PSET 3

Due on **Friday, November 14th, 2025 at 11:59 pm PT**

Ph 220: Quantum Learning Theory, Fall 2025 (Hsin-Yuan Huang) Version: **1.0**

This problem set will cover concepts from quantum threshold search, online learning of quantum states, and shadow tomography. If you find that you are running low on time to finish all the problems, our recommendation is to try to aim for breadth rather than depth – e.g., it is better to complete a few parts of each of the three questions, than to completely solve one of the three questions and skip the others. Below we provide hints for the various problems in this assignment. While these may help you solve the problems more easily, you are not required to follow the hints as long as the proofs you provide are correct.

1 (25 pts.) Learning Parameterized Quantum States

Motivation: Full quantum state tomography requires a number of samples exponential in the number of qubits, which is generally intractable. However, many quantum states encountered in physics and quantum algorithms are not arbitrary but belong to families described by a small number of parameters. In this problem, we will explore how to leverage this structure to learn an unknown state from such a family using only $\tilde{\mathcal{O}}(n)$ copies. This problem will combine several powerful ideas taught in class: covering nets, efficient Gibbs state preparation, online learning, and shadow tomography.

Setup: We are given classical knowledge of a family of n-qubit pure states parameterized by a vector $\vec{x} = (x_1, \ldots, x_k)$ with $\|\vec{x}\|_2 \leqslant 1$, where the number of parameters is $k = \mathcal{O}(\log n)$. The parameterized state is given by a polynomial-depth quantum circuit $U(\vec{x})$ such that $|\psi(\vec{x})\rangle = U(\vec{x})\,|0\rangle^{\otimes n}$. We are also told this family is Lipschitz continuous:

$$\| \, |\psi(\vec{x})\rangle - |\psi(\vec{x}')\rangle \, \|_2 \leqslant \|\vec{x} - \vec{x}'\|_2 \, .$$

We are given $\tilde{\mathcal{O}}(n)$ copies of an unknown state ρ with the promise that $\rho = |\psi(\vec{x}_{\text{true}})\rangle \langle\psi(\vec{x}_{\text{true}})|$ for some unknown $\vec{x}_{\text{true}} \in \mathbb{R}^k$ with $\|\vec{x}\|_2 \leqslant 1$. Our goal is to find a parameter vector $\vec{x}^*$ such that

$$\|\rho - |\psi(\vec{x}^*)\rangle\langle\psi(\vec{x}^*)| \, \|_1 < 0.01 \, ,$$

where $\|\cdot\|_1$ is the trace norm.

1.A. (5 pts.) **(Constructing a Covering Net)** First, we must discretize the continuous parameter space. Show that it is possible to construct a finite set of parameter vectors $\mathcal{N} \subset \{\vec{x} \in \mathbb{R}^k \mid \|\vec{x}\|_2 \leqslant 1\}$, called a net, such that for any potential true parameter vector $\vec{x}_{\text{true}}$, there exists a net vector $\vec{x}_i \in \mathcal{N}$ satisfying $\| \, |\psi(\vec{x}_{\text{true}})\rangle\langle\psi(\vec{x}_{\text{true}})| - |\psi(\vec{x}_i)\rangle\langle\psi(\vec{x}_i)| \, \|_1 \leqslant 0.005$. Prove that the size of this net, $M = |\mathcal{N}|$, is at most polynomial in n. Let the set of states corresponding to this net be $\mathcal{S}_{\text{net}} = \{|\psi_i\rangle\}_{i=1}^M$.

1.B. (5 pts.) **(Computationally Efficient Online Learning for Rank-1 Observables)** Assume access to a $\text{poly}(n)$-time quantum algorithm that can prepare a single copy of the state $\exp(-H)/\operatorname{tr}\exp(-H)$ given any Hermitian operator H with polynomial rank (such an algorithm was developed by Brandão et al. in "Quantum SDP Solvers"). Using this oracle, describe a complete, computationally efficient quantum algorithm for online learning of quantum states when the observables are restricted to be rank 1. Your proof of correctness can cite as a black box any guarantees that were stated in class about matrix multiplicative weights.

