequence of anti-commutativity is that $H^2 = (J_x^2 + J_z^2)\mathbb{1}$, and therefore the thermal expectation of the square of the Hamiltonian is explicitly known (and is a constant) :

$$\langle H^2 \rangle = \langle (J_x^2 + J_z^2)\mathbb{1} + J_x J_z \{X_1 X_2, Z_2 Z_3\} \rangle = J_x^2 + J_z^2 \tag{26}$$

One can therefore write down the remaining three ODEs as follows:

$$\frac{d}{d\beta}\langle H\rangle = -J_x^2 - J_z^2 + \langle H\rangle^2 \tag{27}$$

$$\frac{d}{d\beta}\langle X_1X_2H\rangle = -(J_x^2 + J_z^2)\langle X_1X_2\rangle + \langle H\rangle\langle X_1X_2H\rangle \tag{28}$$

$$\frac{d}{d\beta}\langle Z_2Z_3H\rangle = -(J_x^2 + J_z^2)\langle Z_2Z_3\rangle + \langle H\rangle\langle Z_2Z_3H\rangle \tag{29}$$

So, what we essentially want to do is to simulate the system of five ODEs given by equations (24),(25),(27)-(29) [theory], in order to compare with the simulated analytical solution of the marginals given by equations (22),(23) [truth].

The objective of such a study would be to understand how many points are needed in an RK4 evolution of the ODEs, what initial conditions work well, and whether we can formulate a boundary value problem to determine the optimal initial conditions. Within time constraints, we could only investigate the initial value problem setting with hand-picked initial guesses based on analytical expectations.

3 Numerical simulations

To recall, the single qubit case in the previous section corresponds to evolving :

$$\frac{d}{d\beta}\langle X_1X_2\rangle = -J + J\langle X_1X_2\rangle^2 \tag{30}$$

This is a single uncoupled ODE. The true solution we saw was $\langle X_1X_2\rangle = -\tanh(\beta J)$.

Without loss of generality, we take $J = 10$, and we use the RK4 algorithm (implementation due to Alex Gezerlis [4]) with initial condition $\langle X_1X_2\rangle_{\beta=0.1} = -0.76$ from $\beta = 0.1$ to $\beta = 0.9$. We find good agreement with just over 15 points in the chosen range (Figure 1).

What do we learn from this simple two-qubit example? Given the generality of our introductory discussion, quite a bit actually. First, we learn that for any Hamiltonian defined by a single local marginal, having experimental access to the marginal even at a single temperature allows us to reconstruct the entire evolution of the marginal across temperatures by setting up an ODE according to the general formalism of Equation (9). Second, for different values of J, the trajectory of $\langle X_1X_2\rangle$ is characterized by its slope when $\langle X_1X_2\rangle = 0$: the slope is $-J$. This means even if we didn't know J, we can find J numerically by root-finding. And lastly, if we knew J but not the right initial condition, we can find that too by root-finding.

Now, let us look at the three-qubit case. This time, we have a set of five coupled ODEs to simulate,

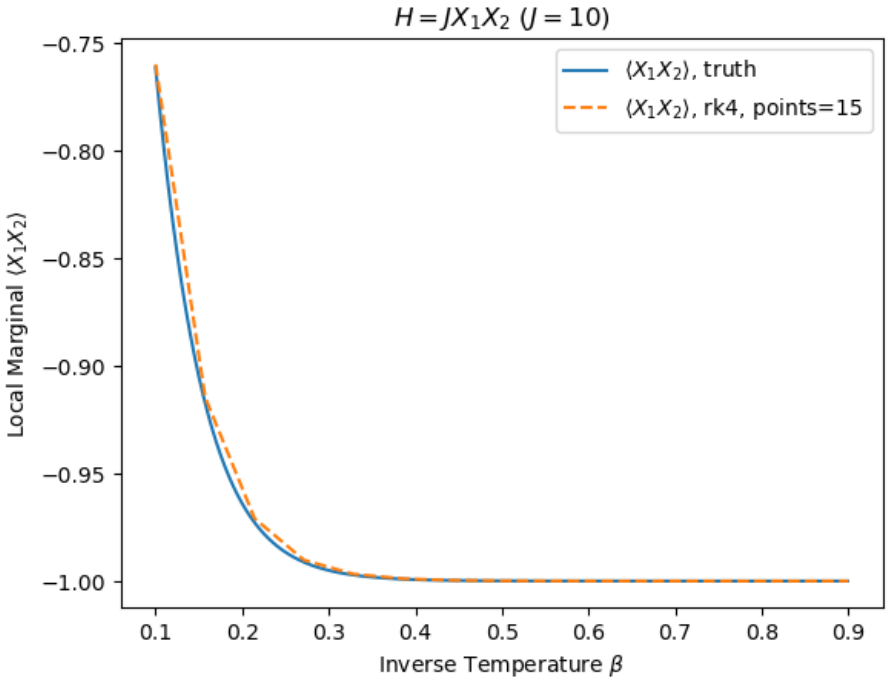

Figure 1: The two-qubit example

which we re-print from the previous section :

$$\frac{d}{d\beta}\langle X_1X_2\rangle = -\langle X_1X_2H\rangle + \langle H\rangle\langle X_1X_2\rangle \tag{31}$$

$$\frac{d}{d\beta}\langle Z_2Z_3\rangle = -\langle Z_2Z_3H\rangle + \langle H\rangle\langle Z_2Z_3\rangle \tag{32}$$

$$\frac{d}{d\beta}\langle H\rangle = -J_x^2 - J_z^2 + \langle H\rangle^2 \tag{33}$$

$$\frac{d}{d\beta}\langle X_1X_2H\rangle = -(J_x^2 + J_z^2)\langle X_1X_2\rangle + \langle H\rangle\langle X_1X_2H\rangle \tag{34}$$

$$\frac{d}{d\beta}\langle Z_2Z_3H\rangle = -(J_x^2 + J_z^2)\langle Z_2Z_3\rangle + \langle H\rangle\langle Z_2Z_3H\rangle \tag{35}$$

Without loss of generality, we take $J_x = 2$ and $J_z = 3$. This time, we don't have an explicit ground truth. However, as we noted earlier, solving Equations (22), (23) with high numerical precision is "like" ground truth. We do so using `numpy` functions and a fine discretization of 0.001 from $\beta = 0.1$ to $\beta = 0.9$.

On the other hand, solving the five coupled ODEs above requires an implementation of the generalized RK4 algorithm (again due to Alex Gezerlis). We use the earlier solution to set the initial conditions for RK4 as $[\langle X_1X_2\rangle, \langle Z_2Z_3\rangle, \langle H\rangle, \langle X_1X_2H\rangle, \langle Z_2Z_3H\rangle]_{\beta=0.1} = [0.0, -0.3, -0.9, 0.1, 3.0]$. Since it gets cluttered to have five plots together, we separate out the plots for relevant marginals $\langle X_1X_2\rangle$ and $\langle Z_2X_3\rangle$ from the other "irrelevant" marginals.

This time (Figure 2), RK4 gets the general dependence correctly, however deviates asymptotically for two of the expectation values: $\langle Z_2Z_3\rangle$ and $\langle H\rangle$ (the total average energy), even though we

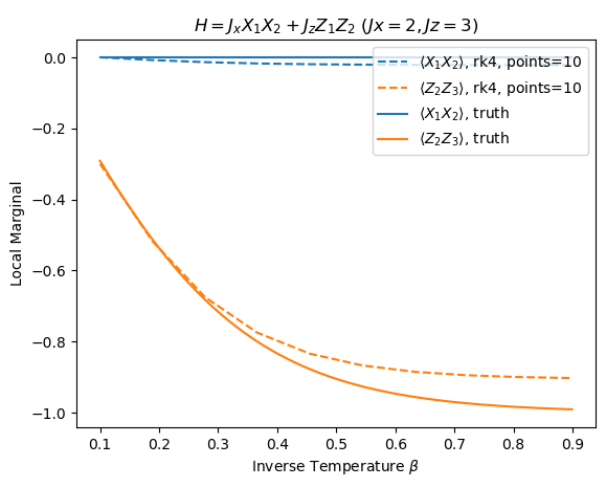

(a) the relevant marginals

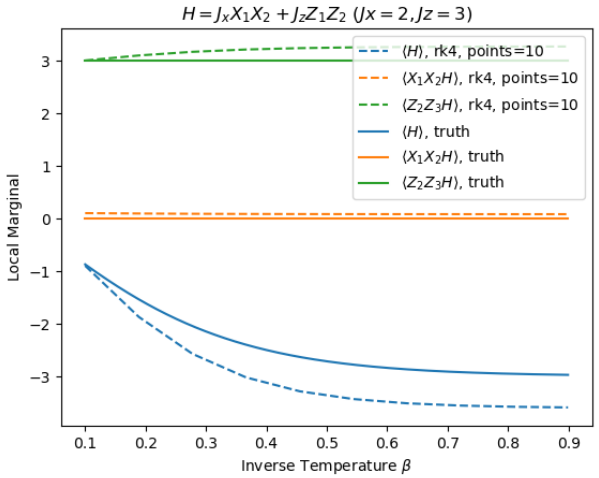

(b) the "irrelevant" marginals

Figure 2: The three-qubit example

supply accurate initial values. We needed just over 10 points in the chosen range. Supplying a lot higher points did not change the asymptotic deviation of the RK4 solution.

It just so happens in this case that $\langle X_1X_2\rangle = 0$ for all temperatures. That however does not diminish the non-linearity of our coupled ODEs. Presumably, an example where the relevant local marginals are never flat zero may circumvent the asymptotic deviation, but we leave this for future exploration. We include a copy of the Python codes used at the end of this essay.

4 Discussion

Although our numerics are limited to simple two and three qubit Hamiltonians, we did so to demonstrate the power of a general theory. The formalism developed in Equation (9) is quite general, and we've managed to verify the fact that evolving that system of equations with a standard intial-value based ODE algorithm like RK4 does reproduce the correct dependence of the local marginals on temperature expected analytically. Besides, once we scale up system size, we may not have easy access to diagonalizable models for verification.

It is worth remarking that the finiteness of the system of ODEs in (9) did depend on the *anti-commutativity* of the canonical observables defining the local structure of interactions, as is normally the case with Pauli operators. The anti-commutativity argument actually does not directly generalize to $n > 3$ qubits, however, this is something we'd like to explore in the near future since we expect tricks of a similar nature to help constrain the system of ODEs.

