Ph 220: Quantum Learning Theory – Lecture Note 4: How to Protect Quantum Data While Using Them?

Hsin-Yuan Huang (Robert)

Caltech

In our previous lecture, we established that to predict M observables O_i that either have bounded Frobenius norm ($\mathrm{Tr}(O_i^2) = \mathcal{O}(1)$) or act on only a few qubits (e.g., k-local with $k = \mathcal{O}(1)$), we can use $N = \mathcal{O}(\log(M)/\varepsilon^2)$ copies of ρ to create classical shadows and post-process them to predict these observables. This approach is highly efficient in sample complexity, but it applies only to these restricted families of observables and necessarily destroys the quantum states in the measurement process.

For general high-weight observables such as $Z^{\otimes n}$ or $P \in \{X, Y, Z\}^{\otimes n}$, we have $\mathrm{Tr}(O_i^2) = 2^n$, and the classical shadow method requires $N = \mathcal{O}(2^n \log M)$ samples, which is exponentially large. In fact, we will see that *any* protocol measuring single copies of ρ one at a time requires $N = \Omega(2^n/\varepsilon^2)$ samples to predict all Pauli observables.

To surpass this fundamental single-copy limitation, we must perform quantum data analysis on many samples simultaneously. This involves using a quantum memory to store multiple copies of the unknown quantum state $\rho^{\otimes N}$ as our quantum dataset, and performing entangling coherent quantum operations across them. A natural question arises: can we protect and reuse our quantum data while we continue to extract new properties from them? This lecture will develop all the key algorithmic tools that enable precisely this capability.

1 Prediction Task under a Strong Promise

Let us first consider a problem with a very strong structural promise.

Q: Given M observables $O_1, \ldots, O_M$ with eigenvalues ± 1 and N samples of an unknown state ρ such that $\mathrm{Tr}(O_i\rho) \in \{-1, 0, 1\}$, how many samples N are needed to predict $\mathrm{Tr}(O_i\rho)$ for all i?

The promise $\mathrm{Tr}(O_i\rho) \in \{-1, 0, 1\}$ is extremely powerful. For an observable O_i with eigenvalues ± 1, we can write its spectral decomposition as $O_i = P_{i,+} - P_{i,-}$, where $P_{i,+}$ and $P_{i,-}$ are projectors onto the $+1$ and -1 eigenspaces, respectively, and $P_{i,+} + P_{i,-} = \mathbb{I}$. We have three cases:

- If $\mathrm{Tr}(O_i\rho) = 1$, then $\mathrm{Tr}(P_{i,+}\rho) = 1$. The state ρ must lie entirely in the $+1$ eigenspace, so $\rho = P_{i,+}\rho P_{i,+}$.
- If $\mathrm{Tr}(O_i\rho) = -1$, then $\mathrm{Tr}(P_{i,-}\rho) = 1$. The state ρ must lie entirely in the -1 eigenspace, so $\rho = P_{i,-}\rho P_{i,-}$.
- If $\mathrm{Tr}(O_i\rho) = 0$, then $\mathrm{Tr}(P_{i,+}\rho) = \mathrm{Tr}(P_{i,-}\rho) = 1/2$.

If we were to measure ρ directly in the eigenbasis of O_i, the state would collapse. In the $\mathrm{Tr}(O_i\rho) = 0$ case, we would obtain outcome $+1$ or -1 with equal probability, and the post-measurement state would be projected and destroyed, rendering it useless for subsequent measurements of O_j. The key insight is to perform a *collective gentle measurement* on the N-copy state $\sigma = \rho^{\otimes N}$.

1.1 Collective Measurement Protocol

For each observable O_i, we define a single-copy controlled unitary U_{O_i} that acts on the data state ρ and one ancilla qubit initialized to $|0\rangle$:

$$U_{O_i}(|\psi\rangle \otimes |0\rangle) = P_{i,+} |\psi\rangle \otimes |1\rangle + P_{i,-} |\psi\rangle \otimes |-1\rangle . \tag{1}$$

Here we use $|1\rangle$ and $|-1\rangle$ to label two orthogonal ancilla states, corresponding to eigenvalues ± 1.

We apply this operation coherently to all N copies of ρ in our quantum memory $\sigma = \rho^{\otimes N}$, using N ancilla qubits initialized to $|0\rangle^{\otimes N}$. The collective unitary is $U_{O_i}^{\otimes N}$:

$$\sigma' = U_{O_i}^{\otimes N}(\sigma \otimes |0\rangle\langle 0|^{\otimes N})U_{O_i}^{\dagger \otimes N}. \tag{2}$$

The resulting joint state σ' between the N data registers and the N ancilla registers is:

$$\sigma' = \sum_{\vec{\lambda},\vec{\lambda}' \in \{\pm 1\}^N} \left(\bigotimes_{k=1}^{N} P_{i,\lambda_k} \rho P_{i,\lambda'_k} \right) \otimes |\vec{\lambda}\rangle\langle\vec{\lambda}'|, \tag{3}$$

where $\vec{\lambda} = (\lambda_1, \ldots, \lambda_N)$ and $|\vec{\lambda}\rangle = |\lambda_1\rangle \otimes \cdots \otimes |\lambda_N\rangle$.

Now, instead of measuring each ancilla individually, we perform a collective three-outcome POVM on the N-ancilla register:

- $E_1 = |1...1\rangle\langle 1...1|$ (projector corresponding to outcome "1"),
- $E_{-1} = |-1...-1\rangle\langle -1...-1|$ (projector corresponding to outcome "-1"),
- $E_0 = \mathbb{I} - E_1 - E_{-1}$ (projector corresponding to outcome "0").

1.2 Analysis of Measurement Outcomes

Let us analyze the probability of each outcome for the three cases in our promise.

Case 1: $\mathrm{Tr}(O_i\rho) = 1$**.** This implies $\rho = P_{i,+}\rho P_{i,+}$ and $P_{i,-}\rho = 0$. The only non-zero term in the sum for σ' is when all $\lambda_k = 1$ and $\lambda'_k = 1$:

$$\sigma' = (P_{i,+}\rho P_{i,+})^{\otimes N} \otimes |1...1\rangle\langle 1...1| = \sigma \otimes E_1. \tag{4}$$

The outcome probabilities are:

- $\Pr[1] = \mathrm{Tr}(\sigma'(\mathbb{I} \otimes E_1)) = \mathrm{Tr}((\sigma \otimes E_1)(\mathbb{I} \otimes E_1)) = 1$.
- $\Pr[-1] = \mathrm{Tr}(\sigma'(\mathbb{I} \otimes E_{-1})) = 0$.
- $\Pr[0] = \mathrm{Tr}(\sigma'(\mathbb{I} \otimes E_0)) = 0$.

We obtain outcome "1" with certainty. The post-measurement state is σ', and tracing out the ancilla gives $\mathrm{Tr}_{\mathrm{anc}}(\sigma \otimes E_1) = \sigma$. The quantum data is completely undisturbed.

Case 2: $\mathrm{Tr}(O_i\rho) = -1$**.** This implies $\rho = P_{i,-}\rho P_{i,-}$ and $P_{i,+}\rho = 0$:

$$\sigma' = (P_{i,-}\rho P_{i,-})^{\otimes N} \otimes |-1...-1\rangle\langle -1...-1| = \sigma \otimes E_{-1}. \tag{5}$$

The outcome probabilities are $\Pr[1] = 0$, $\Pr[-1] = 1$, $\Pr[0] = 0$. We obtain outcome "-1" with certainty, and the state is undisturbed.