1.C. (7 pts.) **(Computationally Efficient Shadow Tomography for Rank-1 Observables)** The complete shadow tomography protocol taught in class is only guaranteed to be sample-efficient; its computational complexity can be exponential. However, by replacing its internal online learning subroutine with your efficient implementation from the previous question, we can make the entire protocol computationally efficient for our specific problem. Describe how to create a quantum algorithm for shadow tomography for observables $O_1 = |\psi_1\rangle\langle\psi_1|, \ldots, O_M = |\psi_M\rangle\langle\psi_M|$ that runs in $\text{poly}(n)$ time, uses $\tilde{\mathcal{O}}(n)$ copies of an unknown state ρ, and predicts the expectation values $\text{tr}(|\psi_i\rangle\langle\psi_i|\,\rho)$ up to any small constant error for all $i \in \{1, \ldots, M\}$. You may use Theorem 115 from the lecture notes as a black box (while we did not stipulate that the blended measurements algorithm is computationally efficient, you may assume that without proof for this problem).

1.D. (8 pts.) **(The Full Learning Algorithm)** You now have all the necessary components. Combine them to develop a final quantum learning algorithm that solves the original problem of learning parametrized quantum states with high probability in $\text{poly}(n)$ time. Your description should first outline the complete sequence of steps. Then, provide a proof of correctness, showing that the output $\vec{x}^*$ satisfies the desired accuracy bound. Your proof should also justify the sample complexity, explaining why the number of copies of ρ needed is only $\tilde{\mathcal{O}}(n)$.

2 (40 PTS.) FUN WITH ONLINE MACHINE LEARNING

Motivation: In class, we introduced the online learning of quantum states, which allows us to learn about an unknown state ρ by sequentially predicting the outcomes of measurements $O^{(t)}$ and receiving feedback. In this problem, we will explore this model, its properties and its implications in more detail.

First, you will derive the Matrix Multiplicative Weights (MMW) update rule from the more general principle of Online Mirror Descent.

Second, you will use the performance guarantees of this algorithm to prove a powerful and surprising result: *every* n-qubit state, no matter how complex, admits a succinct classical description (as a simple Gibbs state) that is good for predicting all Pauli observables.

Finally, you will show that this power has computational limits. You will prove that any *computationally efficient* algorithm that achieves these learning guarantees would imply an efficient quantum algorithm for 3-SAT, which is believed to be impossible.

Setup: We are in the online learning setting. There is an unknown n-qubit state ρ. At each timestep $t = 1, \ldots, T$:

(a) The learner maintains a hypothesis state $\sigma^{(t)}$ without knowledge of $O^{(t)}$.

(b) The adversary presents an observable $O^{(t)}$ with eigenvalues in $[0, 1]$ without knowledge of $\sigma^{(t)}$.

(c) The learner predicts $p_t = \mathsf{Tr}(O^{(t)}\sigma^{(t)})$.

(d) The learner receives the true expectation value $b_t = \mathsf{Tr}(O^{(t)}\rho)$.

(e) The learner suffers a loss $l_t(\sigma^{(t)}) = |p_t - b_t|$, and updates the hypothesis to $\sigma^{(t+1)}$.

2.A. (10 PTS.) **(Deriving the MMW Update Rule)** The Matrix Multiplicative Weights (MMW) algorithm can be derived from the Online Mirror Descent (OMD) framework. The OMD update rule with learning rate $\eta > 0$ is:

$$\sigma^{(t+1)} = \arg\min_{\sigma} \left\{ \eta \cdot \mathsf{Tr}(G_t\sigma) + D(\sigma||\sigma^{(t)}) \right\}$$

where $D(\sigma||\sigma^{(t)}) = \mathsf{Tr}(\sigma(\log\sigma - \log\sigma^{(t)}))$ is the quantum relative entropy, and G_t is the gradient (or a subgradient) of the loss function at $\sigma^{(t)}$. For the absolute loss $l_t(\sigma) = |\mathsf{Tr}(O^{(t)}\sigma) - b_t|$, we can use the subgradient given by $G_t = s_t O^{(t)}$, where $s_t = \mathsf{sign}(\mathsf{Tr}(O^{(t)}\sigma^{(t)}) - b_t)$ is a scalar. The optimization is over all possible density matrices with unit trace.