Finally, it would be very interesting to model the evolution of local marginals with temperature as a *boundary value problem* in order to discover the optimal initial conditions that we used "by hand" in our numerical studies. This can be used for example to limit our data access to fewer local marginals while allowing measurements at multiple temperatures.

Separately, in the past, we explored (incompletely) two related problems: a) *Hamiltonian Learning from Gibbs states*: Suppose we had access to the marginals at a fixed temperature with the interaction strengths *unknown* but the structure of local interactions known. Then is it possible to

efficiently reconstruct the Hamiltonian parameters? b) *Shadow Tomography*: Given access to many copies of an unknown many-qubit state (in general mixed), what is the best sample complexity of determining the expectation values of a set of local observables? We include our notes on these topics in the appendices, but remark that these not structured to directly supplement main content until this section.

5 Acknowledgement

We thank Beni Yoshida and Timothy Hsieh for helpful discussions at different stages of this work. We thank Alexandros Gezerlis for an excellent class on Computational Physics. We thank Perimeter Institute and the University of Waterloo for their kind hospitality and supporting our work.

6 Appendix A: Hamiltonian Learning from Gibbs states

In reference [2], AAKS rigorously proved that the coefficients of very general quantum local Hamiltonians can be learnt to arbitrary precision given polynomially many copies of the system's Gibbs states on an experimental device. In subsequent work by Haah *et. al.* [5], optimal sample and time complexity bounds were obtained in a high-temperature regime. More recently, Bakshi *et. al.* [6] developed a sample and time efficient protocol for arbitrary temperatures. While being of general interest to a broader scientific community, this avenue of research poses interesting problems from complexity-theoretic and physics perspectives [7].

1.1 An important result in AAKS was the strong convexity of quantum log-partition functions. Does the strong convexity (α) of log-partition function have a physical meaning? For commuting Hamiltonians, a simple calculation yields that the largest lower bound on α attainable by any strategy is the minimum variance of the Hamiltonian over a unit sphere in parameter space. The generalization to non-commuting Hamiltonians is complicated, however we take preliminary steps in the right direction.

1.2 In a lengthy derivation, AAKS obtained an $O(1/n)$ lower bound on α at finite temperature for non-commuting Hamiltonians. Unfortunately, this bound would vanish for sufficiently large system-size, thereby badly affecting complexity guarantees as well as sensitivity to noise. Is there a larger bound on α, ideally independent of system size (n)? Does this lead to a sample complexity linear in n? How does it improve time complexity and sensitivity to noise?

Motivated by these questions, we explored the complexity of Hamiltonian Learning with some generality. Our main technical contribution was a factor-of-2 improvement in the sensitivity to noise bound in AAKS.

6.1 Setup and Notation

We are interested in learning from Gibbs states of local Hamiltonians prepared on synthetic many-qubit devices. It is worth noting that although the Gibbs state is a matrix of the same dimension as the Hamiltonian, it has qualitatively different tensor structure than the Hamiltonian itself. There are quantum algorithms capable of preparing the Gibbs state on synthetic devices [8, 9]. Such algorithms rely on the ability to approximate arbitrary smooth functions of a Hermitian matrix by a linear combination of unitaries (LCU) amenable to circuit constructions [10, 11].

Concretely, we are interested in probing copies of Gibbs states of κ-local lattice Hamiltonians H on n-qubits, with $m = O(n)$ local terms that are: mutually orthogonal, traceless, and in general non-commuting (the last criterion being a key hurdle). For simplicity, we assume that local operators have exactly κ-sized support, as opposed to usual studies that allow $\leq \kappa$ support. For convenience of thinking in the energy basis, we assume that the spectrum of H and its local basis are gapped and non-degenerate. Generalizations may be non-trivial, left for future work.

With the above mental picture, we begin our analysis with some basic definitions.

Definition 1 (qubit Hilbert space)**.** *The state space of a single-qubit is a two-dimensional complex euclidean space [c.e.s aka Hilbert space] denoted* $\mathcal{X} = \mathbb{C}^2$.

The state space of n-qubits is $\mathcal{X}^{\otimes n}$. A valid basis for $\mathcal{X}^{\otimes n}$ is $\{|0\rangle, |1\rangle, |2\rangle \cdots |2^n - 1\rangle\}$ when we allow complex linear combinations. We also call it the *standard basis*.

Definition 2 (graph)**.** *A graph denoted* $G = (V, E)$ *is an indexed set of vertices* $V(G)$*, and a set of edges* $E(G)$*. Each edge is a pair of vertices.*

For n-qubits placed on the vertices of a graph, $|V| = n$. The graph basically restricts the allowed set of gates in a quantum computation. Detail: Our graph is finite, undirected, and loopless. Any finite graph has an embedding in $\mathbb{R}^3$, and we restrict visualization to this basic embedding. Next we define a special graph - the lattice

Definition 3 (lattice)**.** *A lattice denoted* $\Lambda = (V, E)$ *is a sparse connected graph with constant interior-degree. Each interior vertex has a constant neighborhood:* $|E(v)| = c \quad \forall v \in int(V)$*. This means* $|E| = O(n)$ *for an n-qubit lattice. We interpret the neighborhood induced by any vertex as being* ***graphically local****.*

Detail: Degree can change at the boundary of Λ - denoted $\partial\Lambda$. Consequently, $\Lambda = \text{int}(\Lambda) \cup \partial\Lambda$. All constant-degree graphs are sparse, but not vice versa.

Observe that our notion of a lattice only requires graphical locality, going beyond its subset - geometric locality. There is nothing "unphysical" about graphical locality in the context of synthetic devices (See for example [12] that realizes long-range-connectivity in a ring-resonator architecture for superconducting qubits). Small changes in euclidean distances can change the interaction strength but **not** the connectivity graph itself. In this lens, the locality of a QC is robust to deformations that preserve the connectivity graph. 2-local or larger measurements are allowed with graphically local qubits even when they are geometrically distant. As such, we will freely use the phrase "spatially contiguous qubits" in the sense of both graphical and geometric locality. We note that Haah *et. al.*'s [5] intro comment about constant-degree expander graphs being low-intersection (graphical) but non-local (geometric) is sensible, however our definition also counts these graphs as lattices, since "distance" is generally well understood purely in terms of number of edges [13]. Perhaps Lieb-Robinson bounds may not hold, but we aren't worried about it at this point.

Definition 4 (κ-local operator (LO))**.** *A Hermitian matrix* O *is called a* κ*-LO if it acts non-trivially on* κ *spatially contiguous qubits. Its tensor structure is supposed to be understood via the operator space:* $\text{Herm}(\mathcal{X}^{\otimes\kappa}) \otimes \mathbb{1}^{\otimes(n-\kappa)}$*. Its support set has size* κ*:* $|\text{Supp}(O)| = \kappa$

Each element of $\left(\text{Herm}(\mathcal{X}^{\otimes\kappa}) \otimes \mathbb{1}^{\otimes(n-\kappa)}\right)$ is a κ-LO that acts non-trivially on a chosen subset $V_0 \subset V(\Lambda)$ of κ spatially contiguous qubits. A note on our notation for Identity operators: when we use $\mathbb{1}$ without a subscript, it is a single-qubit Identity operator (in $\mathbb{C}^{2\times2}$). When we do use a

subscript like in $\mathbb{1}_m$, m is the matrix dimension of I. For example, $\mathbb{1} \equiv \mathbb{1}_2$ and $\mathbb{1}^{\otimes 2} \equiv \mathbb{1}_4$. Notation such as $\mathrm{D}(\cdot)$, $\mathrm{Pos}(\cdot)$, $\mathrm{Herm}(\cdot)$ for operator spaces and $\mathcal{X}$, $\mathcal{Y}$, $\mathcal{Z}$ for Hilbert spaces are adapted from Watrous [14]. $\mathrm{Conv}(\cdot)$ denotes convex hull. $\mathrm{Spec}(\cdot)$ denotes spectrum.

Worth refreshing some facts (good mental exercise, nothing more):

(i) $\mathrm{D}(\mathcal{X}) \subset \mathrm{Pos}(\mathcal{X}) \subset \mathrm{Herm}(\mathcal{X}) \subset \mathrm{Normal}(\mathcal{X})$. Other relevant subspaces of $\mathrm{Normal}(\mathcal{X})$ are $\mathrm{U}(\mathcal{X})$ and $\mathrm{Proj}(\mathcal{X})$. Operator $\hat{O} \in \mathrm{Normal}(\mathcal{X})$ *iff* it permits a unitary diagonalization.

(ii) If we choose $f(\theta) = e^{i\theta}$ for $\theta \in \mathbb{R}$, then $f(\mathrm{Herm}(\mathcal{X})) = U(\mathcal{X})$.

(iii) $D(\mathcal{X}^{\otimes n})$ is the union of a continuum of standard $(2^n - 1)$-simplices $\mathrm{Conv}(\{|v\rangle\langle v|\})$ where $\{|v\rangle\}$ is a chosen ON basis in $\mathcal{X}$. A continuum of ON bases in $\mathcal{X}$ are generated by unitary transforms on $\{|v\rangle\}$.

(iv) An n-qubit "quantum computation" $U_0 \in U(\mathcal{X}^{\otimes n})$ does two things as a map:

- It leaves the unit-sphere in $\mathcal{X}^{\otimes n}$ invariant while "moving patches on its surface". The "fixed points" on the unit-sphere are eigenstates of U_0.
- It leaves $D(\mathcal{X}^{\otimes n})$ invariant while "moving one simplex to another". There is exactly one "fixed simplex", which is the one spanned by the eigenbasis of U_0.

Definition 5 (κ-local Hamiltonian (LH))**.** *Given an n-qubit lattice $\Lambda = (V, E)$, a κ-LH on it is a sum of fixed-size LOs supported on different neighborhoods -*

$$H(x) = \sum_{j=1}^{m} x_j L_j \quad \in \mathrm{Herm}(\mathcal{X}^{\otimes n}) \tag{36}$$

where $m = O(n)$, parameters $x = (x_1\ x_2\ \ldots\ x_m)^\intercal \in \mathbb{R}^m$, and $\{L_j\}$ are basis LOs satisfying some important constraints -

$$\begin{array}{rll} \textit{κ-local:} & |\mathrm{Supp}(L_j)| = \kappa & j \in [m] \\ \textit{mutually orthogonal:} & \mathrm{tr}(L_i L_j) = 2^n \delta_{ij} & i, j \in [m] \\ \textit{non-commuting in general:} & [L_i, L_j] \neq 0 \quad \textit{whenever} \quad \mathrm{Supp}(L_i) \cap \mathrm{Supp}(L_j) \neq \varnothing & i, j \in [m] \\ \textit{Hermitian \& traceless:} & L_j^\dagger = L_j \quad \& \quad \mathrm{tr}(L_j) = 0 & j \in [m] \end{array}$$

All the above properties are characteristic of many local Pauli bases, however in an generalized view, the properties themselves matter more than a specific choice of Pauli basis.

