Ph 220: Quantum Learning Theory
Lecture 10: Quantum Advantage in Learning from Experiments

Hsin-Yuan Huang (Robert)

California Institute of Technology

1 Introduction

We have explored the capabilities of classical and quantum algorithms for learning about quantum systems. Today, we address a complementary and equally fundamental question: what are the limits of *classical* AI agents when interacting with quantum systems? and what advantage could quantum AI agents have over classical AI agents. While quantum agents can leverage entanglement and coherent quantum memory to learn properties such as Pauli observables efficiently (as we have seen from PSET 2), we will establish today that classical agents face severe fundamental limitations. We prove that learning to predict the outcome of a chosen Pauli observable requires an exponential number of experiments for any classical learning algorithm. This result demonstrates a provable exponential quantum advantage in learning from experiments.

2 Modeling Classical Learning Agents

To establish a rigorous lower bound on classical learning, we must first formalize what it means for a classical agent to learn from quantum experiments.

2.1 Memory States and Measurement Strategies

A general classical learning agent operates through repeated cycles of interaction with a quantum system that changes its internal memory state:

1. **State:** The agent maintains a classical memory state M.
2. **Action:** Based on M, the agent selects a measurement to perform on the quantum system described by state ρ.
3. **Outcome:** The agent receives a classical measurement outcome x.
4. **Update:** The agent updates its memory state based on the observed outcome.

Mathematically, the measurement strategy of an agent in state M is described by a Positive Operator-Valued Measure (POVM) $\mathcal{M} = \{F_x\}_x$, where $\sum_x F_x = \mathbb{I}$ and $F_x \succeq 0$ for all outcomes x. The probability of observing outcome x when measuring state ρ is given by the Born rule:

$$p(x) = \mathrm{Tr}(F_x \rho). \tag{1}$$

The outcome causes a memory state transition $M \to M'$.

2.2 Reduction to Rank-1 POVMs

To simplify our analysis without loss of generality, we restrict attention to measurements with rank-1 POVM elements.

Lemma 1 (Rank-1 Simplification)**.** Any classical learning agent using general POVMs can be simulated by an agent using only rank-1 POVMs without changing the distinguishing capability.

Proof. Consider a general POVM element F_x. Since F_x is positive semidefinite, we can perform a spectral decomposition:

$$F_x = \sum_y \alpha_{xy} |\psi_{xy}\rangle\langle\psi_{xy}|, \quad \text{where } \alpha_{xy} \geq 0. \tag{2}$$

We define a refined measurement strategy with outcomes labeled by pairs (x, y). The corresponding POVM elements are:

$$\tilde{F}_{x,y} = \alpha_{xy} |\psi_{xy}\rangle\langle\psi_{xy}|. \tag{3}$$

This constitutes a valid rank-1 POVM. The original agent is simulated by measuring outcome (x, y) and retaining only the x component. Thus, we can restrict our analysis to agents using rank-1 POVM elements of the form $\{w_x |\psi_x\rangle\langle\psi_x|\}_x$, where $w_x \geq 0$ and $\sum_x w_x \langle\psi_x|\psi_x\rangle = 1$. □

2.3 The Learning Tree Framework

The graph describing the memory transition can be a complicated directed graph with many cycles. However, a classical learning agent with a general graph can always be simulated by a classical learning agent with a directed tree that only walks from the root to a leaf. This is because the new classical learning agent can simply record the entire history of the memory state of the original classical agent. he new classical agent will start at the root with depth 0. After N experiments, the new classical agent will arrive at a leaf with depth N. We call it the **learning tree**.

Definition 1 (Learning Tree)**.** A learning tree of depth N consists of:

- A root node r representing the initial state.
- Internal nodes u, each associated with a rank-1 POVM $\{w_x^{(u)} |\psi_x^{(u)}\rangle\}_x$.
- Edges $u \to_x v$ representing the transition from node u to node v upon the observation of a specific measurement outcome x.
- Leaves ℓ representing the final memory states after N experiments.

