Ph 220: Quantum Learning Theory
Lecture 7: Learning Many-Body Hamiltonians

Hsin-Yuan Huang (Robert)

California Institute of Technology

1 Introduction

1.1 Recap: Quantum Sensing

In our first lecture, we explored the fundamental limits that quantum mechanics imposes on parameter estimation. Consider the simplest scenario: a single qubit evolving under the Hamiltonian $H = \omega Z$, where ω is an unknown parameter we wish to estimate to precision ε. The evolution time T required to achieve this precision is governed by two distinct scaling laws:

- **Standard Quantum Limit (SQL):** $T \sim \mathcal{O}(\varepsilon^{-2})$.
- **Heisenberg Limit (HL):** $T \sim \mathcal{O}(\varepsilon^{-1})$. This is the ultimate limit imposed by quantum mechanics, achievable through quantum-enhanced strategies such as long-time coherent evolution or entanglement-based protocols.

The SQL vs HL separation is one of the first known quantum advantage in sensing and learning.

1.2 The Many-Body Challenge

Today we consider a significantly more demanding problem. We consider an n-qubit system governed by an unknown Hamiltonian H describing a quantum system with many constituents. Any such Hamiltonian can be expanded in the Pauli basis:

$$H = \sum_{P \in \{I,X,Y,Z\}^{\otimes n}} \alpha_P P, \tag{1}$$

where $\alpha_P \in \mathbb{R}$ are the unknown coefficients we seek to learn. Unlike the single-qubit case, learning many-body Hamiltonians is much harder. Previous algorithms were fundamentally limited to the SQL scaling of $\mathcal{O}(\varepsilon^{-2})$. The difficulty stems from the interplay between many-body interactions and entanglement: as evolution time t increases, entanglement grows rapidly throughout the system. This entanglement causes different parameters to become inextricably coupled, and the resulting many-body correlations act as an effective decoherence mechanism that destroys the quantum enhancement necessary for achieving the Heisenberg limit.

To circumvent these obstacles, prior approaches restricted themselves to short evolution times, thereby forfeiting any quantum advantage and remaining confined to the SQL.

1.3 A Solution: Reshaping

In this lecture, we introduce the **Reshaping Technique**, an approach that achieves the Heisenberg limit even for interacting many-body Hamiltonians.

The central idea is simple: we strategically interleave the unknown Hamiltonian evolution with carefully chosen unitaries to effectively decouple the system into simple, non-interacting patches. Within each patch, the effective Hamiltonian has known eigenstates, enabling us to apply powerful quantum-enhanced phase estimation protocols. This technique allows us to reduce an intractable many-body problem into a collection of tractable few-body problems that can be solved in parallel, thereby unlocking Heisenberg-limited precision for the many-body setting.

2 The Reshaping Technique

2.1 Conceptual Foundation

The reshaping technique rests on a simple but powerful observation: while we cannot directly control or measure the unknown Hamiltonian H, we can manipulate how the system evolves under H by interleaving this evolution with known unitaries. This manipulation allows us to synthesize an effective Hamiltonian $\widetilde{H}$ with more favorable properties for learning.

2.2 Mathematical Definition

Given an unknown Hamiltonian H, we select a finite set of K unitaries $\{U_k\}_{k=1}^K$ together with associated weights $w_k \geq 0$ satisfying $\sum_{k=1}^K w_k = 1$. We define the **reshaped Hamiltonian** as:

$$\widetilde{H} \triangleq \sum_{k=1}^{K} w_k U_k H U_k^\dagger. \tag{2}$$

This definition encodes a weighted average of the Hamiltonian H in different reference frames determined by the unitaries U_k. The choice of unitaries and weights will depend on the specific learning task, as we illustrate through examples below.

2.3 Physical Implementation via Quantum Simulation

A natural question arises: How can we implement evolution under $\widetilde{H}$ when we have direct access only to the physical evolution $U(t) = e^{-itH}$? The answer lies in leveraging techniques from Hamiltonian simulation.