Recall: Any complex Hermitian H is characterized by a spectral decomposition $H = \sum_{j=1}^{2^n} \lambda_j \Pi_j$ where λ_j are *real* and $\Pi_j \in \mathrm{Proj}(\mathcal{X}^{\otimes n})$ are *complex* valued. λ_j can be positive, negative or zero.

Definition 6 (Gibbs state of a LH)**.** *The Gibbs state ($\beta > 0$) of a κ-LH $H(x) \in \mathrm{Herm}(\mathcal{X}^{\otimes n})$ denoted $\rho(x)$ is defined as*

$$\rho(x) = \frac{e^{-\beta H(x)}}{\mathrm{tr}(e^{-\beta H(x)})} \quad \in \mathrm{D}(\mathcal{X}^{\otimes n}) \tag{37}$$

Henceforth $\rho(\cdot)$ is used exclusively for Gibbs states. Interpretations:

- In the energy basis, the Gibbs state represents a Boltzmann distribution over the eigenstates of H familiar from stat-mech. $Z(x) = \mathrm{tr}(e^{-\beta H(x)})$ is referred to as partition function.
- We have $Z(x) > 0$ always. If $H(x)$ has even one $\lambda \leq 0$, then $Z(x) \geq 1$. Hence, low-energy states ramp-up $Z(x)$, while high-energy states ramp-down $Z(x)$.

- Gibbs states are not just positive semi-definite, but positive definite. Hence they have a full Image and empty Kernel: $\mathrm{Im}(\rho(x)) = \mathcal{X}^{\otimes n}$ while $\mathrm{Ker}(\rho(x)) = \varnothing$.

However, note that $e^{-\beta\lambda_H}$ can be arbitrarily small under appropriate limits. Suppose we define an "ϵ-approximate Kernel" of $\rho \in \mathrm{D}(\mathcal{X}^{\otimes n})$ as follows:

$$\mathrm{Ker}_\epsilon(\rho) = \{|v\rangle \in \mathcal{X}^{\otimes n} : \|\rho|v\rangle\|_2 \leq \epsilon\} \tag{38}$$

Then even for exponentially small ϵ, $\mathrm{Ker}_\epsilon(\rho(x))$ is non-empty for some choices of $H(x)$. For our problem, it is desirable that Gibbs states have *empty approximate Kernels* as well. Non-empty approximate Kernels of $\rho(x)$ are likely when $\lambda_H \gg k_B T$, that is when $H(x)$ has high-energy states that far surpass the thermal energy scale. Hence we desire that $\mathbf{T} > \lambda^*/\mathbf{k_B}$ [where $\lambda^* = \max_{\mathbf{x}} \|\mathbf{H(x)}\|_\infty$] in an experiment for accurate Hamiltonian learning. Full-rank Gibbs states $\rho(x)$ with empty approximate Kernels favor lower sensitivity to noise in learning from their marginals.

We propose a distinction between *stable recovery* and *unique recovery* in the learning problem. One does not guarantee the other. Unique recovery is self-explanatory: the inverse map must be injective. Stable recovery means: reconstructions are robust against small perturbations/noise in their pre-image. In QHLP, this means noise in local marginals must not get amplified in what we learn: LH parameters. It appears that very-low-temperature states ($\beta \to 0$) forbid stable recovery, while very-high-temperature states ($\beta \to \infty$) forbid unique recovery. We will see this heuristically in the proof of Lemma 8 and rigorously in Lemma 14

Problem (Quantum Hamiltonian Learning Problem (QHLP))**.** *Consider characters Alice and Bob. Given an n-qubit lattice $\Lambda = (V, E)$, Alice chooses a κ-LH (strictly obeying Definition 5)*

$$H(\mu) = \sum_{j=1}^{m} \mu_j L_j \qquad \mu \in \mathbb{R}^m$$

with basis $\{L_j\}$ known to Bob but coefficients $\{\mu_j\}$ unknown to him. She tells him: I can give you polynomially many copies of $H(\mu)$'s Gibbs state $\rho(\mu)$ on a QC, can you learn $\{\mu_j\}$ to ϵ-precision in 2-norm? What is your sample and time complexity? Formally, we require Bob to learn an estimate $\hat{\mu}$ of μ such that $\|\hat{\mu} - \mu\|_2 \leq \epsilon$, for arbitrary $\epsilon > 0$ provided by Alice.

Remarkably, AAKS provided a rigorous solution to QHLP with a polynomial sample complexity [2]. As we discussed earlier, they did not investigate time complexity, however, in a high-temperature regime, sample and time optimal bounds utilizing cluster expansions are now known [5].

6.2 Existence of bijective map

The foundational principle underlying QHLP is the existence of a bijective map between the space of LHs and the space of local marginals induced by their Gibbs states. We define local marginals in Definition 7.

That such a map exists for quantum LHs should point to a fundamental info-theoretic property, a possible manifestation of short-range-entanglement (SRE). We believe SRE of quantum Gibbs states is in direct analogy to conditional-independence (CI) of **classical** Gibbs distributions on graphs. CI in any multivariate probability distribution manifests as a simple yet fundamental info-theoretic property: $I(A:C|B) = 0$ (zero CMI) for all tri-partitions (A, B, C) of Λ that have **no**

direct edges from A to C. However, we are not aware of an equivalent info-theoretic manifestation of SRE, since $I(A:C|B) \neq 0$ (non-zero QCMI) in general for similar tri-partitions of Λ for quantum Gibbs states. Intuitively, zero QCMI means that for every pair of vertices $\{v_1, v_2\}$ with no direct edge, the reduced density matrix $\rho_{v_1 v_2}$ is separable. Even without zero QCMI, quantum Gibbs states are locally reconstructible.

Definition 7 (κ-Local Marginal (LM))**.** *For each local basis operator L_j, there is an associated LM $l_j(\mu) \in \mathbb{R}$, given by*

$$l_j(\mu) = \mathrm{tr}(\rho(\mu) L_j) \equiv \langle L_j \rangle_\mu \tag{39}$$

These LMs can be stacked into a column vector like: $l(\mu) = (l_1(\mu)\ l_2(\mu)\ \ldots l_m(\mu))^T \in \mathbb{R}^m$.

Lemma 8 (Bijective map from LH to LMs)**.** *The non-linear mapping $\mathcal{T} : \mathbb{R}^m \to \mathbb{R}^m$ that takes LH parameters μ to LMs $l(\mu)$ is bijective. Notation: $\mathcal{T}[\mu] = l(\mu)$, and $\mathcal{T}^{-1}[l(\mu)] = \mu$*

Proof. Since the domain and co-domain of $\mathcal{T}$ are the same set, if $\mathcal{T}$ is injective, it is also surjective. Hence, it suffices to prove that $\mathcal{T}$ is injective. Injectivity requires us to prove: For $\mu, \lambda \in \mathbb{R}^m$, $\mu \neq \lambda \Rightarrow l(\mu) \neq l(\lambda)$. Equivalently, we require: $l(\mu) = l(\lambda) \Rightarrow \mu = \lambda$ ("consistent marginals implies identical LHs")

Claim 8.1: The von-Neumann entropy of a Gibbs state $\rho(\lambda)$ is given by

$$S(\rho(\lambda)) = \beta \langle H(\lambda) \rangle_\lambda + \log Z(\lambda) \tag{40}$$

Notice how it is consistent with the thermodynamic relation $F = E - TS$ if $\log Z = -\beta F$.

Claim 8.2: The relative entropy of $\rho(\mu)$ w.r.t. $\rho(\lambda)$ is given by

$$S(\rho(\mu) \| \rho(\lambda)) = \beta [\langle H(\lambda) \rangle_\mu - \langle H(\mu) \rangle_\mu] + [\log Z(\lambda) - \log Z(\mu)] \tag{41}$$

Both the above claims follow straightforwardly from the definition of quantum entropies applied to the Gibbs state. Now consider the following chain of thought

$$\begin{aligned}
l(\mu) = l(\lambda) &\implies \langle L_j \rangle_\mu = \langle L_j \rangle_\lambda \quad \forall j \in [m] && \text{(elementwise comparison)} \\
&\implies \langle H(x) \rangle_\mu = \langle H(x) \rangle_\lambda \quad \forall x \in \mathbb{R}^m && \text{(linearity of expectation)} \\
&\implies \langle H(\lambda) \rangle_\mu = \langle H(\lambda) \rangle_\lambda && \text{(choosing } x = \lambda\text{)} \\
&\implies S(\rho(\mu) \| \rho(\lambda)) = S(\rho(\lambda)) - S(\rho(\mu)) && \text{(substituting into eq 41)} \\
&\implies S(\rho(\lambda)) \geq S(\rho(\mu)) && \text{(non-negativity of relative entropy)}
\end{aligned}$$

Recall that the relative entropy is an asymmetric yet non-negative function. Had we swapped the arguments of $S(\cdot\|\cdot)$ in the same chain of thought, we would get $S(\rho(\mu)) \geq S(\rho(\lambda))$. Hence $S(\rho(\mu)) = S(\rho(\lambda))$, implies the relative entropy of two Gibbs states with consistent marginals is exactly zero, implies the Gibbs states would have to be identical to begin with:

$$l(\mu) = l(\lambda) \Rightarrow S(\rho(\mu)) = S(\rho(\lambda)) \Rightarrow S(\rho(\mu) \| \rho(\lambda)) = 0 \Rightarrow \rho(\mu) = \rho(\lambda) \tag{42}$$

The subtlety in this argument is that $S(\rho\|\sigma) = \mathrm{tr}(\rho \log \rho - \rho \log \sigma)$ holds only when $\mathrm{Im}(\rho) \subseteq \mathrm{Im}(\sigma)$ by definition. Our ability to swap ρ and σ and still use the same formula is strictly contingent on $\mathrm{Im}(\rho) = \mathrm{Im}(\sigma)$. Hence, $\rho(\mu)$ **and** $\rho(\lambda)$ **must have identical Kernel and Image spaces** for equation 42 to be considered valid.