For an unknown quantum state ρ that the learning agent tries to learn, the probability of transitioning from node u to child node v is given by the Born rule:

$$\Pr(u \to_x v \mid \rho) = w_x^{(u)} \langle\psi_x^{(u)}| \rho |\psi_x^{(u)}\rangle . \tag{4}$$

The probability of reaching a specific leaf ℓ is the product of transition probabilities along the unique path from the root:

$$\Pr(\ell \mid \rho) = \prod_{(u \to_x v) \in \text{path}(r \to \ell)} \Pr(u \to_x v \mid \rho). \tag{5}$$

3 The Distinguishing Task

To demonstrate the fundamental limitations of classical learning agents, we define a hypothesis testing problem such that being able to solve this problem implies being able to predict any Pauli observables. Hence, by proving a lower bound on the number of experiments required by the classical learning agents for this hypothesis testing problem, we can obtain the same lower bound for predicting Pauli observables.

3.1 World A and World B

We consider two scenarios (worlds) that the agent must distinguish:

- **World A (Pure Noise):** The quantum system is in the maximally mixed state:

$$\rho_A = \frac{\mathbb{I}}{2^n}. \tag{6}$$

- **World B (Signal):** The quantum system is in a state biased along an unknown Pauli operator $P \in \{I, X, Y, Z\}^{\otimes n} \setminus \{I\}$ with random sign:

$$\rho_B = \rho_{P,\sigma} = \frac{\mathbb{I} + \sigma\epsilon P}{2^n}, \tag{7}$$

  where $\sigma \in \{+1, -1\}$ is chosen uniformly at random, and $\epsilon \in (0, 1]$ is a bias parameter.

3.2 Learning Objective

The agent's goal is to distinguish World A from World B with high probability after performing N experiments. Equivalently, given a learning tree of depth N, we ask: how large must N be for the probability distributions over leaves to differ significantly between the two worlds?

3.3 Reduction to Predicting Pauli Observables

It is easy to see that if a classical agent can predict the expectation value $\mathrm{Tr}(P\rho)$ for any P, then the classical agent can distinguish World A and B. Hence, the exponential lower bound for this task implies that the same exponential lower bound for the task of predicting Pauli observables.

4 Proof of Exponential Lower Bound

We now prove that classical agents require exponentially many experiments to distinguish the two worlds. This implies the following theorem via the reduction above.

Theorem 1 (Classical Hardness of Predicting Pauli's)**.** Any classical agent that can learn from N experiments on ρ to predict the expectation values of all Pauli observables $P \in \{I, X, Y, Z\}^{\otimes n}$ to ϵ error with constant success probability requires $N = \Omega(2^n/\epsilon^2)$ measurements.

4.1 Total Variation Distance

The distinguishing capability of any agent is bounded by the total variation distance (TVD) between the probability distributions over leaves induced by the two worlds. Let $p_A(\ell)$ denote the probability of reaching leaf ℓ in World A, and let $p_{B,P,\sigma}(\ell)$ denote the corresponding probability in World B for a specific Pauli P and sign σ. Since the Pauli operator P is unknown to the agent and the sign σ is random, we average over these choices. The relevant distance is:

$$\Delta = \mathbb{E}_{P \neq I} \left[\mathrm{TVD} \left(p_A, \mathbb{E}_{\sigma \in \{+1,-1\}} p_{B,P,\sigma} \right) \right], \tag{8}$$

where

$$\mathrm{TVD}(p, q) = \frac{1}{2} \sum_{\ell} |p(\ell) - q(\ell)| . \tag{9}$$

Note that we actually care only about

$$\Delta' = \mathrm{TVD} \left(p_A, \mathbb{E}_{P \neq I} \mathbb{E}_{\sigma \in \{+1,-1\}} p_{B,P,\sigma} \right) \leq \Delta. \tag{10}$$

Proving an upper bound of Δ always imply the same upper bound for Δ'. However, the other direction may not be true in general. Hence, it is stronger to establish an upper bound on Δ.

Proving an upper bound on Δ implies that after the classical agent finishes doing experiments, if someone revealed the correct P to the classical agent, it would still be hard for the classical agent to distinguish between World A and B. This means even when a significant amount of information is revealed to the classical agent, as long as the classical agent has finished the experiments, then the distinguishing task is hard to solve. In contrast, proving an upper bound on Δ' does not imply that if the correct P is revealed, then the distinguishing task remains hard. Being able to prove an upper bound on Δ means that even predicting one adversarially chosen Pauli P is hard for classical agents (not just hard for predicting all 4^n Pauli observables).

4.2 Computing Path Probabilities

We analyze the probability of reaching a leaf ℓ through a specific path from the root.