While we cannot access $\widetilde{H}$ directly, we can approximate its time evolution $e^{-it\widetilde{H}}$ using the available evolution e^{-itH} combined with our unitaries. A canonical approach is:

Trotterization: We can use deterministic product formulas. The first-order Trotter approximation for time t is given by:

$$e^{-it\widetilde{H}} \approx \left(\prod_{k=1}^{K} e^{-i(t/r) w_k U_k H U_k^\dagger} \right)^r. \tag{3}$$

The key identity enabling this implementation is:

$$e^{-i\tau(UHU^\dagger)} = U e^{-i\tau H} U^\dagger, \tag{4}$$

which shows that evolving under the conjugated Hamiltonian $UHU^\dagger$ is equivalent to sandwiching the physical evolution between the circuits U and $U^\dagger$. This implementation requires only the ability to apply the unknown evolution e^{-itH} and the unitaries U_k, without any knowledge of the Hamiltonian parameters themselves.

3 Warm-Up: Learning Single-Qubit Hamiltonians

To build intuition, let us first apply the reshaping technique to learn a single-qubit Hamiltonian. This example illustrates the core principles in their simplest form.

3.1 The Problem

Consider an arbitrary single-qubit Hamiltonian:

$$H = \lambda_x X + \lambda_y Y + \lambda_z Z, \tag{5}$$

where $\lambda_x, \lambda_y, \lambda_z$ are unknown parameters satisfying $|\lambda_x|, |\lambda_y|, |\lambda_z| \leq 1$. Our goal is to estimate each parameter to accuracy ε with high probability.

3.2 The Obstacle

If all three coefficients were zero except one, say $H = \lambda_z Z$, this problem would be straightforward. The eigenstates of Z are just $|0\rangle, |1\rangle$, and we could directly apply what we learned in Lecture Note 1 (also known as Robust Phase Estimation; see `https://arxiv.org/abs/1502.02677`) to learn λ_z with a total sensing time of $T = \mathcal{O}(\varepsilon^{-1})$. However, for a generic single-qubit Hamiltonian where all three coefficients are nonzero, the eigenstates are unknown. Standard quantum sensing protocols require knowledge of the eigenstates, creating an apparent obstacle: we cannot identify the eigenstates without first learning the Hamiltonian parameters, but we cannot efficiently learn the parameters without knowing the eigenstates.

3.3 The Solution: Strategic Reshaping

We resolve this circularity through reshaping. To isolate and estimate λ_z while eliminating the interference from λ_x and λ_y, we choose unitaries $U_1 = I$ and $U_2 = Z$ with equal weights $w_1 = w_2 = 1/2$. The reshaped Hamiltonian becomes:

$$\widetilde{H} = \frac{1}{2}(I \cdot H \cdot I) + \frac{1}{2}(Z \cdot H \cdot Z) \tag{6}$$

$$= \frac{1}{2}(\lambda_x X + \lambda_y Y + \lambda_z Z) + \frac{1}{2}(-\lambda_x X - \lambda_y Y + \lambda_z Z) \tag{7}$$

$$= \lambda_z Z. \tag{8}$$

The second line follows from the anticommutation relations: $ZXZ = -X$ and $ZYZ = -Y$. These relations cause the X and Y terms to cancel perfectly when we average over the two reference frames, leaving only the desired Z term.

Having isolated λ_z, we can now apply Robust Phase Estimation using the known eigenstates $|0\rangle, |1\rangle$ of Z to achieve Heisenberg-limited precision $T \sim \mathcal{O}(\varepsilon^{-1})$.

3.4 Learning All Parameters

By analogous constructions, we can isolate each parameter:

- To learn λ_x: Use $U_1 = I, U_2 = X$ with weights $1/2$ each, yielding $\widetilde{H} = \lambda_x X$.
- To learn λ_y: Use $U_1 = I, U_2 = Y$ with weights $1/2$ each, yielding $\widetilde{H} = \lambda_y Y$.

This procedure allows us to estimate all three parameters of an arbitrary single-qubit Hamiltonian with total evolution time $T = \mathcal{O}(\varepsilon^{-1} \log(1/\delta))$ and using only $\mathcal{O}(\text{poly}\log(\varepsilon^{-1}) \log(1/\delta))$ experiments, where δ is the failure probability.