Show that this OMD update rule leads to the MMW update rule. That is, starting from a maximally mixed state $\sigma^{(1)} = \mathbb{I}/2^n$, prove that the hypothesis state at the beginning of step t is given by:

$$\sigma^{(t)} = \frac{\exp\left(-\eta \sum_{t'=1}^{t-1} s_{t'} O^{(t')}\right)}{\mathsf{Tr}\left(\exp\left(-\eta \sum_{t'=1}^{t-1} s_{t'} O^{(t')}\right)\right)}.$$

2.B. (15 PTS.) **(Succinct Classical Descriptions for All Quantum States)** A key result for the MMW algorithm is the following mistake bound (which you may assume without proof):

Theorem (Mistake Bound): For any sequence of observables $O^{(t)}$ with eigenvalues in $[0, 1]$, and for any unknown state ρ, the MMW algorithm (with an appropriate η) guarantees that the number of "mistakes", i.e., timesteps t where $|\mathsf{Tr}(O^{(t)}\sigma^{(t)}) - \mathsf{Tr}(O^{(t)}\rho)| > \epsilon$, is at most $M' = \mathcal{O}(n/\epsilon^2)$.

Use this theorem to prove the following two-part claim, which shows that *every* n-qubit state has a simple, approximate, succinct classical description for predicting all 4^n Pauli observables:

(1) For any n-qubit state ρ and any $\epsilon > 0$, there exists an n-qubit "Gibbs state" σ^* of the form

$$\sigma^* = \frac{\exp\left(-\sum_{i=1}^{M} \alpha_i P_i\right)}{\mathsf{Tr}\left(\exp\left(-\sum_{i=1}^{M} \alpha_i P_i\right)\right)}$$

where $M = \mathcal{O}(n/\epsilon^2)$, each P_i is a Pauli operator ($P_i \in \{I, X, Y, Z\}^{\otimes n}$), and $\alpha_i \in \mathbb{R}$. This state σ^* approximates ρ on *all* 4^n Pauli observables, i.e.,

$$|\mathsf{Tr}(P\sigma^*) - \mathsf{Tr}(P\rho)| \leqslant \epsilon \quad \text{for all } P \in \{I, X, Y, Z\}^{\otimes n}$$

(2) Show that the state σ^* from part (1) constitutes a succinct classical description of ρ by proving that it can be specified using only $\tilde{\mathcal{O}}(n^2/\epsilon^2)$ classical bits. This succinct classical description of ρ can be used to predict the expectation values of all 4^n Pauli observables of ρ up to ϵ additive error.

2.C. (15 PTS.) **(Computational Hardness of Efficient Online Learning)** The MMW update rule you derived in part (a) may not be computationally efficient since preparing these Gibbs states may be computationally hard. In this problem, you will show that this is unavoidable. Specifically, you will prove that any *computationally efficient* quantum algorithm that solves the online learning problem (even when restricted to few-body observables) would imply an efficient quantum algorithm for an NP-hard problem.

Assume for contradiction that you have an efficient quantum algorithm $\mathcal{A}$ that runs in $\mathrm{poly}(n, T)$ time. At each time step t, $\mathcal{A}$ can be used to efficiently produce any number of copies of $\sigma^{(t)}$. This algorithm $\mathcal{A}$ comes with the following guarantee. Given any sequence of few-body observables $O^{(1)}, \ldots, O^{(T)}$ with eigenvalues in $[0, 1]$, and feedback $b_1, \ldots, b_T$ that is realizable (i.e., there exists some n-qubit state ρ such that $b_t = \mathrm{Tr}(O^{(t)}\rho)$ for all t), $\mathcal{A}$ will make at most $M = \mathcal{O}(n/\epsilon^2)$ "mistakes" (i.e., predictions $p_t = \mathrm{Tr}(O^{(t)}\sigma^{(t)})$ where $|p_t - b_t| > \epsilon$).

