PSET 5

Due on **Friday, December 12, 2025 at 11:59 pm PT (last day of extension is Dec. 14)**

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

This problem set covers concepts from hardness of learning and quantum advantage in learning. Due to Caltech's required grade submission date, all submissions must be received by **December 14th at 11:59 pm PT**. If you find yourself running short on time to complete all problems, we recommend aiming for breadth rather than depth. For instance, it is better to complete several parts across different problems than to fully solve one problem while skipping the others entirely. Below we provide hints for the various problems in this assignment. While these hints may help you solve the problems more efficiently, you are not required to follow them as long as your proofs are correct.

1 (60 pts.) Hardness of Classifying Quantum Phases of Matter

Introduction. In this problem, you will develop a clearer understanding of the computational complexity underlying the classification of quantum phases of matter. You will show how the path-recording oracle can be used to prove the gluing lemma: two smaller Haar-random unitaries can be glued together to closely approximate a larger Haar-random unitary. Then, you will show how the gluing lemma allows us to construct pseudorandom unitaries (PRUs) in very low circuit depth. Finally, you will use this to establish quantum computational hardness for classifying topological phases of matter.

Part 1: Gluing lemma. Consider two independent Haar-random unitaries V and W acting on three subsystems A, B, C, where V acts on AB and W acts on BC. We would like to understand the random unitary $W \cdot V$ using the path-recording oracle framework covered in class.

1.A. (5 pts.) Write down the pure output state of a quantum algorithm that makes t queries to $W \cdot V$ using two separate path-recording oracles for W and V. There should be four purifying registers W_X, W_Y, V_X, V_Y associated with the two path-recording oracles. Prove that the density matrix obtained by tracing out all purifying registers is close to the density matrix of a quantum algorithm making t queries to $W \cdot V$ (averaged over Haar-random W and V) up to trace distance $\mathcal{O}(t^2/2^{|AB|} + t^2/2^{|BC|})$. Briefly comment on how the pure output state using two path-recording oracles differs from the pure output state of the quantum algorithm that makes t queries to a single path-recording oracle simulating a Haar-random unitary U acting on ABC using purifying registers U_X, U_Y.

1.B. (7.5 pts.) By construction of the path-recording oracle for simulating Haar-random V on AB, the output bitstrings $y_1, \ldots, y_t$ stored in the purifying register V_Y are all distinct on the bits corresponding to subsystem AB. Now consider the projector Π_{V_Y} acting on the purifying register V_Y that projects onto the subspace where these output bitstrings are also distinct when restricted to subsystem B alone (that is, for any two queries $i \neq j$, the bitstrings y_i and y_j differ in at least one bit within subsystem B). Prove that applying the projector Π_{V_Y} to the pure output state after t queries only changes the state by $\mathcal{O}(t^2/2^{|B|})$ in trace distance. (Hint: use the gentle measurement lemma.)

1.C. (10 pts.) Consider the projector Π_{U_Y} acting on the purifying register U_Y associated with a Haar-random unitary U acting on ABC. This projector projects onto the subspace where the output bitstrings $y_1, \ldots, y_t$ are all distinct when restricted to subsystem BC. Prove that applying the projector Π_{U_Y} to the pure output state of a quantum algorithm after t queries to a single path-recording oracle associated with U only changes the state by $\mathcal{O}(t^2/2^{|BC|})$ in trace distance. Then, explicitly construct an isometry acting only on the purifying register that takes the projected pure output state with purifying registers W_X, W_Y, V_X, V_Y from the previous part to the projected pure output state with purifying registers U_X, U_Y (up to a close-to-1 scalar factor).

1.D. (5 pts.) Prove that the density matrix of a quantum algorithm making t queries to $W \cdot V$ (averaged over Haar-random W acting on BC and Haar-random V acting on AB) differs from the density matrix of a quantum algorithm making t queries to U (averaged over Haar-random U acting on ABC) by at most a trace distance of $\mathcal{O}(t^2/2^{|B|})$.

Part 2: Shallow PRUs. We can now use the gluing lemma to construct shallow PRUs on any geometry. For simplicity, we focus on geometries defined by finite-dimensional square lattices.

1.E. (7.5 pts.) We begin by considering n qubits arranged on a 1D chain. Define an n-qubit random unitary using two layers of Haar-random unitaries acting on small blocks of qubits of size 2ξ with overlaps of ξ qubits:

$$(U_2 \cdot U_4 \cdot \ldots)(U_1 \cdot U_3 \cdot U_5 \ldots),$$

where each U_i acts on 2ξ qubits, U_1 acts on qubits $\{1, 2, \ldots, 2\xi\}$, U_2 acts on qubits $\{\xi+1, \xi+2, \ldots, 3\xi\}$, and so on. Use the gluing lemma to prove that the density matrix of a quantum algorithm making t queries to this random unitary (averaged over the randomness) differs from the density matrix of a quantum algorithm making t queries to an n-qubit Haar-random unitary by at most a trace distance of $\mathcal{O}((n/\xi) \cdot t^2/2^{\xi})$.

1.F. (5 PTS.) Prove how the above construction and the resulting statement can be generalized to any geometry defined by a finite-dimensional square lattice.