Case 3: $\mathrm{Tr}(O_i\rho) = 0$**.** This implies $\mathrm{Tr}(P_{i,+}\rho) = 1/2$ and $\mathrm{Tr}(P_{i,-}\rho) = 1/2$. The state σ' on the data and ancilla registers is:

$$\sigma' = \sum_{\vec{\lambda},\vec{\lambda}' \in \{\pm 1\}^N} \left(\frac{1}{2}\right)^N \left(\bigotimes_{k=1}^{N} P_{i,\lambda_k} \rho P_{i,\lambda'_k}\right) \otimes |\vec{\lambda}\rangle\langle\vec{\lambda}'|. \tag{6}$$

The probabilities for our three-outcome POVM on the ancilla register are:

- $\Pr[1] = \mathrm{Tr}_{\mathrm{anc}}(\mathrm{Tr}_{\mathrm{data}}(\sigma')E_1) = \mathrm{Tr}_{\mathrm{anc}}\left(\sum_{\vec{\lambda}} \left(\frac{1}{2}\right)^N |\vec{\lambda}\rangle\langle\vec{\lambda}|\right) E_1$. Only $\vec{\lambda} = (1, ..., 1)$ contributes, so $\Pr[1] = (1/2)^N$.
- $\Pr[-1] = \mathrm{Tr}_{\mathrm{anc}}(\dots)E_{-1}$. Only $\vec{\lambda} = (-1, ..., -1)$ contributes, so $\Pr[-1] = (1/2)^N$.
- $\Pr[0] = 1 - \Pr[1] - \Pr[-1] = 1 - 2 \cdot (1/2)^N = 1 - 2^{1-N}$.

In this case, the correct outcome is "0", which occurs with extremely high probability $1 - 2^{1-N}$.

1.3 Sample Complexity

This measurement is gentle. For any O_i, the correct outcome ("1" if $\mathrm{Tr}(O_i\rho) = 1$, "-1" if $\mathrm{Tr}(O_i\rho) = -1$, or "0" if $\mathrm{Tr}(O_i\rho) = 0$) occurs with probability $p_{\mathrm{correct}} \geq 1 - 2^{1-N}$. The failure probability for a single measurement i on the original state $\sigma_0 = \rho^{\otimes N}$ is $\varepsilon_i = 1 - p_{\mathrm{correct}}$:

$$\varepsilon_i \leq 2^{1-N} \quad \text{for all } i = 1, \dots, M. \tag{7}$$

We can now apply the Quantum Union Bound that we will prove later (Theorem 1). We perform M sequential measurements. Let P_i be the projector corresponding to the *correct* outcome for O_i (i.e., E_1, E_{-1}, or E_0). The probability of obtaining correct outcome for all M measurements is:

$$\Pr[\text{All Correct}] \geq 1 - 4\sum_{i=1}^{M} \varepsilon_i. \tag{8}$$

Using our bound $\varepsilon_i \leq 2^{1-N}$:

$$\Pr[\text{All Correct}] \geq 1 - 4\sum_{i=1}^{M} 2^{1-N} = 1 - 4M2^{1-N} = 1 - M2^{3-N}. \tag{9}$$

We want this to succeed with high probability, say $\geq 1 - \delta$. We need to choose N large enough:

$$M2^{3-N} \leq \delta \implies \frac{M}{\delta} \leq 2^{N-3} \implies N \geq 3 + \log_2\left(\frac{M}{\delta}\right). \tag{10}$$

Therefore, $N = \mathcal{O}(\log(M/\delta))$ samples is sufficient. This demonstrates that by using a quantum memory and collective measurements, we can reuse our quantum data M times with a cost that scales only logarithmically with M. The total disturbance to the state σ, by the Gentle Measurement Lemma (Lemma 1), is also bounded by $\|\sigma_0 - \sigma_M\|_1 \leq 2\sqrt{\sum \varepsilon_i} \leq 2\sqrt{M2^{1-N}}$. We will develop these two theorems in the following sections.

2 Gentle Measurement Lemma

If we perform a sequence of projective measurements, where each individual measurement has a high probability of yielding a certain success outcome, how much is the state disturbed after conditioning on the whole sequence of successes? The sequential version of the Gentle Measurement Lemma bounds this disturbance.

Lemma 1 (Gentle Measurement Lemma)**.** *Let ρ be a state. We sequentially perform M projective measurements, $\mathcal{E}_i = \{P_i, \mathbb{I} - P_i\}$, for $i = 1, \ldots, M$. Suppose for each measurement i, the desired outcome P_i happens with high probability on the original state: $\mathrm{Tr}(P_i \rho) \geq 1 - \varepsilon_i$.*

Let ρ_{post} be the normalized state after conditioning on observing the sequence of outcomes $P_1, \ldots, P_M$. That is, if $V = P_M \ldots P_1$, then $\rho_{post} = \frac{V \rho V^\dagger}{\mathrm{Tr}(V \rho V^\dagger)}$. Then the fidelity between the original state and the final post-measurement state is bounded by:

$$F(\rho, \rho_{post}) \geq 1 - \sum_{i=1}^{M} \varepsilon_i, \tag{11}$$

where $F(\rho, \sigma) = (\mathrm{Tr}\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}})^2$ is the fidelity. The trace distance is bounded by:

$$\frac{1}{2}\|\rho - \rho_{post}\|_1 \leq \sqrt{1 - F(\rho, \rho_{post})} \leq \sqrt{\sum_{i=1}^{M} \varepsilon_i}. \tag{12}$$

Proof. The proof for the pure state case was developed in PSET 1, Problem 2.c. We extend this to mixed states using purifications. Let $|\psi_\rho\rangle$ be a purification of ρ on AE. Let $\tilde{P}_i = P_i \otimes \mathbb{I}_E$. Let $|\psi_{\mathrm{post}}\rangle = (\tilde{P}_M \ldots \tilde{P}_1) |\psi_\rho\rangle / \| \ldots \|$. From the pure state result in PSET 1, we have $F(|\psi_\rho\rangle\langle\psi_\rho|, |\psi_{\mathrm{post}}\rangle\langle\psi_{\mathrm{post}}|) \geq 1 - \sum \varepsilon_i$. By the properties of fidelity under partial trace:

$$F(\rho, \rho_{\mathrm{post}}) = F(\mathrm{Tr}_E(|\psi_\rho\rangle\langle\psi_\rho|), \mathrm{Tr}_E(|\psi_{\mathrm{post}}\rangle\langle\psi_{\mathrm{post}}|)) \geq F(|\psi_\rho\rangle\langle\psi_\rho|, |\psi_{\mathrm{post}}\rangle\langle\psi_{\mathrm{post}}|) \geq 1 - \sum_{i=1}^{M} \varepsilon_i. \tag{13}$$

The trace distance bound then follows directly from the Fuchs-van de Graaf inequalities. □

3 The Quantum Union Bound

The Gentle Measurement Lemma tells us about the state disturbance when conditioned on success. The Quantum Union Bound complements this by giving a lower bound on the probability that a specific sequence of *desired* outcomes occurs when performing sequential projective measurements.