Lucky for us, the *exact* Kernel of all finite-temp Gibbs states are empty: $\mathrm{Ker}(\rho(x)) = \varnothing \quad \forall x \in \mathbb{R}^m$. But we have also seen in the last section that *approximate* Kernels may be non-empty if $H(x)$ has high-energy states that surpass the thermal energy scale $k_B T$. We can preclude this possibility by assuming a sufficiently large temperature $\mathbf{T} > \lambda^*/\mathbf{k_B}$ where $\lambda^* = \max_x \lambda_{max}(H(x))$

Aside: We expect sensitivity $\frac{\|\rho(\mu)-\rho(\lambda)\|_2}{\|l(\mu)-l(\lambda)\|_2}$ to be an increasing function of the constant $\beta\lambda^*$.

Given that, we now have to prove: $\rho(\mu) = \rho(\lambda) \Rightarrow \mu = \lambda$.

$$
\begin{aligned}
\rho(\mu) = \rho(\lambda) &\implies \log \rho(\mu) = \log \rho(\lambda) \\
&\implies \beta H(\mu) + \log Z(\mu)\mathbb{1}^{\otimes n} = \beta H(\lambda) + \log Z(\lambda)\mathbb{1}^{\otimes n} \\
&\implies \log\left(\frac{Z(\mu)}{Z(\lambda)}\right)\mathbb{1}^{\otimes n} = \beta \sum_{j=1}^{m} (\mu_j - \lambda_j) L_j \\
&\implies \log\left(\frac{Z(\mu)}{Z(\lambda)}\right)\mathrm{tr}(L_k) = \beta(\mu_k - \lambda_k)2^n \quad \forall k \in [m] \qquad \text{(Ortho: } \mathrm{tr}(L_j L_k) = 2^n \delta_{jk}\text{)} \\
&\implies \mu = \lambda \quad \text{for } \beta > 0 \qquad \text{(Traceless: } \mathrm{tr}(L_k) = 0\text{)}
\end{aligned}
$$

Hence, fixing basis $\{L_j\}$, each finite-temperature Gibbs state is generated by a unique LH

- When $\beta \to 0$ ($T \to \infty$), $\mu \neq \lambda$ is allowed $\Rightarrow$ unique recovery forbidden. This conforms to the statement: *nothing is learnable from a maximally mixed state.*
- When $\beta \to \infty$ ($T \to 0$), $\mu = \lambda$ if basis $\{L_j\}$ is traceless and $\log(Z(\mu)/Z(\lambda))$ scales as $O(\beta)$. Since $\frac{\lambda_{max}}{k_B T} \to \infty$, the ground state is extremely sensitive to noise in local marginals, and so are the LH parameters $\Rightarrow$ stable recovery forbidden. This conforms to the statement: *learning from ground states is highly sensitive to noise*

We see that $\mathcal{T}$ is certainly injective at finite-temperature. As explained in the start, this also means $\mathcal{T}$ is bijective. This concludes our proof. □

PS: The existence of a bijective map is solely concerned with *uniqueness*. As discussed, it does not by itself account for the lack of *stability* at low temperatures.

6.3 Explicit form of inverse map & sensitivity to noise

Thus far, we have shown the existence of a bijective non-linear map $\mathcal{T}$ from κ-LHs to the κ-LMs induced by their Gibbs states. We also have an explicit form for the forward map as computing local expectation values. We would now like to present an explicit form for the inverse map $\mathcal{T}^{-1}(\cdot)$ going from κ-LMs back to κ-LH parameters, as well as understand the sensitivity of $\mathcal{T}^{-1}$ to noise in the marginals, which is understood via a bound on the condition number of $\mathcal{T}$. We present a small factor-of-2 improvement over the condition number obtained by AAKS.

Let us recall what we know. We have a synthetic many-qubit device on which we prepare copies of the Gibbs states $\rho(\mu)$ of a chosen local Hamiltonian $H(\mu)$. We then make local measurements on a certain number of copies (sample complexity) of $\rho(\mu)$ to obtain the full set of local marginals denoted $l(\mu) = (l_1(\mu) \quad l_2(\mu) \ldots l_m(\mu))^\intercal$. The "inputs" to the inverse map are the local marginals $l(\mu) \in \mathbb{R}^m$, and the "outputs" of the inverse map are the LH parameters $\mu \in \mathbb{R}^m$. Without loss of generality, we can assume that μ lies in a 2-norm ball $\|x\|_2 \leq 1$ to reduce our search space.

Lemma 9 (Entropy-Free Energy duality)**.** *Consider a primal optimization task as follows:*

$$\begin{aligned} p^* = \max_{\sigma \in \mathrm{Pos}(\mathcal{X}^{\otimes n})} \quad & S(\sigma) = -\operatorname{tr}(\sigma \log \sigma) \\ s.t. \quad & \operatorname{tr}(\sigma L_j) = l_j(\mu) \quad \forall j \in [m] \\ & \operatorname{tr}(\sigma) = 1 \end{aligned}$$

Its Lagrange dual is the following minimization

$$\begin{aligned} d^* = \min_{x \in \mathbb{R}^m} \quad & g(x) = \log Z(x) + \beta x^\top l(\mu) \\ s.t. \quad & \|x\|_2 \leq 1 \end{aligned}$$

For fixed temperature $\beta > 0$, the unique solution to the above problem is: $(\sigma^*, x^*) = (\rho(\mu), \mu)$. Additionally, strong duality holds: $p^* = d^* = S(\rho(\mu))$

Proof. For clarity, view the primal objective as minimizing negative-entropy $\operatorname{tr}(\sigma \log \sigma)$, which is convex. We then construct a Lagrangian and find its pointwise infimum to be $-g(x)$, which should be concave [pointwise infimum of affine functions]. Maximizing $-g(x)$ then is same as minimizing $g(x)$, a convex function. Also note that the primal form only has equality constraints, so dual variables are unconstrained (except the artificial $\|x\|_2 \leq 1$)

One example of a state that qualifies the marginal constraints $\operatorname{tr}(\sigma L_j) = l_j(\mu) \quad \forall j \in [m]$ while maximizing entropy is the Gibbs state $\rho(\mu) = e^{-\beta H(\mu)}/Z(\mu)$. This makes the primal form strictly feasible, hence strong duality holds [Slater condition].

The primal feasible set has many states, however *the only Gibbs state in the primal feasible set is* $\rho(\mu)$. This follows from Lemma 8, where we proved a bijection between Gibbs states and their LMs. Since maximum-entropy states with linear constraints are Gibbs, we expect that $\rho(\mu)$ is in fact the unique primal solution. This intuition will be made rigorous.

Use dual variables $\beta x_j \in \mathbb{R}$ for marginal constraints and $\nu \in \mathbb{R}$ for the trace constraint to construct the following Lagrangian

$$\mathcal{L}(\sigma, x, \nu) = \operatorname{tr}(\sigma \log \sigma) + \beta \left[\operatorname{tr}(\sigma H(x)) - x^\top l(\mu)\right] + \nu \left[\operatorname{tr}(\sigma) - 1\right] \tag{43}$$

The domain of $\mathcal{L}$ is: $\sigma \in \mathrm{Pos}(\mathcal{X}^{\otimes n})$, $x \in \mathbb{R}^m_+$, $\nu \in \mathbb{R}_+$. By the KKT theorem, σ^* is the primal solution iff (σ^*, x^*, ν^*) is a saddle point of $\mathcal{L}(\cdot)$. The necessary (and often sufficient) conditions for the same are:

(i) Stationarity w.r.t. σ: $\nabla_\sigma \mathcal{L}(\sigma^*, x^*, \nu^*) = 0$

(ii) Primal feasibility: $\operatorname{tr}(\sigma^* L_j) = l_j(\mu) \quad \forall j \in [m]$, and $\operatorname{tr}(\sigma^*) = 1$

(iii) Dual feasibility: $x^* \in \mathbb{R}^m$ and $\nu^* \in \mathbb{R}$

(iv) Complementary slackness: Void for all-equality-constraint primal form

Dual feasibility and complementary slackness are trivial in our problem.

The stationarity condition yields the following

$$\nabla_\sigma \mathcal{L}(\sigma^*, x^*, \nu^*) = 0 \implies \log \sigma^* + \beta H(x^*) + (\nu^* + 1)\mathbb{1}_{2^n} = 0 \tag{44}$$

$$\implies \sigma^* = \frac{e^{-\beta H(x^*)}}{e^{(\nu^* + 1)}} = \rho(x^*) \tag{45}$$

We see that σ^* is indeed a Gibbs state. We have seen that $\rho(\mu)$ is the unique Gibbs state in the primal feasible set. Hence

$$(\sigma^*, x^*, \nu^*) = (\rho(\mu),\ \mu,\ \log Z(\mu) - 1) \tag{46}$$

Hence σ^* is primal feasible. The more general stationarity condition away from the saddle point: $\nabla_\sigma \mathcal{L}(\sigma^*, x, \nu) = 0$ yields the solution

$$\sigma^*(x, \nu) = \frac{e^{-\beta H(x)}}{e^{(\nu+1)}} = \rho(x) \tag{47}$$

Hence the dual objective for minimization is given by

$$\begin{aligned} g(x) = -\inf_\sigma \mathcal{L}(\sigma, x, \nu) &= -\mathcal{L}(\sigma^*(x,\nu), x, \nu) \\ &= S(\rho(x)) - \beta\left[\langle H(x)\rangle_x - x^\top l(\mu)\right] \\ &= \cancel{\beta\langle H(x)\rangle_x} + \log Z(x) - \cancel{\beta\langle H(x)\rangle_x} + \beta x^\top l(\mu) \\ &= \log Z(x) + \beta x^\top l(\mu) \end{aligned}$$

To verify strong duality, note that

$$\begin{aligned} d^* = g(\mu) &= \log Z(\mu) + \beta\mu^\top l(\mu) \\ &= \log Z(\mu) + \beta\langle H(\mu)\rangle_\mu && (\langle H(\mu)\rangle_\mu = \mu^\top l(\mu) \text{ using Definition 7}) \\ &= S(\rho(\mu)) = p^* && \text{(using Claim 8.1)} \end{aligned}$$

This concludes our proof □

Corollary 10 (Inverse map)**.** *As a result of Lemma 9, we have*

$$\begin{aligned} \mu = \mathcal{T}^{-1}(l(\mu)) = \arg\min_{x\in\mathbb{R}^m} \quad & [\log Z(x) + \beta x^\top l(\mu)] \\ s.t. \quad & \|x\|_2 \le 1 \end{aligned}$$

Provided we choose a μ inside the Unit 2-norm Ball (U2B): $\|x\|_2 \le 1$ in $\mathbb{R}^m$, the above inverse map gives a clear prescription for learning μ: Measure LMs $l(\mu)$, and minimize $[\log Z(x) + \beta x^\top l(\mu)]$ numerically inside U2B.