World A: In the maximally mixed state $\rho_A = \mathbb{I}/2^n$, the probability is:

$$p_A(\ell) = \prod_{(u \to_x v) \in \mathrm{path}(\ell)} w_x^{(u)} \langle \psi_x^{(u)} | \frac{\mathbb{I}}{2^n} | \psi_x^{(u)} \rangle = \prod_{(u \to_x v) \in \mathrm{path}(\ell)} \frac{w_x^{(u)}}{2^n}, \tag{11}$$

where the product is taken over all edges $(u \to_x v)$ on the path from the root to leaf ℓ.

World B: In the biased state $\rho_{B,P,\sigma} = (\mathbb{I} + \sigma\epsilon P)/2^n$, we have:

$$p_{B,P,\sigma}(\ell) = \prod_{(u \to_x v) \in \mathrm{path}(\ell)} w_x^{(u)} \langle \psi_x^{(u)} | \frac{\mathbb{I} + \sigma\epsilon P}{2^n} | \psi_x^{(u)} \rangle \tag{12}$$

$$= \prod_{(u \to_x v) \in \mathrm{path}(\ell)} \frac{w_x^{(u)}}{2^n} \left(1 + \sigma\epsilon \langle \psi_x^{(u)} | P | \psi_x^{(u)} \rangle \right) \tag{13}$$

$$= p_A(\ell) \prod_{(u \to_x v) \in \mathrm{path}(\ell)} (1 + \sigma\epsilon \langle P \rangle_{u,x}) , \tag{14}$$

where we define the expectation value $\langle P \rangle_{u,x} \equiv \langle \psi_x^{(u)} | P | \psi_x^{(u)} \rangle$.

4.3 Bounding the Distinguishing Advantage

Averaging over the random sign $\sigma \in \{+1, -1\}$, we obtain:

$$\mathbb{E}_\sigma[p_{B,P,\sigma}(\ell)] = p_A(\ell) \cdot \mathbb{E}_\sigma \left[\prod_{(u \to_x v) \in \text{path}(\ell)} (1 + \sigma\epsilon\langle P\rangle_{u,x}) \right] \tag{15}$$

$$= p_A(\ell) \cdot Q_P(\ell), \tag{16}$$

where we define

$$Q_P(\ell) = \mathbb{E}_{\sigma\in\{+1,-1\}} \prod_{(u \to_x v) \in \text{path}(\ell)} (1 + \sigma\epsilon\langle P\rangle_{u,x}). \tag{17}$$

Let $\delta_{u,x} = \epsilon\langle P\rangle_{u,x}$. Note that $|\delta_{u,x}| \leq \epsilon \leq 1$ since $|\langle P\rangle_{u,x}| \leq 1$ for any normalized Pauli operator. We can expand:

$$Q_P(\ell) = \frac{1}{2} \left[\prod_{(u \to_x v) \in \text{path}(\ell)} (1 + \delta_{u,x}) + \prod_{(u \to_x v) \in \text{path}(\ell)} (1 - \delta_{u,x}) \right]. \tag{18}$$

4.4 Application of AM-GM Inequality

We apply the arithmetic mean-geometric mean (AM-GM) inequality, which states that $\frac{a+b}{2} \geq \sqrt{ab}$ for non-negative a, b:

$$Q_P(\ell) \geq \sqrt{\left(\prod_{(u \to_x v) \in \text{path}(\ell)} (1 + \delta_{u,x}) \right) \left(\prod_{(u \to_x v) \in \text{path}(\ell)} (1 - \delta_{u,x}) \right)} \tag{19}$$

$$= \sqrt{\prod_{(u \to_x v) \in \text{path}(\ell)} (1 - \delta_{u,x}^2)}. \tag{20}$$

Using the inequality $\prod_{i=1}^N (1 - a_i) \geq 1 - \sum_{i=1}^N a_i$ for $a_i \in [0, 1)$, we obtain:

$$Q_P(\ell) \geq \sqrt{1 - \sum_{(u \to_x v) \in \text{path}(\ell)} \delta_{u,x}^2}. \tag{21}$$

For $y \in [0, 1]$, we have $1 - \sqrt{1 - y} \leq y$. Therefore:

$$|1 - Q_P(\ell)| = 1 - Q_P(\ell) \tag{22}$$

$$\leq 1 - \sqrt{1 - \sum_{(u \to_x v) \in \text{path}(\ell)} \delta_{u,x}^2} \tag{23}$$

$$\leq \sum_{(u \to_x v) \in \text{path}(\ell)} \delta_{u,x}^2 \tag{24}$$

$$= \sum_{(u \to_x v) \in \text{path}(\ell)} \epsilon^2 \langle P\rangle_{u,x}^2. \tag{25}$$