Theorem 1 (Quantum Union Bound)**.** *Let ρ be an unknown state. Let $P_1, \ldots, P_M$ be projectors representing desired ("yes") outcomes. We sequentially perform the measurements $\mathcal{E}_i = \{P_i, \mathbb{I} - P_i\}$ for $i = 1, \ldots, M$. Let $\varepsilon_i = \mathrm{Tr}((\mathbb{I} - P_i)\rho)$ be the probability of the undesired ("no") outcome $\mathbb{I} - P_i$ on the original quantum state ρ. The probability of observing the "yes" outcome P_i for all M sequential measurements is bounded below by the following:*

$$\Pr[\textit{Observe } P_1, \ldots, P_M \textit{ sequentially}] \geq \left(\frac{1 - \sum_{i=1}^{M} \varepsilon_i}{1 + \sum_{i=1}^{M} \varepsilon_i}\right)^2 \geq 1 - 4\sum_{i=1}^{M} \varepsilon_i. \tag{14}$$

The second inequality uses $\left(\frac{1-x}{1+x}\right)^2 \geq 1 - 4x$ for $x \geq 0$.

3.1 A Succinct Proof of Quantum Union Bound

Let $|\psi\rangle$ be a purification of $\rho = \mathrm{Tr}_E(|\psi\rangle\langle\psi|)$ on a joint system AE. Lift the projectors to $\tilde{P}_t = P_t \otimes \mathbb{I}_E$. The probability of the "no" outcome for $\tilde{P}_t$ on the pure state $|\psi\rangle$ is $\langle\psi|\,(\mathbb{I}-\tilde{P}_t)\,|\psi\rangle = \|(\mathbb{I}-\tilde{P}_t)\,|\psi\rangle\,\|^2 = \varepsilon_t$. Let $|\phi_0\rangle = |\psi\rangle$. Define the sequence of (unnormalized) states corresponding to obtaining the "yes" outcome $\tilde{P}_t$ at each step t:

$$|\phi_t\rangle = \tilde{P}_t\,|\phi_{t-1}\rangle = (\tilde{P}_t \dots \tilde{P}_1)\,|\psi\rangle\,. \tag{15}$$

The probability of observing all M "yes" outcomes sequentially is $p_{\text{all}} = \langle\phi_M|\phi_M\rangle = \|\phi_M\|^2$.

We start by relating $1 - \|\phi_M\|$ to the overlaps using a telescoping sum. Since $\|\psi\| = 1$ and $\|\phi_M\| \geq |\,\langle\psi|\phi_M\rangle\,|$, we have $1 - \|\phi_M\| \leq 1 - |\,\langle\psi|\phi_M\rangle\,|$. Using $\langle\psi|\phi_0\rangle = \langle\psi|\psi\rangle = 1$, we can write:

$$1 - |\,\langle\psi|\phi_M\rangle\,| = |\,\langle\psi|\phi_0\rangle\,| - |\,\langle\psi|\phi_M\rangle\,| = \sum_{t=1}^{M} \left(|\,\langle\psi|\phi_{t-1}\rangle\,| - |\,\langle\psi|\phi_t\rangle\,|\right). \tag{16}$$

Combining these gives:

$$1 - \|\phi_M\| \leq \sum_{t=1}^{M} \left(|\,\langle\psi|\phi_{t-1}\rangle\,| - |\,\langle\psi|\phi_t\rangle\,|\right). \tag{17}$$

Now, consider the term inside the sum. For any projector Π and subnormalized states $|\tilde{\psi}\rangle, |\tilde{\phi}\rangle$, the triangle inequality gives $|\,\langle\tilde{\psi}|\tilde{\phi}\rangle\,| \leq |\,\langle\tilde{\psi}|\Pi|\tilde{\phi}\rangle\,| + |\,\langle\tilde{\psi}|(\mathbb{I}-\Pi)|\tilde{\phi}\rangle\,|$. Applying the Cauchy-Schwarz inequality to the second term yields:

$$|\,\langle\tilde{\psi}|\tilde{\phi}\rangle\,| - |\,\langle\tilde{\psi}|\Pi|\tilde{\phi}\rangle\,| \leq |\,\langle\tilde{\psi}|(\mathbb{I}-\Pi)|\tilde{\phi}\rangle\,| \leq \|(\mathbb{I}-\Pi)\,|\tilde{\psi}\rangle\,\|\|(\mathbb{I}-\Pi)\,|\tilde{\phi}\rangle\,\|. \tag{18}$$

Apply this inequality (18) to each term in the sum (17) with $|\tilde{\psi}\rangle = |\psi\rangle$, $|\tilde{\phi}\rangle = |\phi_{t-1}\rangle$, and $\Pi = \tilde{P}_t$, noting $|\phi_t\rangle = \tilde{P}_t\,|\phi_{t-1}\rangle$:

$$|\,\langle\psi|\phi_{t-1}\rangle\,| - |\,\langle\psi|\phi_t\rangle\,| = |\,\langle\psi|\phi_{t-1}\rangle\,| - |\,\langle\psi|\tilde{P}_t|\phi_{t-1}\rangle\,| \leq \|(\mathbb{I}-\tilde{P}_t)\,|\psi\rangle\,\|\|(\mathbb{I}-\tilde{P}_t)\,|\phi_{t-1}\rangle\,\|. \tag{19}$$

Substituting this back into Equation (17):

$$1 - \|\phi_M\| \leq \sum_{t=1}^{M} \|(\mathbb{I}-\tilde{P}_t)\,|\psi\rangle\,\|\|(\mathbb{I}-\tilde{P}_t)\,|\phi_{t-1}\rangle\,\|. \tag{20}$$

Apply the Cauchy-Schwarz inequality to the sum on the right-hand side, $(\sum a_t b_t)^2 \leq (\sum a_t^2)(\sum b_t^2)$:

$$\left(\sum_{t=1}^{M} \|(\mathbb{I}-\tilde{P}_t)\,|\psi\rangle\,\|\|(\mathbb{I}-\tilde{P}_t)\,|\phi_{t-1}\rangle\,\|\right)^2 \leq \left(\sum_{t=1}^{M} \|(\mathbb{I}-\tilde{P}_t)\,|\psi\rangle\,\|^2\right)\left(\sum_{t=1}^{M} \|(\mathbb{I}-\tilde{P}_t)\,|\phi_{t-1}\rangle\,\|^2\right). \tag{21}$$

Let $S = \sum_{t=1}^{M} \varepsilon_t = \sum_{t=1}^{M} \|(\mathbb{I}-\tilde{P}_t)\,|\psi\rangle\,\|^2$. Recall the identity $\sum_{t=1}^{M} \|(\mathbb{I}-\tilde{P}_t)\,|\phi_{t-1}\rangle\,\|^2 = \|\phi_0\|^2 - \|\phi_M\|^2 = 1 - p_{\text{all}}$. Substituting these into (21):

$$\left(\sum_{t=1}^{M} \|(\mathbb{I}-\tilde{P}_t)\,|\psi\rangle\,\|\|(\mathbb{I}-\tilde{P}_t)\,|\phi_{t-1}\rangle\,\|\right)^2 \leq S(1-p_{\text{all}}). \tag{22}$$

Taking the square root and combining with (20):

$$1 - \|\phi_M\| \leq \sum_{t=1}^{M} \|(\mathbb{I}-\tilde{P}_t)\,|\psi\rangle\,\|\|(\mathbb{I}-\tilde{P}_t)\,|\phi_{t-1}\rangle\,\| \leq \sqrt{S(1-p_{\text{all}})}. \tag{23}$$