So far, we proved the existence of a bijection between LHs and LMs in Lemma 8. The forward map $\mathcal{T}$ was trivially known by Definition 7. The inverse map $\mathcal{T}^{-1}$ is now known due to Lemma 9 and Corollary 10. Before worrying about sample/time complexity, the last piece in this analysis is understanding the sensitivity of $\mathcal{T}^{-1}$ to noise in LMs.

We will see that the sensitivity of $\mathcal{T}^{-1}$ to noise in LMs depends exactly on two factors:

- Temperature: We observed heuristically in Lemma 8 that *lower the temperature, larger the sensitivity to noise in learning.* We will confirm this rigorously.
- Curvature in $\log Z(x)$: The curvature in $\log Z(x)$ is a positive constant α that depends only on temperature and the local basis $\{L_j\}$ chosen in Definition 5. The *lower this curvature, larger the sensitivity to noise in learning.*

Fix inverse temperature β. Then for commuting LHs, we will see that the curvature in $\log Z(x)$ is proportional to a Minimum-Energy-Variance (MEV) over U2B Gibbs states. Also recall from our discussion about empty-approximate-Kernels the quantity $\lambda^* = \max_{x \in \text{U2B}} \|H(x)\|_\infty$. It would be an interesting exercise to study the relationships between α, MEV, and λ^* numerically. At an intuitive level, for commuting LHs, we expect

$$\alpha \propto \text{MEV} \propto \frac{1}{\lambda^*} \tag{48}$$

The sensitivity to noise of $\mathcal{T}^{-1}$ is defined by a bound on its condition number. Central to bounding the condition number of $\mathcal{T}^{-1}$ is a sharp understanding of strongly convex multivariate functions, in our case applied to $\log Z(x)$. For this reason, it is worth taking a small interlude into some general results from convex optimization to guide our thinking

Definition 11 (Convexity characterizations)**.** *A twice-differentiable function $f : \mathbb{R}^m \to \mathbb{R}$ is convex in $\mathcal{D} \subset \mathbb{R}^m$ under the following equivalent criteria -*

(i) $f(y) \geq f(x) + \nabla f(x)^\intercal (y - x) \quad \forall x, y \in \mathcal{D}$ (linear lower bound)

(ii) $(\nabla f(x) - \nabla f(y))^\intercal (x - y) \geq 0 \quad \forall x, y \in \mathcal{D}$ (gradient monotone)

(iii) $\nabla^2 f(x) \succeq 0 \quad \forall x \in \mathcal{D}$ (positive semidefinite Hessian)

Definition 12 (Strong convexity characterizations)**.** *A twice-differentiable function $f : \mathbb{R}^m \to \mathbb{R}$ is said to be α-strongly convex ($\alpha > 0$) in $\mathcal{D} \subset \mathbb{R}^m$ under the following equivalence -*

(i) $g(x) = f(x) - \frac{\alpha}{2}\|x\|_2^2$ is convex in $\mathcal{D}$

(ii) $f(y) \geq f(x) + \nabla f(x)^\intercal (y - x) + \frac{\alpha}{2}\|x - y\|_2^2 \quad \forall x, y \in \mathcal{D}$ (quadratic lower bound)

(iii) $(\nabla f(x) - \nabla f(y))^\intercal (x - y) \geq \alpha \|x - y\|_2^2 \quad \forall x, y \in \mathcal{D}$ (strong gradient monotone)

(iv) $\nabla^2 f(x) \succeq \alpha \mathbb{1}_m \quad \forall x \in \mathcal{D}$ (positive definite Hessian)

To a reader familiar with convex optimization, these are fairly standard results [15, 16]. A somewhat non-trivial characterization is the gradient monotone. It follows by simply swapping x and y in the linear lower bound and adding the two inequalities. A rough geometric interpretation is - if the input to f is changed along some direction, the gradient of f should also change in the same/similar direction. The strong gradient monotone in Definition 12 is understood similarly.

Lemma 13 (Variant of strong gradient monotone)**.** *Let $f : \mathbb{R}^m \to \mathbb{R}$ be an α-strongly convex function ($\alpha > 0$). Then the following is true -*

$$\|x - y\|_2 \leq \frac{1}{\alpha} \|\nabla f(x) - \nabla f(y)\|_2 \quad \forall x, y \in \mathcal{D} \tag{49}$$

Proof. Write the strong gradient monotone condition in reverse

$$0 \leq \alpha \|x - y\|_2^2 \leq (\nabla f(x) - \nabla f(y))^\intercal (x - y)$$

Notice that the RHS is a non-negative inner product, can be upper bounded by Cauchy-Schwarz

$$(\nabla f(x) - \nabla f(y))^\intercal (x - y) \leq \|\nabla f(x) - \nabla f(y)\|_2 \|x - y\|_2$$

Hence the result follows

$$\alpha\|x-y\|_2^2 \le \|\nabla f(x)-\nabla f(y)\|_2\|x-y\|_2$$
$$\implies \quad \|x-y\|_2 \le \frac{1}{\alpha}\|\nabla f(x)-\nabla f(y)\|_2 \quad \forall x,y\in\mathcal{D} \quad \square$$

We are now prepared to present the sensitivity to noise bound for QHLP. Recall that the inverse map $\mathcal{T}^{-1}$ takes LMs $l(\mu)\in\mathbb{R}^m$ back to LH parameters $\mu\in\mathbb{R}^m$

Lemma 14 (Bound on condition number of $\mathcal{T}^{-1}$)**.** *Let $\mathcal{D}\subset\mathbb{R}^m$ and $\nabla^2\log Z(x)\ge\alpha\mathbb{1}_m \quad \forall x\in\mathcal{D}$. For any pair of LH parameters $\mu,\lambda\in\mathcal{D}$, let $l(\mu),l(\lambda)\in\mathbb{R}^m$ be the corresponding local marginals. Then we have -*

$$\|\mu-\lambda\|_2 \le \frac{\beta}{\alpha}\|l(\mu)-l(\lambda)\|_2 \tag{50}$$

Proof. The proof follows by applying Lemma 13 to $f(x)=\log Z(x)$ evaluated at μ,λ.

We first note that since μ,λ minimize inverse problem objectives as described by Corollary 10, they obey first derivative conditions -

$$\nabla\log Z(\mu)=-\beta l(\mu) \qquad \nabla\log Z(\lambda)=-\beta l(\lambda) \tag{51}$$

We then have

$$\|\mu-\lambda\|_2 \le \frac{1}{\alpha}\|\log Z(\mu)-\log Z(\lambda)\|_2 = \frac{\beta}{\alpha}\|l(\mu)-l(\lambda)\|_2 \quad \square$$

Denote the absolute condition number of $\mathcal{T}^{-1}$ as $\mathrm{CoNum}(\mathcal{T}^{-1})$. For each $l(\mu)\in\mathbb{R}^m$, consider a small variation $\delta l(\mu)$ about $l(\mu)$, which results in a variation $\delta\mu$ about μ. It follows that

$$\mathrm{CoNum}(\mathcal{T}^{-1}) = \lim_{\Delta\to 0}\sup_{\|\delta l(\mu)\|_2\le\Delta}\frac{\|\delta\mu\|_2}{\|\delta l(\mu)\|_2} \le \frac{\beta}{\alpha} \tag{52}$$

AAKS [2] obtained a sensitivity bound of $\frac{2\beta}{\alpha}$. In the above simpler proof, we see that the sensitivity bound can be improved by a factor-of-2 to $\frac{\beta}{\alpha}$

The above improvement - while technically interesting, only leads to a constant factor-of-4 improvement in sample complexity, which is asymptotically inconsequential. However, we can now rigorously understand the conditioning of the learning problem as a function of temperature. Specifically, when $\beta\to 0$, the bound guarantees that the learning problem is well-conditioned (allows stable recovery), while when $\beta\to\infty$, the learning problem is ill-conditioned (forbids stable recovery).

7 Appendix B: Shadow Tomography

In this appendix, we collect notes on Shadow Tomography [3] : a prominent modern tool for quantum state tomography, and Unitary Designs [17] : a crucial element in the success of Shadow Tomography.

Consider N qudits, local dimension $q\ge 2$, Hilbert space dimension $d=q^N$. They are arranged as lattice in space of spatial dimension $D\ge 1$.

Definition 15 (standard basis)**.** *In $\mathbb{C}^q$, the standard basis is denoted $\{|0\rangle, \ldots, |q-1\rangle\}$ representing $\{(10\ldots0)^T, (010\ldots0)^T, \ldots, (0\ldots01)^T\}$*

Let $\omega = \exp(\frac{i2\pi}{q})$. We have $\omega^q = 1,\ \omega \neq 1 \Leftrightarrow 1 + \omega + \omega^2 + \cdots + \omega^{q-1} = 0$

Definition 16 (shift, clock, hadamard)**.** *Shift, Clock, and Hadamard are unitaries $X, Z, W \in \mathcal{U}(q)$ defined by*

Shift: $X = \begin{bmatrix} |1\rangle & |2\rangle \ldots |q-1\rangle & |0\rangle \end{bmatrix} = \sum_{j=0}^{q-1} |(j+1) \mod q\rangle\langle j|$

Clock: $Z = \mathrm{diag}\{1,\ \omega,\ \omega^2, \ldots \omega^{q-1}\} = \sum_{j=0}^{q-1} \omega^j |j\rangle\langle j|$

Hadamard: $W = \frac{1}{\sqrt{q}} \begin{bmatrix} 1 & 1 & \ldots & 1 \\ 1 & \omega^{q-1} & \ldots & \omega^{(q-1)^2} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \omega & \ldots & \omega^{q-1} \end{bmatrix} = \frac{1}{\sqrt{q}} \sum_{j,k=0}^{q-1} \frac{1}{\omega^{jk}} |j\rangle\langle k|$

Properties

i) Unitary but not Hermitian in general: $X \neq X^\dagger, Z \neq Z^\dagger, W \neq W^\dagger$ for $q > 2$

ii) $X^q = Z^q = \mathbb{1}_q$, $\mathrm{tr}(X) = \mathrm{tr}(Z) = 0$

iii) $ZX = \omega XZ$. More generally $Z^n X^m = \omega^{mn} X^m Z^n$

iv) $Z = WXW^\dagger$

v) Complete set of q^2 Paulis: $P(m,n) = X^m Z^n = \sum_{j=0}^{q-1} \omega^{nj} |(j+m) \mod q\rangle\langle j|$ where $m, n \in \mathbb{Z}_q$.

vi) $(ZX)^q = (XZ)^q = (-1)^{q+1}\mathbb{1}_q$. If odd q it is $\mathbb{1}_q$, if even q it is $-\mathbb{1}_q$. Here, $(-1)^{q+1}$ appears as the product of q-th roots of unity.