You will use this hypothetical $\mathcal{A}$ to solve the following NP-hard problem:

> *Theorem (Unique 3-SAT): Given a 3-SAT formula ϕ on n variables with $K = \mathrm{poly}(n)$ clauses, which is promised to have a* unique *satisfying assignment $z^* \in \{0,1\}^n$, finding z^* is NP-hard under randomized reduction.*

Describe a complete, polynomial-time quantum algorithm that uses the hypothetical quantum online learning algorithm $\mathcal{A}$ as a subroutine to find the unique solution z^*.

(Hint: Map each 3-SAT clause C_j to a 3-local, Z-diagonal observable O_j. What happens if you feed $\mathcal{A}$ a sequence consisting of these observables repeated many, many times?)

Your answer should explain:

(a) How you define the true state ρ, the sequence of observables $\{O^{(t)}\}$, and the feedback $\{b_t\}$ feed to $\mathcal{A}$.

(b) Why this learning problem is guaranteed to be realizable, allowing you to use $\mathcal{A}$'s guarantee (note: you can achieve this by producing the correct feedback even without knowing z^*).

(c) How you use the mistake bound (and the pigeonhole principle) to argue that the learner's state must converge to one that has low error on all the clause observables.

(d) How you use the learner's final, converged state to extract the unknown bits of z^*.

3 (35 PTS.) QUANTUM THRESHOLD DECISION AND HYPOTHESIS SELECTION

Motivation: In class, we have seen how to perform quantum threshold search for m observables with $O(\log^2(m))$ copies of an n-qubit quantum state ρ. In this problem, we will look at a decision version of quantum threshold search that can be solved using only $O(\log(m))$ copies. This will allow us to identify and learn an unknown quantum state $\rho \in C$ that is chosen from a discrete set of candidates $C = \{\sigma_i\}_{i=1}^m$ with sample complexity $O(\log(m))$. This matches the intuition that we need at least $\Omega(\log m)$ bits of information to identify an element from m possibilities.

Setup: Let $\epsilon, \delta \in (0, 1/2)$. We are given a set of "event/threshold" pairs $(A_i, \theta_i), 1 \leqslant i \leqslant m$, where $0 \leqslant A_i \leqslant I$ are n-qubit observables and $\theta_i \in [0, 1]$ are the threshold values. We are further given gap values $0 \leqslant \eta_1, \ldots, \eta_m \leqslant 1$. The quantum threshold decision problem is to take N copies of an n-qubit state ρ, and correctly output

(a) "Yes, there exists $1 \leqslant j \leqslant m$ such that $|\operatorname{tr}(\rho A_j) - \theta_j| > \eta_j$."; or else,

(b) "No, $|\operatorname{tr}(\rho A_i) - \theta_i| \leqslant \eta_i + \epsilon$ for all $1 \leqslant i \leqslant m$.",

with failure probability at most δ. In the following, we will construct an algorithm that solves quantum threshold decision with $N = O(\log(m/\delta)/\epsilon^2)$ samples.

3.A. (8 PTS.) **(Multi-copy measurement)** For any $1 \leqslant i \leqslant m$, use the one-copy measurement $(A_i, I - A_i)$ to construct an N-copy measurement $(A_i', I - A_i')$ such that

(a) if $|\operatorname{tr}(\rho A_i) - \theta_i| > \eta_i + \epsilon$, then $\operatorname{tr}(\rho A_i') \geqslant 1 - \delta/2$; and

(b) if $|\operatorname{tr}(\rho A_i) - \theta_i| \leqslant \eta_i$, then $\operatorname{tr}(\rho A_i') \leqslant \delta^3/(16m)$.

Calculate the number of copies N you need and show that $N = O(\log(m/\delta)/\epsilon^2)$ suffices.