4 Generalizing to Few-Qubit Hamiltonians

The single-qubit example might seem special, but the reshaping technique generalizes naturally to learning Hamiltonians over a few qubits. This is an important step before tackling the full many-body problem.

4.1 The Two-Qubit Challenge

Consider a generic two-qubit Hamiltonian:

$$H = \sum_{A,B \in \{I,X,Y,Z\}} \lambda_{AB} \, A \otimes B. \tag{9}$$

Suppose we want to estimate the parameter λ_{XZ}, corresponding to the interaction term $X \otimes Z$. Just like in the single-qubit case, the presence of other non-commuting terms (e.g., $Y \otimes Y$ or $Z \otimes X$) scrambles the eigenstates, preventing direct estimation.

4.2 Reshaping for Eigenstate Accessibility

To isolate λ_{XZ}, we apply reshaping using the group of unitaries generated by the target Pauli operators. Specifically, we choose the set of 4 unitaries:

$$\mathcal{U} = \{I \otimes I, \quad X \otimes I, \quad I \otimes Z, \quad X \otimes Z\}. \tag{10}$$

We set the weights equal, $w_k = 1/4$. Let us analyze the effect of this reshaping. The averaging process eliminates any Pauli term that anticommutes with elements of $\mathcal{U}$.

- A term like $Y \otimes I$ is eliminated because $(X \otimes I)(Y \otimes I)(X \otimes I) = -Y \otimes I$.
- A term like $Z \otimes Z$ is eliminated because $(X \otimes I)(Z \otimes Z)(X \otimes I) = -Z \otimes Z$.

The only terms that survive are those that commute with all generators in $\mathcal{U}$. The resulting reshaped Hamiltonian is:

$$\widetilde{H} = \lambda_{XZ} X \otimes Z + \lambda_{XI} X \otimes I + \lambda_{IZ} I \otimes Z + \lambda_{II} I \otimes I. \tag{11}$$

4.3 Why This Solves the Problem

At first glance, $\widetilde{H}$ still appears complex as it contains multiple terms. However, it possesses a structural property: **on each qubit, only one type of non-identity Pauli operator appears.**

- Qubit 1 only sees operators X and I.
- Qubit 2 only sees operators Z and I.

This implies that $\widetilde{H}$ is diagonal in a known product basis. Specifically, the eigenstates are tensor products of the eigenstates of the local Pauli operators:

$$|\psi_{a,b}\rangle = |x_a\rangle \otimes |z_b\rangle \quad \text{where } a, b \in \{0, 1\}, \tag{12}$$

and $X|x_a\rangle = (-1)^a |x_a\rangle$, $Z|z_b\rangle = (-1)^b |z_b\rangle$. Because we explicitly know these eigenstates, we can prepare them and evolve them under the reshaped Hamiltonian $\widetilde{H}$.

Reconstructing the Parameters: The eigenvalues of $\widetilde{H}$ corresponding to these states are linear combinations of the unknown parameters. For example, the energy of the state $|x_0\rangle |z_0\rangle$ (corresponding to the $+1$ eigenvectors) is:

$$E_{00} = \lambda_{XZ}(+1)(+1) + \lambda_{XI}(+1) + \lambda_{IZ}(+1) + \lambda_{II}. \tag{13}$$

By preparing the system in superpositions of eigenstates and evolving under $\widetilde{H}$, we can measure energy differences between pairs of eigenstates using robust phase estimation. These energy differences form a linear system that can be inverted to recover the individual Pauli coefficients $\lambda_{XZ}, \lambda_{XI}, \lambda_{IZ}, \lambda_{II}$.

4.4 General Few-Qubit Case

This construction extends to any constant number of qubits. For a k-qubit Hamiltonian where $k = \mathcal{O}(1)$, we can choose an appropriate set of 2^k unitaries to reshape the Hamiltonian into a form where only one type of Pauli operator acts on each qubit. The resulting effective Hamiltonian has eigenstates that are product states of local Pauli eigenstates, which we can prepare and manipulate directly. The total evolution time required to learn all parameters to precision ε remains $T = \mathcal{O}(\varepsilon^{-1} \log(1/\delta))$ for any constant $k = \mathcal{O}(1)$.