Intuitively $P(m,n)$ can be thought of as "shift with phase" on basis states.

Paulis form a group $\mathcal{P}_1(q)$ for a single qudit. It is a discrete group of size q^2 (quotient-ing out phase). For $q = 2$, it is the familiar single-qubit Pauli group $\{\mathbb{1}_2, X, Z, Y\}$. For $N > 1$ number of q-state qudits ("qu-q-its"), the Pauli group is $\mathcal{P}_N(q) \equiv \mathcal{P}(q)^{\otimes N}$.

Exercise: Show that $\mathcal{P}_N(q)$ and $\mathcal{P}_1(q^N)$ are different in general, even with same size q^{2N}. For example, a ququart [18] Pauli is not a tensor product of two qubit-Paulis generally.

Let $\mathbb{Z}_n = \{0, 1, \ldots n-1\}$. For simplicity, we denote $\mathcal{P}(2) \equiv \mathcal{P}$, and $\mathcal{P}_N(2) \equiv \mathcal{P}_N$. As such, whenever we skip the argument, it is to be understood as a single/many qubit group. The N-qubit Pauli group has size $|\mathcal{P}_N| = 4^N = d^2$.

Suppose local dimension is q. Define $\omega = \exp(i\frac{2\pi}{q})$, and

$$\tilde{\omega} = \begin{cases} \omega & \text{odd q} \\ \sqrt{\omega} & \text{even q} \end{cases} \tag{53}$$

Then the general hierarchy for qu-q-its followed in Ref. [17] is:

i) Pauli group: $\tilde{\mathcal{P}}_N(q) = \langle \tilde{\omega}\mathbb{1}, X, Z\rangle^{\otimes N} = \{\tilde{\omega}^\alpha X(a)Z(b) : a, b \in \mathbb{Z}_q^N\}$.
Size q^{2N+1} for odd q, and $2q^{2N+1}$ for even q. It may be noted that for odd q, $\langle X, Z\rangle$ are sufficient generators since they naturally generate $\{\mathbb{1}, \omega\mathbb{1} \ldots \omega^{q-1}\mathbb{1}\}$

ii) Paulis modulo phase:
$\mathcal{P}_N(q) = \tilde{\mathcal{P}}_N(q)/\langle\tilde{\omega}\mathbb{1}\rangle = \{X(a)Z(b)\} = \{\otimes_{j\in[N]} X^{a_j}Z^{b_j}\} = \{\otimes_{j\in[N]} P(a_j, b_j)\}$. Size q^{2N} regardless of parity

iii) Paulis modulo phase and identity: $\hat{P}_N(q) = \mathcal{P}_N(q) \setminus \{\mathbb{1}\} = \{X(a)Z(b)\} \setminus \{\mathbb{1}\}$. Size $= (q^{2N} - 1)$.

iv) Non-identity Paulis with order[2] dividing q: $\overline{\mathcal{P}}_N(q) = \{\omega^l p : p \in \hat{P}_N(q), l \in \mathbb{Z}_q\}$
Size $= q(q^{2N} - 1)$

We have $\hat{\mathcal{P}}_N \subset \mathcal{P}_N \subset \tilde{\mathcal{P}}_N$. A powerful fact is that $\mathcal{P}_N(q)$ forms a basis for $\mathrm{L}(q^N)$ with complex coefficients, and for $\mathrm{Herm}(q^N)$ with real coefficients.

Theorem 17 (commutation classes)**.** *Each non-identity Pauli P partitions $\mathcal{P}_N(q)$ into q subsets by commutation relation: $\mathcal{P}^{(l)} = \{Q \in \mathcal{P}_N(q) : PQ = \omega^l QP,\ l \in \mathbb{Z}_q\}$. We have $\mathcal{P}_N(q) = \mathcal{P}^{(1)} \cup \mathcal{P}^{(2)} \cup \cdots \cup \mathcal{P}^{(q)}$*

A proof follows straightforwardly from Definition 2. This has an important corollary for qubits -

Corollary 18. *Every non-Identity Pauli commutes with exactly half the elements in a qubit Pauli group, and necessarily anti-commutes with the other half.*

The notion of commuting and anti-commuting Paulis has a nice equivalence in terms of operator size, specifically the non-commuting overlap.

Theorem 19 (non-commuting overlap)**.** *Let $P, Q \in \mathcal{P}_N(2)$, $n \geq 1$. And $P \neq \mathbb{1}^{\otimes N}, Q \neq \mathbb{1}^{\otimes N}$. Then it is the case that -*

- $[P, Q] = 0 \Leftrightarrow |\mathrm{supp}(PQ|_{P\cap Q})|$ *is even*
- $\{P, Q\} = 0 \Leftrightarrow |\mathrm{supp}(PQ|_{P\cap Q})|$ *is odd*

$P \cap Q$ indicates the overlapping support of P and Q. $PQ|_{P\cap Q}$ conveys the idea of focusing the product PQ on the overlapping support $P \cap Q$. Finally, $\mathrm{supp}(PQ|_{P\cap Q})$ measures the size of non-commuting overlap.

In other words, two Paulis commute if and only if their non-commuting overlap has even parity. Likewise, two Paulis anti-commute if and only if their non-commuting overlap has odd parity.

However, theorem 19 does not generalize straightforwardly for local dimension $q > 2$ in terms of remainder classes. We are not aware of a simple generalization.

The Clifford group is defined as the unitary normalizer of the extended Pauli group. It is made finite by quotient-ing out continuous phase $\langle e^{i\theta}\mathbb{1}\rangle$. In particular,

$$\tilde{Cl}_N = \{c \in \mathcal{U}(2^N) : cPc^\dagger \in \tilde{\mathcal{P}}_N \quad \forall P \in \tilde{\mathcal{P}}_N\} \tag{54}$$

$$Cl_N = \tilde{Cl}_N/\langle e^{i\theta}\mathbb{1}\rangle \tag{55}$$

We have $\mathcal{P}_N \subset Cl_N$, and $\mathcal{P}_N \subset \tilde{\mathcal{P}}_N$. However, $\tilde{\mathcal{P}}_N \notin Cl_N$ since Cl_N is modulo global phase. the smallest group that Cl_N

[2] smallest integer k for $P^k = \mathbb{1}$

In fact, $\mathcal{P}_N = \tilde{P}_N \cap Cl_N$. The generators of Cl_N are known to be the familiar gates $\langle H, S, CX\rangle$. It is the case that

$$(X,Y,Z) \xrightarrow{H} (Z,-Y,X) \qquad (X,Y,Z) \xrightarrow{S} (Y,-X,Z) \quad (X_1,X_2,Z_1,Z_2) \xrightarrow{CX} (X_1X_2,X_2,Z_1,Z_1Z_2) \tag{56}$$

For a single qubit ($N = 1$), we have $|\mathcal{P}_1| = 4$ while $|Cl_1| = 6 \cdot 4 = 24$ [19]. Ozols also gives a simple argument for the recursion

$$|Cl_N| = 2(4^N - 1)4^N |Cl_{N-1}| \tag{57}$$

Solving the recursion, it follows that

$$|Cl_N| = 2^{N^2+2N} \prod_{j=1}^{N} (2^{2j} - 1) \tag{58}$$

The symplectic structure dictates that $Cl_N/\mathcal{P}_N \simeq Sp_{2N}(\mathbb{F}_2)$, which are $2N \times 2N$ symplectic matrices over $\mathbb{F}_2$ [20]. We have $|Sp_{2N}(\mathbb{F}_2)| = 2^{N^2} \prod_{j=1}^{N} (2^{2j} - 1)$.

Given an integer $1 \le i \le |Cl_N|$, consider the task of outputting a unique representation of an element in Cl_N on a classical computer. Paulis are uniquely specified by mapping i to pairs of N-bit-strings, an O(1) task[3]. However, mapping the remaining bits of i to unique $2N \times 2N$ Symplectic matrices is non-trivial, requiring an $O(N^3)$ time recursive algorithm [21].

Even if we have efficient binary representations of elements in $\mathcal{P}_N$ and Sp_{2N}, there is the open-ended challenge of *optimal Clifford circuit synthesis*. Only recently the $N = 6$ qubits case was solved [22]. Prior understanding was limited to $N \le 4$ qubits [23, 24, 25]. It is worth stressing that the phrase "optimal synthesis" by itself is most strongly captured by minimum CNOT/2-qubit gate count. A separate body of work studies "depth-optimal synthesis" based on circuit-depth, better understood for general N [26].

We now define unitary twirls as a starting point to understand unitary designs. Recall $d = q^N$. We'll use W_π to denote permutation operators for $\pi \in S_k$.