Substitute $\|\phi_M\| = \sqrt{p_{\text{all}}}$:

$$1 - \sqrt{p_{\text{all}}} \leq \sqrt{S(1 - p_{\text{all}})}. \tag{24}$$

Squaring both sides (as $1 - \sqrt{p_{\text{all}}} \geq 0$):

$$(1 - \sqrt{p_{\text{all}}})^2 \leq S(1 - p_{\text{all}}) = S(1 - \sqrt{p_{\text{all}}})(1 + \sqrt{p_{\text{all}}}). \tag{25}$$

Assuming $p_{\text{all}} < 1$, divide by the positive term $(1 - \sqrt{p_{\text{all}}})$, yielding:

$$1 - \sqrt{p_{\text{all}}} \leq S(1 + \sqrt{p_{\text{all}}}) = S + S\sqrt{p_{\text{all}}}. \tag{26}$$

Rearranging gives:

$$1 - S \leq \sqrt{p_{\text{all}}}(1 + S). \tag{27}$$

$$\sqrt{p_{\text{all}}} \geq \frac{1 - S}{1 + S} = \frac{1 - \sum \varepsilon_t}{1 + \sum \varepsilon_t}. \tag{28}$$

Squaring both sides results in the key result:

$$p_{\text{all}} \geq \left(\frac{1 - \sum \varepsilon_t}{1 + \sum \varepsilon_t} \right)^2. \tag{29}$$

This completes the proof.

3.2 Application: Prediction under a Discretized Promise

Section 1 showed that under the extremely strong promise $\text{Tr}(O_i\rho) \in \{-1, 0, 1\}$, prediction requires only $N = \mathcal{O}(\log(M/\delta))$ copies. Now, consider a weaker, discretized promise.

Q: Given M observables $O_1, \ldots, O_M$ (eigenvalues ± 1) and access to copies of ρ. Suppose that for each i, $\text{Tr}(O_i\rho) \in \mathcal{G}_\varepsilon = \{-1, -1 + \varepsilon, -1 + 2\varepsilon, \ldots, 1 - \varepsilon, 1\}$. How many samples N do we need to predict the correct value $v_i = \text{Tr}(O_i\rho)$ from $\mathcal{G}_\varepsilon$ for all i with a failure probability at most δ?

A: After using the same approach as in Section 1 to create the state:

$$\sigma' = U_{O_i}^{\otimes N}(\sigma \otimes |0\rangle\langle 0|^{\otimes N})U_{O_i}^{\dagger \otimes N}, \tag{30}$$

we can perform an $|\mathcal{G}_\varepsilon|$-outcome POVM on the N-qubit ancilla to produce an output in $\mathcal{G}_\varepsilon$ by counting how many 1's are present. Because of the promise, as long as $N = \mathcal{O}(\log(M/\delta)/\varepsilon^2)$, we can guarantee that the correct answer will be produced with probability at least $1 - \delta/(4M)$. Hence, by the quantum union bound, we can guarantee that we obtain the correct answer for all M observables using only logarithmic number of samples of ρ.

What happens if the promise $\text{Tr}(O_i\rho) \in \mathcal{G}_\varepsilon$ is removed? The algorithm will very likely fail. This is because we can no longer obtain the correct output with high probability. The probability distribution over the outcomes could span across a few outcomes. When measured, the quantum data will be severely disturbed because of measurement collapse.

To build towards a prediction mechanism that does not destroy our quantum data, we need to develop two powerful tools: (1) **Quantum Threshold Search** and (2) **Online Learning of Quantum States**. They will be presented in the following two sections.

4 Quantum Threshold Search

Quantum threshold search addresses the following question:

Q: Given M projectors $P_1, \ldots, P_M$ and N copies of ρ. Suppose there exists j such that $\mathrm{Tr}(P_j\rho) \geq 2/3$. Find *an* index k such that $\mathrm{Tr}(P_k\rho) \geq 1/3$. How many samples N do we need?

We do not need to find the "best" j (the one maximizing $\mathrm{Tr}(P_j\rho)$), just a "good enough" k (one above the lower threshold $1/3$). The challenge is to do this using a sample complexity N that scales polylogarithmically in M. The answer, remarkably, is $N = \mathcal{O}(\log^2 M)$. Notice that the algorithm itself works for any projectors $P_1, \ldots, P_M$; the promise is used in the analysis to guarantee finding a suitable k. This problem is studied and addressed in the paper *Improved Quantum Data Analysis* by Bădescu and O'Donnell.

4.1 Algorithm Idea: Coherent Counting and Stable Reporting

The algorithm uses the following two subroutines:

(1) Coherent Counting. Instead of measuring N copies individually to estimate $\mathrm{Tr}(P_j\rho)$, we use the N-copy state $\sigma = \rho^{\otimes N}$ held in quantum memory. We apply a counting unitary for each projector P_j. This unitary interacts the N-copy data state with an ancilla register (e.g., $\lceil \log(N+1) \rceil$ qubits) initially set to $|0\rangle$. The unitary estimates the number of successes, k, among the N virtual copies based on P_j. It maps the state as follows:

$$\sigma \otimes |0\rangle\langle 0|_{\mathrm{anc}} \xrightarrow{U_{\mathrm{count}}(P_j)} \sum_{k,k'=0}^{N} C_{k,k'} \sigma_{k,k'} \otimes |k\rangle\langle k'|_{\mathrm{anc}}, \tag{31}$$

where $\sigma_{k,k'} = P_j^{(k)} \sigma (P_j^{(k')})^\dagger$, $P_j^{(k)}$ is the projector onto the subspace where exactly k copies satisfy projector P_j, and $C_{k,k'}$ are coherence terms. The ancilla register is now entangled with the data state, holding a superposition representing the distribution of possible counts k. Measuring the ancilla in the computational basis would yield a count k with probability $p_k = \mathrm{Tr}(\sigma_{k,k})$, which follows $\mathrm{Binom}(N, p_j)$ where $p_j = \mathrm{Tr}(P_j\rho)$, but this measurement would collapse and potentially disturb the data state σ.

This coherent counting unitary can be implemented efficiently on a fault-tolerant quantum computer by using the coherent N-copy measurements (which form a unitary, similar to those in Section 1 and Section 3.2), followed by an arithmetic operation to add up all the outcomes in the N-qubit ancilla registers into the counter $|k\rangle_{\mathrm{anc}}$, and concluded by running the inverse (dagger) of the coherent N-copy measurement unitary.

(2) χ^2-Stable Reporting. The naive approach for reporting whether the counter k exceeds a threshold (like $N/2$) is to simply apply a two-outcome projective measurement on the counter $|k\rangle_{\mathrm{anc}}$. However, this could destroy our quantum data σ. The central innovation is to measure the ancilla register *gently* using a procedure inspired by techniques from adaptive data analysis and differential privacy. Instead of learning the exact count k, we perform a noisy threshold test: Is k significantly large (e.g., $k \gtrsim N/2$)? This test is designed such that if the answer is no, the quantum state σ is minimally disturbed. This property is mathematically captured by χ^2-stability, meaning the post-measurement state is close to the pre-measurement state under a χ^2-divergence measure.

4.2 The χ^2-Stable Reporting Mechanism

This is based on a classical statistical procedure, detailed in Section 3 of the paper *Improved Quantum Data Analysis*, which can be implemented as a quantum measurement (POVM) on the counting ancilla register.

Classical Mechanism. Suppose we have a classical count $K \sim \text{Binom}(N, p)$, where $p = \text{Tr}(P\rho)$. We want to test whether $p \geq 1/2$. We add random noise by sampling $X \sim \text{Exponential}(\lambda)$, where the mean is $\mathbb{E}[X] = 1/\lambda$. We then check whether the noisy count exceeds a threshold: define the event $B \equiv \{K + X \geq N/2\}$. The output is binary: did event B occur?