1.G. (10 PTS.) Prove that (without using computational assumptions) a polynomial-time quantum algorithm making polynomially many queries to a Haar-random unitary can be simulated efficiently by a polynomial-time quantum algorithm. (Hint: you may assume as a black box that there exist efficient quantum circuit implementations of the path-recording oracle.) Use this to prove the following: given $\xi = \omega(\log n)$, replacing the small Haar-random unitaries acting on blocks of 2ξ qubits by small PRUs on 2ξ qubits (which are indistinguishable from Haar-random unitaries by any $2^{o(\xi)}$-time quantum algorithm[1]) yields an n-qubit PRU that is indistinguishable from an n-qubit Haar-random unitary by any $\mathrm{poly}(n)$-time quantum algorithm. Combine this with the fact that PRUs on k qubits can be constructed in $\mathrm{poly}(k)$ depth to prove that n-qubit PRUs (indistinguishable from n-qubit Haar-random unitaries by any $\mathrm{poly}(n)$-time quantum algorithm) can be constructed in $\mathrm{poly}(\log n)$ depth.

Part 3: Classifying topological phases of matter. Consider an unknown state $|\psi\rangle$ that is obtained by applying a 2D shallow quantum circuit either to $|0^n\rangle$ (trivial phase) or to the toric code ground state (topological phase).

1.H. (5 PTS.) Prove that when the shallow quantum circuit has depth $d = \mathcal{O}(1)$, the classification can be performed in polynomial time. (Hint: use the trivial state learning algorithm covered in class.)

1.I. (5 PTS.) Prove that when the shallow quantum circuit has depth $d = \mathrm{poly}(\log n)$, the classification cannot be performed in polynomial time, assuming subexponential hardness for quantum-secure pseudorandom functions.

2 (40 PTS.) QUANTUM ADVANTAGE IN QUANTUM PCA

Introduction. Quantum principal component analysis (PCA) is one of the first quantum machine learning algorithms developed. While it was initially believed to offer significant quantum speedups for analyzing classical data, Ewin Tang later showed that quantum PCA can at most offer a modest polynomial speedup in that setting. However, quantum PCA may still offer significant advantages for analyzing quantum data. In this problem, you will prove that this is the case.

Part 1: Quantum Principal Component Analysis. Read the seminal paper *Quantum principal component analysis* (https://arxiv.org/abs/1307.0401) to answer the following questions.

2.A. (3 PTS.) Given a quantum state ρ, what is the principal component $|\psi\rangle$ of ρ?

2.B. (3 PTS.) Consider the n-qubit state $\rho = \frac{1}{2}|0\rangle\langle 0| \otimes \frac{I}{2^{n-1}} + \frac{1}{2}|1\rangle\langle 1| \otimes |\phi\rangle\langle\phi|$, where the first register is a single qubit and $|\phi\rangle$ is an arbitrary unknown pure state on the remaining $n-1$ qubits. How many copies of ρ are needed to estimate the expectation value $\langle\psi|Z_1|\psi\rangle$ (where Z_1 acts on the first qubit) up to a small constant error using quantum PCA? What is the quantum circuit depth required by quantum PCA?

2.C. (4 PTS.) What are the conditions on ρ required by quantum PCA to estimate properties $\langle\psi|O|\psi\rangle$ of the principal component $|\psi\rangle$ of ρ to a small constant error for an efficiently measurable observable O in polynomial time?

Part 2: Quantum Advantage. We are now ready to establish quantum advantage in quantum PCA. We will first establish the necessary mathematical tools. Then, we will prove a quantum advantage in distinguishing pure states from maximally mixed states. Finally, we will prove quantum advantage in quantum PCA.

2.D. (10 PTS.) Prove by induction on t that for any collection of n-qubit pure states $|\psi_1\rangle, \ldots, |\psi_t\rangle$,

$$\sum_{\sigma \in S_t} \mathrm{Tr}\left(\sigma\left(\bigotimes_{i=1}^{t} |\psi_i\rangle\langle\psi_i|\right)\right) \geqslant 1,$$

where σ is a permutation over t copies of n-qubit systems and S_t denotes the symmetric group on t elements.

2.E. (10 PTS.) Consider an unknown n-qubit state ρ that is either a Haar-random state or the maximally mixed state $I/2^n$. Prove that a classical agent that can measure a single sample of ρ to obtain classical data, process the classical data, and choose the next measurement based on past outcomes requires $2^{\Omega(n)}$ samples of ρ to distinguish between the two cases with high probability.

(Hint: $\mathbb{E}_{|\psi\rangle\sim\text{Haar}}|\psi\rangle\langle\psi|^{\otimes t} = \binom{2^n+t-1}{t}^{-1} \sum_{\sigma\in S_t} \sigma$.)

2.F. (5 PTS.) Prove that a classical agent requires superpolynomial time to distinguish between an unknown state ρ that is either a short-range-entangled state generated by applying a $\mathrm{poly}(\log n)$-depth 1D quantum circuit to $|0^n\rangle$ or the maximally mixed state $I/2^n$.

2.G. (5 PTS.) Under the conditions on ρ required by quantum PCA to solve the problem in polynomial time, prove that a classical agent requires superpolynomial time to estimate the expectation value $\langle\psi|Z_1|\psi\rangle$ of the principal component $|\psi\rangle$ of an unknown efficiently-preparable state ρ to a small constant error.

[1] This is known as subexponential hardness, which is stronger than the polynomial hardness covered in class. Most cryptographic schemes are believed to satisfy subexponential hardness, meaning that one needs exponential time (not just polynomial time) to break the scheme.