Definition 20 (k-fold Haar twirl)**.** *For an N qu-q-it system, the k-fold[4] Haar twirl is a channel $\Phi^{(k)}_{Haar} : \mathrm{L}(d^k) \to \mathrm{L}(d^k)$ such that*

$$\Phi^{(k)}_{Haar}(O) = \int_{\mathcal{U}(d)} dU \; U^{\otimes k} O U^{\dagger \otimes k} \quad \textit{for} \quad \forall O \in \mathrm{L}(d^k) \tag{59}$$

where the integral is over the Haar measure on $\mathcal{U}(d)$
Properties

1. *Linearity:* $\Phi^{(k)}_{Haar}(\lambda O) = \lambda \Phi^{(k)}_{Haar}(O) \qquad \forall \lambda \in \mathbb{C}$
2. *Unitary invariance:* $\Phi^{(k)}_{Haar}(O) = V^{\dagger \otimes k} \Phi^{(k)}_{Haar}(O) V^{\otimes k} = \Phi^{(k)}_{Haar}(V^{\dagger \otimes k} O V^{\otimes k}) \qquad \forall V \in U(d^k)$
3. *Schur-Weyl duality:* $\Phi^{(k)}_{Haar}(O) = \sum_{\pi \in S_k} W_\pi u_\pi(O)$ *for some scalars* $\{u_\pi(O) \in \mathbb{C} : \pi \in S_k\}$

[3] simply focus on any $2N$ bits in the binary representation of $1 \le i \le |Cl_N|$, an $O(N^2)$ bit integer
[4] integer $k \ge 1$

Further constraining Property 3 above using linearity and unitary invariance gives the following canonical form [27] -

$$\Phi^{(k)}_{Haar}(O) = \sum_{\pi,\sigma\in S_k} W_\pi \ (Q^{-1})_{\pi\sigma} \operatorname{tr}(W_\sigma O) \qquad \text{where} \qquad Q_{\pi\sigma} = d^{\#\text{cycles}(\pi\sigma)} \tag{60}$$

The inverse of Q-matrix is known as Weingarten matrix, and is guaranteed to exist if $k \leq d$ (?).

One simple example of the Q-matrix is

$$Q = \begin{pmatrix} d^2 & d \\ d & d^2 \end{pmatrix} \quad \text{for} \quad k = 2 \tag{61}$$

Inverting one ultimately gets

$$\Phi^{(2)}_{Haar}(O) = \frac{1}{d^2-1}\left[\mathbb{1}_{d^2}\left(\operatorname{tr}(O) - \frac{\operatorname{tr}(WO)}{d}\right) + W\left(\operatorname{tr}(WO) - \frac{\operatorname{tr}(O)}{d}\right)\right] \tag{62}$$

Above W = SWAP. One other case that is simple even for the general k-fold twirl is when the image of O lies in the symmetric subspace of its k tensor factors.

$$\text{Im}(O) \subseteq \text{sym}(d^k) \quad \Rightarrow \quad \Phi^{(k)}_{Haar}(O) = \frac{\operatorname{tr}(O)}{\binom{d+k-1}{k}} \Pi_{sym(d^k)} \tag{63}$$

e.g. let $d = k = 2$, and $O = 2|00\rangle\langle 00| + 3|11\rangle\langle 11|$, then $\Phi^{(2)}_{Haar}(O) = \frac{5}{6}(\mathbb{1}_4 + W)$. What if $O = |01\rangle\langle 01|$? $O = (|01\rangle + |10\rangle)(\langle 01| + \langle 10|)$?

Besides Haar measure, the notion of "twirling" extends to arbitrary (discrete or continuous) ensembles of unitaries. Of particular interest is the k-fold uniform Clifford twirl defined as

$$\Phi^{(k)}_{Cliff}(O) = \frac{1}{|Cl_N|} \sum_{U\in Cl_N} U^{\otimes k} O U^{\dagger\otimes k} \tag{64}$$

Definition 21 (unitary k-design)**.** *A finite ensemble $\xi = (p_j, U_j)$, where $\{p_j\}$ is a probability distribution over $\{U_j\}$, forms a unitary k-design if the k-fold Ensemble and Haar twirls match -*

$$\Phi^{(k)}_{\xi}(O) = \Phi^{(k)}_{Haar}(O) \quad \forall O \in \text{L}(d^k) \tag{65}$$

A k-design continues to be k'-design for $1 \leq k' \leq k$, as seen by choosing an operator that is identity in $(k - k')$ sites.

It is not hard to see that uniform $\mathcal{P}_N(2)$ forms a unitary 1-design. It was proven in 2015 that uniform $Cl_N(2)$ forms an exact unitary 3-design, and fails to be a 4-design [17, 28]. We review a proof here.

First note that both $\Phi^{(3)}_{Cliff}(.)$ and $\Phi^{(3)}_{Haar}(.)$ are linear maps, hence completely specified by their action on Pauli operators.

The Paulis in this case take the form $P\otimes Q\otimes R$ for arbitrary $P, Q, R \in \mathcal{P}_N(2)$. We would like to know how $\Phi^{(3)}_{Haar}(P\otimes Q\otimes R)$ and $\Phi^{(3)}_{Cliff}(P\otimes Q\otimes R)$ compare for every possible choice of $P, Q, R \in \mathcal{P}_N$. However, it suffices to only compute the Clifford twirl explicitly, while checking a special condition.

Specifically, if we can show that the Clifford twirl lies in the span of three-fold permutation operators span$\{W_\pi\}$ for $\pi \in \mathcal{S}_3$, we are effectively done. Why? It is easy to see that permutation operators are fixed by the Haar twirl: $\Phi^{(k)}_{Haar}(W_\pi) = W_\pi$ for $\pi \in S_k$. So any operator that lives in span$\{W_\pi\}$ is fixed by the Haar twirl. Hence, if $\Phi^{(3)}_{Cliff}(O) \in \text{span}\{W_\pi\} \quad \forall O \in \mathrm{L}(d^3)$, then $\Phi^{(3)}_{Haar}(\Phi^{(3)}_{Cliff}(O)) = \Phi^{(3)}_{Cliff}(O)$. But $\Phi^{(3)}_{Haar}(\Phi^{(3)}_{Cliff}(O)) = \Phi^{(3)}_{Haar}(O)$ simply by unitary invariance of Haar measure. It follows that $\Phi^{(3)}_{Cliff}(O) = \Phi^{(3)}_{Haar}(O)$ whenever $\Phi^{(3)}_{Cliff}(O) \in \text{span}\{W_\pi\}$. If the latter holds $\forall O \in \mathrm{L}(d^3)$, we have a unitary 3-design.[5]. More generally,

$$\varepsilon = (p_j, U_j) \quad \text{is an exact unitary k-design} \quad \underline{\text{iff}} \quad \Phi^{(k)}_\varepsilon(O) \in \text{span}\{W_\pi\} \quad \forall O \in \mathrm{L}(d^k), \pi \in S_k \tag{66}$$

Now, for any choice of $P, Q, R \in \mathcal{P}_N(2)$, we have two relevant cases -
i) $PQR \neq i^\alpha \mathbb{1}_d$. Choose $S \in \mathcal{P}_N$ so that $\{S, PQR\} = 0$. Then it is easy to see that

- $[S,P] = 0 \Rightarrow \{S,QR\} = 0 \Rightarrow [S,Q] = 0$ & $\{S,R\} = 0 \quad$ (or) $\quad [S,R] = 0$ & $\{S,Q\} = 0$. Exactly one of P, Q, R anti-commutes with S
- $\{S,P\} = 0 \Rightarrow [S,QR] = 0 \Rightarrow [S,Q] = [S,R] = 0 \quad$ (or) $\quad \{S,Q\} = \{S,R\} = 0$.
  Either one or all three of P, Q, R anti-commute with S

Basically, any $S \in \mathcal{P}_N(2)$ that anti-commutes with PQR also anti-commutes with an odd number of Paulis out of P, Q, R. This implies $\Phi^{(3)}_{Cliff}(P \otimes Q \otimes R) = \Phi^{(3)}_{Cliff}(SPS^\dagger \otimes SQS^\dagger \otimes SRS^\dagger) = -\Phi^{(3)}_{Cliff}(P \otimes Q \otimes R)$. The first equality is by $Cl_N(2) \cdot S \equiv Cl_N(2)$ [upto global phase that cancels in twirl], second due to odd number of anti-commutes. It follows that $\Phi^{(3)}_{Cliff}(P \otimes Q \otimes R) = 0$

ii) $PQR = i^\alpha \mathbb{1}_d$. There are three sub cases -

- $P = Q = R = \mathbb{1}_d$. In this case, $\Phi^{(3)}_{Cliff}(\mathbb{1}_{d^3}) = \mathbb{1}_{d^3}$
- $P = Q \neq \mathbb{1}_d$ and $R = \mathbb{1}_d$. In this case (and two similar combinations), $\Phi^{(2)}_{Cliff}(P \otimes P \otimes \mathbb{1}_d) = \frac{1}{d^2-1}(d \cdot W_{213} - \mathbb{1}_{d^3})$
- $P \neq Q \neq R \neq \mathbb{1}_d$. In this case, we have Paulis of the form $i^\alpha \cdot P \otimes PR \otimes R$. Additionally, it can be that $\{P,R\} = 0$ or $[P,R] = 0$. Clifford evolution would preserve these commutation relations. Insight into these twirls comes from how the order-3 permutations W_{231} and W_{312} decompose in the Pauli basis.

$$W_{231} = \sum_{P,Q,R \in \mathcal{P}_N} \frac{\delta_{Q,RP}}{d^2} P \otimes Q \otimes R \quad \equiv \Omega + \Theta - \mu \tag{67}$$

$$W_{312} = \sum_{P,Q,R \in \mathcal{P}_N} \frac{\delta_{Q,PR}}{d^2} P \otimes Q \otimes R \quad \equiv \Omega + \Theta + \mu \tag{68}$$

where

$$\Omega = \frac{1}{d^2}\left(\mathbb{1} \otimes \mathbb{1} \otimes \mathbb{1} + \sum_{P \neq \mathbb{1}} (P \otimes P \otimes \mathbb{1} + \mathbb{1} \otimes P \otimes P + P \otimes \mathbb{1} \otimes P)\right) = \frac{W_{123}}{d^2} + \frac{1}{d}(W_{213} + W_{132} + W_{321}) \tag{69}$$

[5] the argument by itself gives a sufficient condition, but necessity also follows directly from Schur-Weyl duality. So we can also invalidate certain ensembles ε as not designs if $\Phi^{(k)}_\varepsilon(O) \notin \text{span}\{W_\pi\}$ for some $O \in \mathrm{L}(d^k)$

and for $P \neq R \neq \mathbb{1}$ we have

$$\Theta = \frac{1}{d^2} \sum_{[P,R]=0} P \otimes PR \otimes R \qquad \mu = \frac{1}{d^2} \sum_{\{P,R\}=0} P \otimes PR \otimes R \tag{70}$$

It follows from that (67) and (68) that

$$\Theta = \frac{1}{2}(W_{312} + W_{231}) - \Omega \qquad \mu = \frac{1}{2}(W_{312} - W_{231}) \tag{71}$$

where Ω has already been expressed in terms of permutations in equation (69). Now observe that for Paulis $P \neq R \neq \mathbb{1}$, the triplet (P, PR, R) under Clifford evolution maps to either (A, AB, B) or $(-A, -AB, B)$ or $(A, -AB, -B)$ or $(-A, AB, -B)$ for some Paulis $A \neq B \neq \mathbb{1}$. But nicely enough, the signs cancel in a tensor product, so it effectively reduces to one case and not four.