5 Scaling to Many-Body Systems: Divide and Conquer

The few-qubit protocol achieves Heisenberg-limited scaling, but its computational cost grows exponentially with the number of qubits involved. For an n-qubit system, directly applying this approach would be intractable. The key idea that enables scaling to large system size n is a divide-and-conquer strategy: we reshape the Hamiltonian to decouple it into many non-interacting patches, each involving only a few qubits.

5.1 Illustrative Example: The Heisenberg Chain

Consider a one-dimensional chain of n qubits with nearest-neighbor interactions, described by an inhomogeneous Heisenberg Hamiltonian:

$$H = \sum_{j=1}^{n-1} \left(J_j^x X_j X_{j+1} + J_j^y Y_j Y_{j+1} + J_j^z Z_j Z_{j+1} \right) + \sum_{j=1}^{n} h_j Z_j, \tag{14}$$

where J_j^x, J_j^y, J_j^z, h_j are unknown coupling constants we wish to learn.

5.2 Learning a Single Bond

Suppose we want to learn the couplings between qubits 1 and 2, namely J_1^x, J_1^y, J_1^z. The challenge is that qubits 1 and 2 are entangled with the rest of the chain through the coupling between qubits 2 and 3. To break this entanglement, we employ a **Pauli twirl** on qubit 3. We choose the four Pauli operators as our reshaping unitaries:

$$U_1 = I, \quad U_2 = X_3, \quad U_3 = Y_3, \quad U_4 = Z_3, \tag{15}$$

with equal weights $w_k = 1/4$ for $k = 1, 2, 3, 4$. The key property of Pauli twirling is the following identity. For any single-qubit operator σ acting on qubit j, we have:

$$\frac{1}{4} \sum_{P \in \{I,X,Y,Z\}} (P_j \otimes I) \cdot (\sigma_j \otimes A) \cdot (P_j \otimes I) = \mathrm{tr}(\sigma) \cdot I_j \otimes A \tag{16}$$

for any operator A on the other qubits. Applying this to H with the twirl on qubit 3:

$$\widetilde{H} = \frac{1}{4} \sum_{P \in \{I,X,Y,Z\}} P_3 H P_3 \tag{17}$$

$$= (J_1^x X_1 X_2 + J_1^y Y_1 Y_2 + J_1^z Z_1 Z_2) + h_1 Z_1 + h_2 Z_2 \tag{18}$$

$$+ \sum_{j=4}^{n-1} \left(J_j^x X_j X_{j+1} + J_j^y Y_j Y_{j+1} + J_j^z Z_j Z_{j+1} \right) + \sum_{j=4}^{n} h_j Z_j. \tag{19}$$

All terms involving qubit 3 have vanished. The reshaped Hamiltonian decomposes as:

$$\widetilde{H} = \widetilde{H}_{1,2} + \widetilde{H}_{\geq 4}, \tag{20}$$

where $\widetilde{H}_{1,2}$ acts only on qubits 1 and 2, while $\widetilde{H}_{\geq 4}$ acts only on qubits 4 and beyond. Since these sectors commute, they evolve independently.

With qubits 1 and 2 decoupled from the rest of the system, we can apply the few-qubit learning protocol to estimate all parameters on this bond with Heisenberg-limited scaling.

5.3 Parallel Learning Through Strategic Decoupling

The power of this approach emerges when we decouple multiple bonds simultaneously. By applying Pauli twirls on qubits at regular intervals (say, qubits 3, 6, 9, ...), we break the chain into isolated two-qubit islands. Specifically, we reshape using:

$$U_1 = I, \quad U_2 = \bigotimes_{k=1}^{\lfloor n/3 \rfloor} X_{3k}, \quad U_3 = \bigotimes_{k=1}^{\lfloor n/3 \rfloor} Y_{3k}, \quad U_4 = \bigotimes_{k=1}^{\lfloor n/3 \rfloor} Z_{3k}, \tag{21}$$

with weights $w_k = 1/4$. The resulting Hamiltonian decomposes into non-interacting patches:

$$\widetilde{H} = \widetilde{H}_{1,2} + \widetilde{H}_{4,5} + \widetilde{H}_{7,8} + \cdots, \tag{22}$$

where each patch $\widetilde{H}_{j,j+1}$ describes an isolated two-qubit system.