Quantum Implementation. This classical mechanism can be implemented on a quantum computer as a two-outcome POVM $\{M_B, M_{\bar{B}}\}$ acting only on the ancilla register (which is in an entangled state involving $|k\rangle\langle k'|$ after coherent counting). M_B corresponds to the noisy count being above the threshold (we refer to this as outcome "1"), and $M_{\bar{B}} = \mathbb{I}_{\text{anc}} - M_B$ corresponds to it being below (we refer to this as outcome "0").

Stability (Classical). A key classical result (Theorem 1.2 in the paper) states that if the probability $\mathbb{P}[B]$ of the noisy count exceeding the threshold is small ($\mathbb{P}[B] < 1/4$), then conditioning on the opposite event $\bar{B}$ (outcome "0") barely changes the distribution of the original count K. This closeness is measured by the χ^2-divergence:

$$d_{\chi^2}((K|\bar{B}), K) \leq \mathbb{P}[B]^2 \left(\frac{N}{(\mathbb{E}[X])^2} \right). \tag{32}$$

Hence, the disturbance is small if the mean of the exponentially-decaying noise $\mathbb{E}[X]$ is sufficiently large, specifically $\mathbb{E}[X] \gtrsim \sqrt{N}$.

Gentleness (Quantum). Because the ancilla is entangled with the data state σ, measuring the ancilla affects σ. Corollary 3.5 in the paper uses the classical χ^2-stability (32) to bound the disturbance to the data state σ when the measurement outcome is $\bar{B}$ (outcome "0"). The post-measurement data state $\sigma|_{\bar{B}}$ remains close to the pre-measurement state σ:

$$\|\sigma - \sigma|_{\bar{B}}\|_1 \leq 2 d_{\text{Bures}}(\sigma, \sigma|_{\bar{B}}) \leq 2\sqrt{d_{\chi^2}((K|\bar{B}), K)} \leq 2 \cdot \mathbb{P}[B] \cdot \frac{\sqrt{N}}{\mathbb{E}[X]}, \tag{33}$$

using $d_{\text{Bures}} \leq \sqrt{d_{\chi^2}}$ and the Fuchs-van de Graaf inequalities. If we choose the noise mean $\mathbb{E}[X] = c\sqrt{N}$ for some constant $c \geq 1$, the trace distance disturbance is $\mathcal{O}(\mathbb{P}[B]/c)$. Hence, the disturbance scales linearly with the probability $\mathbb{P}[B]$ that the noisy threshold is exceeded, with the proportionality constant depending on the noise level.

4.3 The Full Algorithm and its Analysis

This leads to the iterative algorithm presented below.

Algorithm: Quantum Threshold Search

1. **Initialization:** Take $N = \Theta(\log^2 M)$ copies of ρ into quantum memory $\sigma_0 = \rho^{\otimes N}$. Set the mean exponential noise $\mathbb{E}[X] = \Theta(\log M)$ (e.g., $\mathbb{E}[X] = c\sqrt{N}$ with $N = \Theta(\log^2 M)$ implies $\mathbb{E}[X] = \Theta(\log M)$). Let $\sigma_{\text{current}} = \sigma_0$.

2. **Iteration:** For $j = 1$ to M:

   (a) Apply the coherent counting unitary $U_{\text{count}}(P_j)$ on σ_{current} and a fresh ancilla register initialized to $|0\rangle$.

   (b) Perform the χ^2-stable reporting POVM $\{M_{B_j}, M_{\bar{B}_j}\}$ on the ancilla register, corresponding classically to the event $B_j \equiv \{K + X \geq N/2\}$ where $K \sim \text{Binom}(N, \text{Tr}(P_j\rho))$.

   (c) **If outcome is B_j (report "1"):** Halt and output the index j.

   (d) **If outcome is $\bar{B}_j$ (report "0"):** The disturbance on the state caused by the measurement is proportional to the probability $\mathbb{P}[B_j]$ of observing "1". The state σ_{current} is now updated to the post-measurement state conditioned on outcome $\bar{B}_j$. Continue to the next iteration $j + 1$.

3. **Failure:** If the loop finishes without halting, output failure.

Theorem 2 (Quantum Threshold Search)**.** *Using $N = \mathcal{O}(\log^2 M)$ samples, the quantum threshold search algorithm finds an index j such that $\text{Tr}(P_j\rho) \geq 1/3$ with constant success probability, provided there exists at least one k with $\text{Tr}(P_k\rho) \geq 2/3$.*

Proof Idea. The analysis balances the probability of stopping correctly against the accumulated disturbance to the quantum state (our precious quantum data). Let $p_j = \mathbb{P}[B_j]$ be the probability of outcome "1" for projector P_j on the *initial* quantum state σ_0.

Failure Mode (a): Never Stop. Does the algorithm run through all M projectors without ever obtaining outcome "1"? Let q_M be the probability of this event. Because the promise ensures $\exists k$ such that $\text{Tr}(P_k\rho) \geq 2/3$, we know at least one p_k is large. The gentleness property (33) ensures that the sequential probability q_M does not deviate too much from $\prod_{j=1}^{M}(1 - p_j)$. Together, one can show that q_M is sufficiently smaller than 1.

Failure Mode (b): Stop at Bad j. Does the algorithm stop at an index j where $\text{Tr}(P_j\rho) < 1/3$? Let $\mathcal{B}$ denote the set of such bad indices. For $j \in \mathcal{B}$, the binomial count K is concentrated below $N/3$. The probability $p_j = \mathbb{P}[B_j] = \mathbb{P}[K + X \geq N/2]$ is exponentially small in N due to the concentration of binomial random variables and the choice of threshold. The probability of stopping at *any* bad j is $\sum_{j\in\mathcal{B}} s_j$, where s_j is the probability of the sequence $0, 0, \ldots, 0, 1$ (stopping at j). Using the quantum union bound or related arguments, one can show that $\sum_{j\in\mathcal{B}} s_j$ is small.

Success. Since the total probability is 1, and the probability of never stopping (q_M) and the probability of stopping at a bad j ($\sum_{j\in\mathcal{B}} s_j$) are both sufficiently bounded away from 1, the probability of stopping at a good index j (where $\text{Tr}(P_j\rho) \geq 1/3$) must be bounded below by a positive constant. The key insight is that the disturbance to the state predominantly accumulates only when the algorithm correctly identifies a potential candidate (i.e., obtains outcome "1" for a j with large $\text{Tr}(P_j\rho)$), but then it halts. The disturbance accumulated during the "0" outcomes for bad j's is minimal due to the χ^2-stability (33) and the exponentially small values of p_j for bad j.

This adaptive procedure, enabled by coherent processing and gentle measurements, allows the quantum data to be reused M times with only polylogarithmic sample cost $N = \mathcal{O}(\log^2 M)$. □

5 Online Machine Learning in the Quantum World

In online machine learning, we make sequential predictions, receive feedback in the form of losses, and update our model in real time. We begin by developing the classical theory of prediction with expert advice and demonstrate how it generalizes elegantly to the quantum domain.