Hint: Use the Chernoff bound for IID Bernoulli variables $X_i, 1 \leqslant i \leqslant N$: $\mathbf{Pr}[|\frac{1}{N}\sum_{i=1}^N X_i - \mathbb{E}[X_i]| \geqslant t] \leqslant 2e^{-2t^2N}$. Use triangle inequality to lower bound the distance between the empirical mean and expectation value via the threshold value.

3.B. (2 PTS.) **(Make a decision that is not too optimistic)** Intuitively, the Yes case should correspond to at least one A'_i happening, while the No case corresponds to all A'_i not happening. Following this intuition, we make a decision by the measurement $(B, I-B)$ where B is the projector onto the span of eigenvectors of $\sum_{i=1}^m A'_i$ with eigenvalue at least $v \in (0, m)$. Show that for any ρ, $\mathrm{tr}(\rho B) \leqslant \mathrm{tr}(\rho \sum_{i=1}^m A'_i)/v$.

3.C. (5 PTS.) **(Not too pessimistic either)** Show that for any ρ, $\mathrm{tr}(\rho B) \geqslant \max_{1\leqslant i\leqslant m} \mathrm{tr}(\rho A_i) - 2\sqrt{v}$.
Hint: Let A_j be a maximizer of $\mathrm{tr}(\rho A'_j)$. Define $\delta = 1 - \mathrm{tr}(\rho A'_j)$ and $\beta = \mathrm{tr}(\rho(I-B)A'_j(I-B))$. First, show that $\beta \leqslant v$. Then, use the matrix inequality $\mathrm{tr}(\rho XY) \leqslant \sqrt{\mathrm{tr}(\rho X)}\sqrt{\mathrm{tr}(\rho Y^\dagger XY)}, \forall X, Y \geqslant 0$, to show that $\mathrm{tr}(\rho(I-B)) \leqslant \sqrt{1-\delta}\sqrt{\beta} + \sqrt{\delta}\sqrt{\mathrm{tr}(\rho(I-B)) - \beta}$.

3.D. (8 PTS.) **(Quantum threshold decision)** Choose a suitable v and give a quantum algorithm that solves quantum threshold decision with success probability $1-\delta$ using $N = O(\log(m/\delta)/\epsilon^2)$ samples.

Now we use quantum threshold decision to study the problem of hypothesis selection. Consider a discrete set of n-qubit quantum states $C = \{\sigma_i\}_{i=1}^m$ where the minimal distance between two candidate states $\alpha = \min_{i\neq j} d_{\mathrm{tr}}(\sigma_i, \sigma_j) > 0$ can be calculated. Here, we define trace distance as $d_{\mathrm{tr}}(\sigma_i, \sigma_j) = \frac{1}{2}\|\sigma_i - \sigma_j\|_1 \in [0, 1]$. Given copies of an unknown state $\rho \in C$, our task is to identify its label $\rho = \sigma_{i^*}$ with success probability at least $1-\delta$.

3.E. (2 PTS.) **(Holevo-Helstrom measurements)** We need observables that provide information about how good each hypothesis is. Recall that Holevo-Helstrom theorem states that there exists a set of n-qubit observables $0 \leqslant A_{ij} \leqslant I, 1 \leqslant i, j \leqslant m, i \neq j$, such that $\mathrm{tr}(\sigma_i A_{ij}) - \mathrm{tr}(\sigma_j A_{ij}) = d_{\mathrm{tr}}(\sigma_i, \sigma_j)$. We take $A_{ji} = I - A_{ij}$ and these observables can be calculated from the set C that we know. We can characterize how bad a hypothesis σ_k is by $\Delta_k = \max_{i<j} |\mathrm{tr}(\rho A_{ij}) - \mathrm{tr}(\sigma_k A_{ij})|$. Show that $\Delta_{i^*} = 0$ and $\Delta_k \geqslant \alpha$ for any $k \neq i^*$.

3.F. (10 PTS.) **(Quantum hypothesis selection)** Choose suitable event/threshold pairs and gap values and use quantum threshold decision multiple times to solve the hypothesis selection problem. Show that $N = O(\log(m/\delta)/\alpha^2)$ samples suffice.
Hint: You may find the quantum union bound useful.