In $\bar{P}_N = \{\pm P : P \in \mathcal{P}_N, P \neq \mathbb{1}\}$, the number of pairs of anti-commuting Paulis is $(2d^2 - 2)d^2$, while the number of pairs of commuting Paulis (non-identity and non-equal) is $(2d^2-2)(d^2-4)$

It follows that for Paulis $P \neq R \neq \mathbb{1}$

$$\{P, R\} = 0 : \qquad \Phi^{(3)}_{Cliff}(i^\alpha \cdot P \otimes PR \otimes R) = \frac{i^\alpha}{2d^2(d^2-1)} \sum_{\{A,B\}=0} A \otimes AB \otimes B \tag{72}$$

$$= \frac{i^\alpha}{4(d^2-1)}(W_{312} - W_{231}) \tag{73}$$

$$[P, R] = 0 : \qquad \Phi^{(3)}_{Cliff}(i^\alpha \cdot P \otimes PR \otimes R) = \frac{i^\alpha}{2(d^2-1)(d^2-4)} \sum_{[A,B]=0} A \otimes AB \otimes B \tag{74}$$

$$= \frac{i^\alpha d^2}{4(d^2-1)(d^2-4)}(W_{312} + W_{231} - 2\Omega) \tag{75}$$

$$= \frac{i^\alpha d^2}{4(d^2-1)(d^2-4)} \left(W_{312} + W_{231} - \frac{2}{d}(W_{213} + W_{132} + W_{321}) - \frac{2}{d^2} W_{123} \right) \tag{76}$$

It follows that $\Phi^{(3)}_{Cliff}(P \otimes Q \otimes R) \in \text{span}\{W_\pi\}$ for all $P, Q, R \in \mathcal{P}_N(2)$, $\pi \in S_3$. Hence the Clifford group forms a unitary 3-design.

We are now well equipped to understand Shadow Tomography.

Consider a setting where $\rho \in \text{Pos}(2^N)$ is an unknown N-qubit state known to Alice but unknown to Bob. Alice prepares copies of ρ in an experimental quantum device presented to Bob. She then asks Bob to estimate the expectation values of a set of observables $\{O_1, O_2, \ldots O_M\}$ drawn from $\mathcal{P}_N(2)$ on the unknown state ρ to ϵ-precision each.

Bob is only allowed to measure each copy of ρ once in a basis drawn from an N-qubit Unitary Ensemble ξ. Bob is asked to provide an analytical upper bound on the number of copies needed for his protocol (Sample complexity N_s)

Substitute Huang, Kueng and Preskill (HKP) in place of Bob. HKP find that

$$N_s = O\left(\frac{\log M}{\epsilon^2} \max_j \|O_j\|^2_\xi \right) \tag{77}$$

where $\|\cdot\|_\xi$ is called the Shadow Norm of ensemble ξ. The Shadow Norm is defined in terms of the ensemble 3-twirl $\Phi_\xi^{(3)}(\cdot)$ as follows -

$$\|O\|_\xi^2 = \max_{\rho \in \text{Pos}(2^N)} \text{tr}\left[(\rho \otimes \mathcal{M}_\xi^{-1}(O) \otimes \mathcal{M}_\xi^{-1}(O)) \cdot \Phi_\xi^{(3)}(\sum_b |bbb\rangle\langle bbb|)\right] \tag{78}$$

Here $\mathcal{M}_\xi(\cdot)$ is the Measurement Channel for ensemble ξ, and is defined in terms of the ensemble 2-twirl $\Phi_\xi^{(2)}(\cdot)$ as follows -

$$\mathcal{M}_\xi(O) = \text{tr}_1\left[(O \otimes \mathbb{1})\Phi_\xi^{(2)}(\sum_b |bb\rangle\langle bb|)\right] \tag{79}$$

The interested reader should go through reference [3] for greater detail on the protocol. What is important to us is that whenever the ensemble ξ is Haar-random or a composition of Haar-random gates, the k-twirls $\Phi_\xi^{(k)}(\cdot)$ are in-principle tractable analytically.

HKP found that if $\xi = Cl(2)^{\otimes N}$, then $2^N \leq \|O_A\|_\xi^2 \leq 3 \cdot 2^N$, while if $\xi = Cl(2^N)$ then $\|O_A\|_\xi^2 = 3^{|O_A|}$, where $|O_A|$ is the locality of O_A for all $O_A \in \mathcal{P}_N(2)$

In other words, a "global Clifford" ensemble is efficient for local observable estimation, whereas both "local Clifford" and "global Clifford" ensembles are bad for global observable estimation

Since we learnt earlier that the Clifford group on qubits is an exact unitary 3-design, both $\mathcal{M}_\xi(\cdot)$ and $\|\cdot\|_\xi$ for Clifford based ensembles are exactly equivalent to that of Haar-gate based ensembles.

These notes on Shadow Tomography are not complete by any means. There are several interesting directions pertaining to Shadow Tomography we would like to explore in the future.

References

[1] David Poulin and Pawel Wocjan. "Sampling from the thermal quantum Gibbs state and evaluating partition functions with a quantum computer". In: *Physical review letters* 103.22 (2009), p. 220502.

[2] Anurag Anshu et al. "Sample-efficient learning of quantum many-body systems". In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2020, pp. 685–691.

[3] Hsin-Yuan Huang, Richard Kueng, and John Preskill. "Predicting many properties of a quantum system from very few measurements". In: *Nature Physics* 16.10 (2020), pp. 1050–1057.

[4] Alex Gezerlis. *Numerical methods in physics with Python*. Cambridge University Press, 2023.

[5] Jeongwan Haah, Robin Kothari, and Ewin Tang. "Optimal learning of quantum Hamiltonians from high-temperature Gibbs states". In: *arXiv preprint arXiv:2108.04842* (2021).

[6] Ainesh Bakshi et al. "Learning quantum Hamiltonians at any temperature in polynomial time". In: *arXiv preprint arXiv:2310.02243* (2023).

[7] Anurag Anshu. "Some Recent Progress in Learning Theory: The Quantum Side". In: *Harvard Data Science Review* 4.1 (Jan. 27, 2022). https://hdsr.mitpress.mit.edu/pub/3x2sd8nq. DOI: 10.1162/99608f92.9c47e402. URL: https://hdsr.mitpress.mit.edu/pub/3x2sd8nq.

[8] Fernando GSL Brandão et al. "Quantum SDP solvers: Large speed-ups, optimality, and applications to quantum learning". In: *arXiv preprint arXiv:1710.02581* (2017).

[9] Anirban Narayan Chowdhury and Rolando D Somma. "Quantum algorithms for Gibbs sampling and hitting-time estimation". In: *arXiv preprint arXiv:1603.02940* (2016).
[10] Sathyawageeswar Subramanian, Stephen Brierley, and Richard Jozsa. "Implementing smooth functions of a Hermitian matrix on a quantum computer". In: *Journal of Physics Communications* 3.6 (2019), p. 065002.
[11] Andrew M Childs and Nathan Wiebe. "Hamiltonian simulation using linear combinations of unitary operations". In: *arXiv preprint arXiv:1202.5822* (2012).
[12] Sumeru Hazra et al. "Ring-Resonator-Based Coupling Architecture for Enhanced Connectivity in a Superconducting Multiqubit Network". In: *Physical Review Applied* 16.2 (2021), p. 024018.
[13] Matthew B Hastings. "Locality in quantum systems". In: *Quantum Theory from Small to Large Scales* 95 (2010), pp. 171–212.
[14] John Watrous. *The theory of quantum information*. Cambridge university press, 2018.
[15] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
[16] Xingyu Zhou. "On the fenchel duality between strong convexity and lipschitz continuous gradient". In: *arXiv preprint arXiv:1803.06573* (2018).
[17] Zak Webb. "The Clifford group forms a unitary 3-design". In: *arXiv preprint arXiv:1510.02769* (2015).
[18] Eleonora Nagali et al. "Experimental generation and characterization of single-photon hybrid ququarts based on polarization and orbital angular momentum encoding". In: *Physical Review A* 81.5 (2010), p. 052317.
[19] Maris Ozols. "Clifford group". In: *Essays at University of Waterloo, Spring* (2008).
[20] Scott Aaronson and Daniel Gottesman. "Improved simulation of stabilizer circuits". In: *Physical Review A* 70.5 (2004), p. 052328.
[21] Robert Koenig and John A Smolin. "How to efficiently select an arbitrary Clifford group element". In: *Journal of Mathematical Physics* 55.12 (2014), p. 122202.
[22] Sergey Bravyi, Joseph A Latone, and Dmitri Maslov. "6-qubit optimal Clifford circuits". In: *npj Quantum Information* 8.1 (2022), p. 79.
[23] Farrokh Vatan and Colin Williams. "Optimal quantum circuits for general two-qubit gates". In: *Physical Review A* 69.3 (2004), p. 032315.
[24] Oleg Golubitsky and Dmitri Maslov. "A study of optimal 4-bit reversible Toffoli circuits and their synthesis". In: *IEEE Transactions on Computers* 61.9 (2011), pp. 1341–1353.
[25] Vadym Kliuchnikov and Dmitri Maslov. "Optimization of Clifford circuits". In: *Physical Review A* 88.5 (2013), p. 052307.
[26] Sergey Bravyi and Dmitri Maslov. "Hadamard-free circuits expose the structure of the Clifford group". In: *IEEE Transactions on Information Theory* 67.7 (2021), pp. 4546–4563.
[27] Daniel A Roberts and Beni Yoshida. "Chaos and complexity by design". In: *Journal of High Energy Physics* 2017.4 (2017), pp. 1–64.
[28] Huangjun Zhu. "Multiqubit Clifford groups are unitary 3-designs". In: *Physical Review A* 96.6 (2017), p. 062336.