5.1 The Prediction with Experts Problem

Consider a canonical problem in online learning. Suppose you wish to make predictions with access to a panel of d experts. At each time step $t \in \{1, \ldots, T\}$, the following sequence occurs:

1. **You predict:** You choose a probability distribution $q^{(t)}$ over the d experts, representing your trust in each expert. Your prediction is a weighted average of their individual predictions.
2. **Experts predict:** The experts' predictions $s_1^{(t)}, \ldots, s_d^{(t)} \in [0, 1]$ are revealed.
3. **Truth is revealed:** The true answer $a^{(t)}$ is revealed. You incur a loss defined as the absolute difference between your blended prediction and the true answer:

$$l^{(t)}(q^{(t)}) = \left| \sum_{x=1}^{d} q^{(t)}(x) s_x^{(t)} - a^{(t)} \right|. \tag{34}$$

Our objective is to design an algorithm for selecting the probability distribution $q^{(t)}$ at each time step that minimizes the total cumulative loss $\sum_{t=1}^{T} l^{(t)}(q^{(t)})$.

5.1.1 Regret: A Measure of Performance

To evaluate an algorithm's performance, we compare it against the best possible strategy in hindsight. Specifically, we compare our performance to the single best fixed distribution p^* that minimizes the total loss over all T steps:

$$p^* = \arg\min_{p} \sum_{t=1}^{T} l^{(t)}(p). \tag{35}$$

Since we cannot use p^* from the beginning without knowledge of future losses, we instead aim to minimize the **regret** R_T, defined as the difference between our algorithm's total loss and the best-in-hindsight fixed strategy's loss:

$$R_T = \sum_{t=1}^{T} l^{(t)}(q^{(t)}) - \sum_{t=1}^{T} l^{(t)}(p^*). \tag{36}$$

A successful algorithm achieves *sublinear regret*, meaning R_T grows slower than T (e.g., as $\sqrt{T}$). This ensures that the average regret per step, R_T/T, vanishes as $T \to \infty$.

5.1.2 A Naive Algorithm: Follow the Leader

A natural initial strategy is Follow the Leader (FTL). At each step t, FTL selects the expert or distribution that has incurred the minimum cumulative loss up to time $t - 1$.

Question: How can we construct an adversarial scenario where FTL fails catastrophically?

Answer: Consider a simple case with $d = 2$ experts where the world is adversarial and the experts alternate in their failures. Suppose:

- Expert 1 has losses: $(1, 0, 0, 1, 1, 0, 0, \ldots)$
- Expert 2 has losses: $(0, 1, 1, 0, 0, 1, 1, \ldots)$

The FTL algorithm selects selects Expert 1 and 2 uniformly for $t = 1$. The FTL algorithm observes that Expert 2 performed better at $t = 1$ (loss 0) and thus selects Expert 2 for $t = 2$. However, Expert 2 incurs loss 1 at $t = 2$. The algorithm then selects Expert 1 and 2 uniformly for $t = 3$. The algorithm then switches to Expert 1 for $t = 4$, selects uniform for $t = 5$, switches to Expert 2 for $t = 6$, selects uniform for $t = 7$, and this pattern continues. The total loss is $\frac{3}{4}T$. Meanwhile, the best fixed strategy in hindsight is $p^* = (1/2, 1/2)$, which achieves loss $1/2$ at every step, for a total loss of $T/2$. The regret for FTL is thus $R_T = 3T/4 - T/2 = T/4$, which is linear in T. This linear regret is unacceptable, indicating the need for a more stable algorithm.

5.2 Follow the Regularized Leader

The fundamental problem with FTL is excessive reactivity: it switches completely to the new best strategy. We address this by adding a regularization term that penalizes large deviations from the previous distribution. This approach yields the **Follow the Regularized Leader (FTRL)** algorithm. *Notably, for many years, this algorithm formed the core of Google Ads' machine learning infrastructure, demonstrating its immense practical importance.*

Algorithm (FTRL):

- Initialize with the uniform distribution: $q^{(1)} = (1/d, \ldots, 1/d)$.
- For each time step t, compute the next distribution $q^{(t+1)}$ by solving:

$$q^{(t+1)} = \arg\min_q \left[\hat{l}^{(t)}(q) + \lambda R(q \| q^{(t)}) \right], \tag{37}$$

  where:

  - $\hat{l}^{(t)}(q)$ is a linearized proxy for the loss at step t. Instead of using the full, potentially nonlinear loss function, we use a first-order approximation based on the (sub-)gradient $\nabla^{(t)} = \nabla l^{(t)}(q^{(t)})$ at the current point:

$$\hat{l}^{(t)}(q) = l^{(t)}(q^{(t)}) + \nabla^{(t)} \cdot (q - q^{(t)}). \tag{38}$$

    Since $q^{(t)}$, $l^{(t)}(q^{(t)})$, and $\nabla^{(t)} \cdot q^{(t)}$ are constant, minimizing this is equivalent to minimizing $\langle q, \nabla^{(t)} \rangle$. For our specific loss $l^{(t)}(q) = |\langle q, s^{(t)} \rangle - a^{(t)}|$, the gradient is

$$\nabla^{(t)} = \text{sign}(\langle q^{(t)}, s^{(t)} \rangle - a^{(t)}) \cdot s^{(t)}. \tag{39}$$

  - $R(q \| q^{(t)})$ is the **regularization term**, measuring the divergence from the new distribution q to the old distribution $q^{(t)}$. An optimal choice is the **Kullback-Leibler (KL) divergence**:

$$R(q \| q^{(t)}) = D_{\text{KL}}(q \| q^{(t)}) = \sum_x q(x) \log \left[\frac{q(x)}{q^{(t)}(x)} \right]. \tag{40}$$

  - $\lambda > 0$ is the regularization parameter (or inverse learning rate), controlling the trade-off between stability (staying close to $q^{(t)}$) and agility (minimizing current loss).

5.3 Multiplicative Weight Update

When we use the KL divergence as the regularizer in FTRL, the minimization problem admits a closed-form solution known as the **Multiplicative Weight Update (MWU)** rule:

$$q^{(t+1)}(x) = \frac{q^{(t)}(x) \exp\left(-\frac{1}{\lambda}\nabla_x^{(t)}\right)}{\sum_{x'} q^{(t)}(x') \exp\left(-\frac{1}{\lambda}\nabla_{x'}^{(t)}\right)}. \tag{41}$$

This algorithm has an intuitive interpretation. The weight of each expert x is multiplied by a factor $\exp(-\nabla_x^{(t)}/\lambda)$. If an expert's gradient component $\nabla_x^{(t)}$ is large and positive (indicating they contributed to a large loss), their weight is exponentially suppressed. Conversely, if their gradient is negative, their weight is exponentially boosted. The denominator normalizes to ensure $q^{(t+1)}$ remains a valid probability distribution.

The underlying principle is compelling: rather than abruptly changing beliefs, we gradually and multiplicatively accumulate new information.

Theorem 3 (FTRL/MWU Regret Bound)**.** *The regret of the FTRL/MWU algorithm is bounded. Assuming the gradient components satisfy $|\nabla_x^{(t)}| \le 1$, we have*

$$R_T \le \frac{T}{\lambda} + \lambda \log(d). \tag{42}$$

By choosing the optimal regularization parameter $\lambda = \sqrt{T/\log d}$ to balance these two contributions, we achieve a sublinear regret:

$$R_T \le 2\sqrt{T \log d}. \tag{43}$$

Proof. The proof employs a potential function analysis, a standard technique in online learning theory. Define the unnormalized potential function Φ_t as

$$\Phi_t := \sum_{x=1}^{d} \exp\left(-\frac{1}{\lambda}\sum_{\tau=1}^{t-1}\nabla_x^{(\tau)}\right) = \sum_{x=1}^{d} w_t(x), \tag{44}$$

where $w_t(x)$ denotes the unnormalized weight of expert x at time t. Note that the algorithm's distribution is $q^{(t)}(x) = w_t(x)/\Phi_t$.