Because these patches are non-interacting, we can perform parameter estimation on all of them simultaneously. This parallelization is the key to achieving system-size-independent complexity.

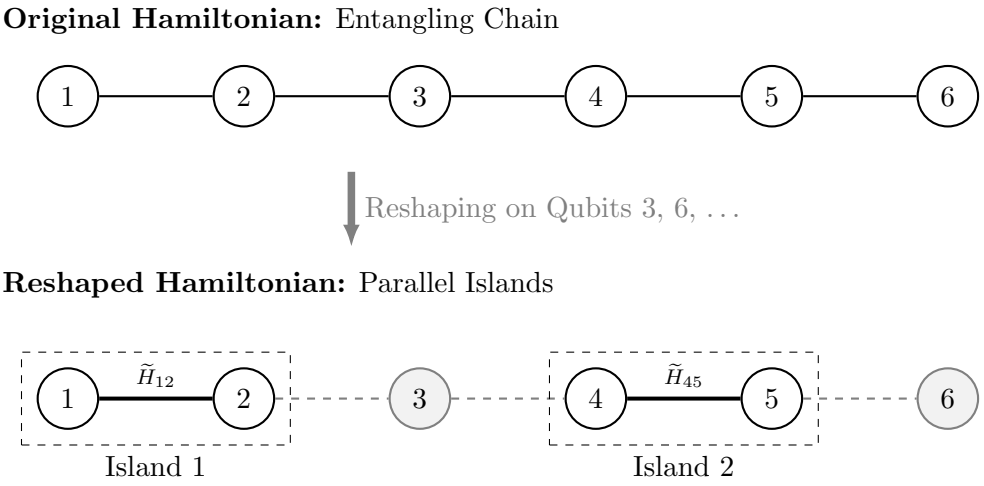

Figure 1: The divide-and-conquer strategy. By applying Pauli twirls on buffer qubits (shown in gray), we decouple the system into isolated islands. These islands can be learned in parallel.

5.4 Extension to General Geometries

The divide-and-conquer strategy extends beyond one-dimensional chains. For any geometry, where each qubit participates in at most $d = \mathcal{O}(1)$ interaction terms, we can identify an appropriate set of buffer qubits to decouple the system.

Graph Coloring: We can formalize this using graph theory. Construct an interaction graph where vertices represent interaction terms and edges connect terms that share qubits. A proper vertex coloring of this graph with χ colors allows us to partition the terms into χ non-interacting groups. By twirling appropriate qubits, we can isolate any one of these groups. For geometrically local Hamiltonians in D dimensions, the interaction graph has bounded chromatic number $\chi = \mathcal{O}(1)$ independent of system size. This ensures that we can always decouple the system into a constant number of non-interacting patches, regardless of how large n grows.

6 Final Result

We now state our main result (the SPAM robustness comes from the SPAM robustness of Robust Phase Estimation; this will be developed in detail in PSET 4).

Theorem 1 (Heisenberg-Limited Learning)**.** *Consider an unknown n-qubit Hamiltonian $H = \sum_{a=1}^{M} \lambda_a E_a$ where each E_a is a Pauli operator acting on $\mathcal{O}(1)$ qubits, each qubit is acted upon by $\mathcal{O}(1)$ terms, and $|\lambda_a| \leq 1$ for all a. There exists an algorithm robust to state preparation and measurement errors that, for any individual parameter λ_a, produces an estimate $\hat{\lambda}_a$ satisfying:*

$$\Pr\left[|\hat{\lambda}_a - \lambda_a| \leq \varepsilon\right] \geq 1 - \delta \tag{23}$$

using total evolution time:

$$T = \mathcal{O}\left(\frac{1}{\varepsilon} \log\left(\frac{1}{\delta}\right)\right). \tag{24}$$

This scaling is independent of the system size n and the number of parameters $M = \Theta(n)$.