4.5 Computing the Expected Distance

The total variation distance for a fixed Pauli P is:

$$\Delta_P = \mathrm{TVD}(p_A, \mathbb{E}_\sigma p_{B,P,\sigma}) \tag{26}$$

$$= \frac{1}{2}\sum_\ell |p_A(\ell) - p_A(\ell)Q_P(\ell)| \tag{27}$$

$$= \frac{1}{2}\sum_\ell p_A(\ell)\,|1 - Q_P(\ell)| \tag{28}$$

$$\leq \frac{1}{2}\sum_\ell p_A(\ell) \sum_{(u\to_x v)\in \mathrm{path}(\ell)} \epsilon^2 \langle P\rangle_{u,x}^2. \tag{29}$$

Averaging over all non-identity Paulis:

$$\Delta = \mathbb{E}_{P\neq I}[\Delta_P] \tag{30}$$

$$\leq \frac{\epsilon^2}{2}\sum_\ell p_A(\ell) \sum_{(u\to_x v)\in \mathrm{path}(\ell)} \mathbb{E}_{P\neq I}\left[\langle P\rangle_{u,x}^2\right] \tag{31}$$

$$= \frac{\epsilon^2}{2}\sum_\ell p_A(\ell)\cdot |\mathrm{path}(\ell)|\cdot \mathbb{E}_{P\neq I}\left[\langle P\rangle^2\right], \tag{32}$$

where $|\mathrm{path}(\ell)|$ denotes the length (depth) of the path to ℓ, and we use the fact that all measurement vectors have the same expected squared overlap with a random Pauli.

4.6 Expected Squared Overlap with Random Pauli

We bound the expectation $\mathbb{E}_{P\neq I}[\langle P\rangle^2]$ for any normalized pure state $|\psi\rangle$. The set of all n-qubit Pauli operators $\{P_0 = I, P_1, \ldots, P_{4^n-1}\}$ forms an orthogonal basis for operators on n qubits. For any pure state $|\psi\rangle$, we have the completeness relation:

$$\sum_{j=0}^{4^n-1} |\langle\psi|\,P_j\,|\psi\rangle|^2 = 2^n \mathrm{Tr}(|\psi\rangle\langle\psi|^2) = 2^n. \tag{33}$$

Separating the identity contribution:

$$|\langle\psi|I|\psi\rangle|^2 + \sum_{P\neq I} |\langle\psi|\,P\,|\psi\rangle|^2 = 2^n. \tag{34}$$

Since $\langle\psi|I|\psi\rangle = \langle\psi|\psi\rangle = 1$, we obtain:

$$\sum_{P\neq I} |\langle\psi|\,P\,|\psi\rangle|^2 = 2^n - 1. \tag{35}$$

Therefore:

$$\mathbb{E}_{P\neq I}\left[|\langle\psi|\,P\,|\psi\rangle|^2\right] = \frac{2^n-1}{4^n-1} = \frac{2^n-1}{(2^n-1)(2^n+1)} = \frac{1}{2^n+1} \leq \frac{1}{2^n}. \tag{36}$$

4.7 Final Bound

Substituting this expectation into our bound:

$$\Delta \leq \frac{\epsilon^2}{2 \cdot 2^n} \sum_{\ell} p_A(\ell) \cdot |\text{path}(\ell)| \tag{37}$$

$$= \frac{\epsilon^2}{2^{n+1}} \sum_{\ell} p_A(\ell) \cdot |\text{path}(\ell)| \,. \tag{38}$$

The sum $\sum_\ell p_A(\ell) \cdot |\text{path}(\ell)|$ equals the expected depth of the tree, which is at most N (the maximum depth). Together, we have:

$$\Delta \leq \frac{N\epsilon^2}{2^{n+1}}. \tag{39}$$

4.8 Conclusion of the Proof

For the agent to distinguish the two worlds with constant success probability, we require $\Delta = \Omega(1)$. This necessitates:

$$N = \Omega\left(\frac{2^n}{\epsilon^2}\right). \tag{40}$$

Because being able to predict Pauli observables to ϵ error implies the ability to distinguish between the two worlds, the lower bound applies to the Pauli prediction task. This completes the proof of Theorem 1. You can now check the accompanying slides for physical experiments conducted on Google's quantum processor demonstrating quantum advantage in learning from experiments.