Consider how Φ evolves from t to $t+1$:

$$\Phi_{t+1} = \sum_{x=1}^{d} w_t(x) \exp\left(-\frac{1}{\lambda}\nabla_x^{(t)}\right). \tag{45}$$

Using the inequality $e^{-y} \le 1 - y + y^2$ for $|y| \le 1$ (which follows from Taylor expansion), we obtain

$$\begin{aligned}
\Phi_{t+1} &\le \sum_x w_t(x)\left(1 - \frac{1}{\lambda}\nabla_x^{(t)} + \frac{1}{\lambda^2}(\nabla_x^{(t)})^2\right) && (46)\\
&\le \left(1 + \frac{1}{\lambda^2}\right)\sum_x w_t(x) - \frac{1}{\lambda}\sum_x w_t(x)\nabla_x^{(t)} \quad (\text{using } |\nabla_x^{(t)}| \le 1) && (47)\\
&= \left(1 + \frac{1}{\lambda^2}\right)\Phi_t - \frac{1}{\lambda}\Phi_t \sum_x q^{(t)}(x)\nabla_x^{(t)} && (48)\\
&\le \Phi_t \exp\left(\frac{1}{\lambda^2} - \frac{1}{\lambda}\langle q^{(t)}, \nabla^{(t)}\rangle\right) \quad (\text{using } 1 + y \le e^y). && (49)
\end{aligned}$$

Applying this recursively from $t = 1$ to T, and noting $\Phi_1 = d$, we obtain

$$\Phi_{T+1} \leq d \cdot \exp\left(\frac{T}{\lambda^2} - \frac{1}{\lambda}\sum_{t=1}^{T}\langle q^{(t)}, \nabla^{(t)}\rangle\right). \tag{50}$$

For a lower bound on Φ_{T+1}, observe that for any fixed distribution p,

$$\Phi_{T+1} = \sum_x \exp\left(-\frac{1}{\lambda}\sum_{t=1}^{T}\nabla_x^{(t)}\right) \geq \exp\left(-\frac{1}{\lambda}\sum_{t=1}^{T}\langle p^*, \nabla^{(t)}\rangle\right), \tag{51}$$

where p^* is the best fixed distribution in hindsight.

Combining the upper and lower bounds and taking logarithms yields

$$-\frac{1}{\lambda}\sum_{t=1}^{T}\langle p^*, \nabla^{(t)}\rangle \leq \log d + \frac{T}{\lambda^2} - \frac{1}{\lambda}\sum_{t=1}^{T}\langle q^{(t)}, \nabla^{(t)}\rangle. \tag{52}$$

Rearranging and multiplying by λ gives

$$\sum_{t=1}^{T}\langle \nabla^{(t)}, q^{(t)} - p^*\rangle \leq \lambda \log d + \frac{T}{\lambda}. \tag{53}$$

Finally, using the convexity of the loss function, we have for any convex function that $l(p^*) \geq l(q^{(t)}) + \nabla l(q^{(t)}) \cdot (p^* - q^{(t)})$. Summing over t yields

$$R_T = \sum_{t=1}^{T}(l^{(t)}(q^{(t)}) - l^{(t)}(p^*)) \leq \sum_{t=1}^{T}\langle \nabla^{(t)}, q^{(t)} - p^*\rangle. \tag{54}$$

Thus, we obtain the regret bound $R_T \leq \lambda \log d + T/\lambda$, completing the proof. □

5.4 Quantum Online Learning

We now extend these ideas to the quantum domain. What is the quantum analogue of a probability distribution q over d experts? It is a density matrix σ on a d-dimensional Hilbert space.

Question: Consider an n-qubit system with dimension $d = 2^n$. At each time step $t \in \{1, \ldots, T\}$:

1. **You propose** a hypothesis state $\sigma^{(t)}$, an n-qubit density matrix.
2. **World reveals** an observable $O^{(t)}$ (a Hermitian operator with eigenvalues in $[0, 1]$) and the true expectation value $a^{(t)} = \mathrm{Tr}(O^{(t)}\rho)$, where ρ is the true unknown state.
3. **You incur loss** based on your prediction $\mathrm{Tr}(O^{(t)}\sigma^{(t)})$:

$$l^{(t)}(\sigma^{(t)}) = |\mathrm{Tr}(O^{(t)}\sigma^{(t)}) - a^{(t)}|. \tag{55}$$

What is an achievable regret $R_T = \sum_{t=1}^{T} l^{(t)}(\sigma^{(t)}) - \min_\rho \sum_{t=1}^{T} l^{(t)}(\rho)$?

Answer: The framework generalizes perfectly. The regret bound has the same form, with the dimension replaced appropriately:

$$R_T \leq \mathcal{O}(\sqrt{T \log d}) = \mathcal{O}(\sqrt{Tn}). \tag{56}$$

This result is remarkable: the regret scales only with the square root of the number of qubits, not the exponential dimension of the Hilbert space.

Algorithm (Quantum FTRL / Matrix Multiplicative Weight):

The algorithm is a direct matrix generalization of the classical case:

- Initialize with the maximally mixed state: $\sigma^{(1)} = I/2^n$.
- For $t = 1, \ldots, T$:
  - The gradient becomes a gradient operator $V^{(t)}$. For our loss function, this is

$$V^{(t)} = \text{sign}(\text{Tr}(O^{(t)}\sigma^{(t)}) - a^{(t)}) \cdot O^{(t)}. \tag{57}$$

  - The update rule is the **Matrix Multiplicative Weight** update, derived from FTRL using the **quantum relative entropy** $S(\sigma\|\rho) = \text{Tr}(\sigma(\log\sigma - \log\rho))$ as the regularizer:

$$\sigma^{(t+1)} = \frac{\exp\left(-\frac{1}{\lambda}\sum_{\tau=1}^{t} V^{(\tau)}\right)}{\text{Tr}\left(\exp\left(-\frac{1}{\lambda}\sum_{\tau=1}^{t} V^{(\tau)}\right)\right)}. \tag{58}$$

Observe the structure of this state. Defining the effective Hamiltonian $H_t = \sum_{\tau=1}^{t} V^{(\tau)}$, our hypothesis state becomes

$$\sigma^{(t+1)} = \frac{e^{-H_t/\lambda}}{\text{Tr}(e^{-H_t/\lambda})}. \tag{59}$$

This is precisely a **Gibbs state** for the Hamiltonian H_t with effective temperature $T_{\text{eff}} \propto \lambda$. *This connection implies that if we possess an efficient quantum algorithm for preparing Gibbs states, we automatically obtain an efficient algorithm for online learning of quantum states.*

Theorem 4. *Using this online learning algorithm, we can learn n-qubit observables while making only $\mathcal{O}(n/\varepsilon^2)$ large errors.*

Proof sketch. This result follows from the sublinear regret property. Consider running the online algorithm and counting steps where the loss is large: $l^{(t)}(\sigma^{(t)}) > \varepsilon$. The total loss accumulated from these large error steps is at least $\#\{\text{large errors}\} \cdot \varepsilon$. This sum must also be bounded by the total regret R_T (plus the best-in-hindsight loss, which is at most 0 if ρ is the true state):

$$\#\{\text{large errors}\} \cdot \varepsilon \leq \sum_{t \in \text{large errors}} l^{(t)}(\sigma^{(t)}) \leq R_T \leq \mathcal{O}(\sqrt{Tn}). \tag{60}$$

This yields

$$\#\{\text{large errors}\} \leq \mathcal{O}\left(\frac{\sqrt{Tn}}{\varepsilon}\right) = \mathcal{O}\left(\frac{n}{\varepsilon^2}\right), \tag{61}$$

where we have set $T = \mathcal{O}(n/\varepsilon^2)$ to optimize the bound. This demonstrates that our algorithm learns and converges: the number of large errors is finite and scales only linearly in n. □

6 Full-Fledged Shadow Tomography

Suppose we can collect only a finite number of samples of an unknown state ρ. This constraint is extremely practical for two reasons:

1. **Sample scarcity:** Collecting quantum data samples from real-world physical systems (e.g., complex molecules) may be exceedingly difficult.
2. **Sample cost:** Producing high-fidelity copies of a ground state on a fault-tolerant quantum computer is a slow, challenging process requiring long cooling times.

Consider a large-scale problem where we would like to conduct many experiments to extract many different properties of ρ. Each experiment might represent an entire quantum simulation by generating excitations on the ground state ρ, evolving under the Hamiltonian to produce dynamics of the excitations, and performing measurements to extract outcomes of the simulation experiment. All of these experiments can be written as some highly complex observables $O_1, \ldots, O_M$ we would like to measure on ρ. Given the sample scarcity and cost, we would like to use as few samples of ρ as possible. This leads to the following question.

Question: Given M observables $O_1, \ldots, O_M$ (with eigenvalues in $[0, 1]$) and N copies of ρ, we wish to predict $\mathrm{Tr}(O_i\rho)$ for all M observables. How many samples of ρ are required to predict all M expectation values to within error ε?

Answer: The naive approach would require $\mathcal{O}(\min(4^n, M))$ samples, as we would either to perform full quantum state tomography with sample complexity $\mathcal{O}(4^n)$ or direct measurements of the M observables. However, we can achieve a dramatically better performance of only $\mathcal{O}(n \log^2 M)$.

Recall the basic Quantum Threshold Search (QTS):

Theorem 5 (Quantum Threshold Search, Section 4)**.** *Using $N = \mathcal{O}(\frac{\log^2 M}{\varepsilon^2})$ samples of ρ, for any sequence of projectors $A_1, \ldots, A_M$, the QTS algorithm finds an index i such that $\mathrm{Tr}(A_i\rho) \geq 1/3$ (with high probability), provided one exists with $\mathrm{Tr}(A_k\rho) \geq 2/3$.*

This is powerful, but somewhat restrictive: it only applies to projectors and fixed thresholds $(1/3, 2/3)$. How can we generalize this to arbitrary observables O_i and arbitrary thresholds θ_i?

6.1 Advanced Quantum Threshold Search

This generalization yields **Advanced Quantum Threshold Search**, which we obtain as a corollary of the basic QTS by cleverly defining our projectors.

Corollary 1 (Advanced Quantum Threshold Search)**.** *Using $N_{\text{total}} = \mathcal{O}(\frac{\log^2(M/\delta)}{\varepsilon^2})$ samples of ρ, there exists an algorithm that, for any sequence of observables $O_1, \ldots, O_M$ (with eigenvalues in $[0, 1]$) and any sequence of thresholds $\theta_1, \ldots, \theta_M$, guarantees with probability $1 - \delta$:*

1. ***Detection:** If the algorithm stops at index i, then the true value is far from the threshold:*

$$|\mathrm{Tr}(O_i\rho) - \theta_i| > \varepsilon. \tag{62}$$

2. ***Consistency:** If the algorithm does not stop at index i, then the value is close to its threshold:*

$$|\mathrm{Tr}(O_i\rho) - \theta_i| \leq \varepsilon. \tag{63}$$

Proof sketch: Reduction to Basic QTS. We reduce this problem to the basic QTS by cleverly constructing projectors $A_1, \ldots, A_M$ with the required properties.

For each pair (O_i, θ_i), we design a check projector A_i such that:

1. If $|\mathrm{Tr}(O_i\rho) - \theta_i| \geq 2\varepsilon/3$ (the far case), then $\mathrm{Tr}(A_i\rho^{\otimes N_{\text{check}}}) \geq 2/3$.
2. If $|\mathrm{Tr}(O_i\rho) - \theta_i| \leq \varepsilon/3$ (the close case), then $\mathrm{Tr}(A_i\rho^{\otimes N_{\text{check}}}) \leq 1/3$.

The projector A_i is an algorithmic projector acting on $N_{\text{check}} = \mathcal{O}(1/\varepsilon^2)$ copies of ρ:

1. Take $N_{\text{check}} = \mathcal{O}(1/\varepsilon^2)$ copies of ρ.

2. Run the phase estimation algorithm on these copies to estimate $p = \mathrm{Tr}(O_i\rho)$. This number of copies suffices to obtain an estimate $\tilde{p}$ satisfying $|\tilde{p} - p| \leq \varepsilon/3$ with high probability.

3. Write this estimate $\tilde{p}$ into an ancilla register.

4. Perform a computation on this ancilla to check whether $|\tilde{p} - \theta_i| > \varepsilon/2$.

5. Define A_i as the projector onto the subspace of ancillas corresponding to the "yes" outcome.

Analysis:

- **Far case ($|p - \theta_i| \geq 2\varepsilon/3$):** Our estimate $\tilde{p}$ is close to p. By the triangle inequality,

$$|\tilde{p} - \theta_i| \geq |p - \theta_i| - |\tilde{p} - p| \geq \frac{2\varepsilon}{3} - \frac{\varepsilon}{3} = \frac{\varepsilon}{3}. \tag{64}$$

  We tune the phase estimation parameters (specifically N_{check} and the threshold $\varepsilon/2$) to ensure this "yes" outcome occurs with probability $\geq 2/3$.

- **Close case ($|p - \theta_i| \leq \varepsilon/3$):** Again, $\tilde{p}$ is close to p. By the triangle inequality,

$$|\tilde{p} - \theta_i| \leq |\tilde{p} - p| + |p - \theta_i| \leq \frac{\varepsilon}{3} + \frac{\varepsilon}{3} = \frac{2\varepsilon}{3}. \tag{65}$$

  The probability of the check firing (outputting "yes") is low and can be tuned to be $\leq 1/3$.

Thus, we have constructed projectors A_i satisfying the $1/3, 2/3$ promise. We now run the basic QTS algorithm on these algorithmic projectors $A_1, \ldots, A_M$. The total sample complexity is the cost of the outer QTS loop, $\mathcal{O}(\log^2 M)$, times the cost of each inner check, $N_{\text{check}} = \mathcal{O}(1/\varepsilon^2)$, yielding the final bound $N_{\text{total}} = \mathcal{O}(\frac{\log^2 M}{\varepsilon^2})$. □

6.2 Online Learning + Quantum Threshold Search = Shadow Tomography

We now have all the ingredients to achieve our final, remarkable goal: full-fledged shadow tomography. We will show how to combine the **Online Learning** framework (which provides a classical model $\sigma^{(t)}$ and a guaranteed bound on updates) with **Advanced Quantum Threshold Search** (which acts as a powerful, gentle "consistency checker") to solve the prediction problem.

The Goal: Given M observables $O_1, \ldots, O_M$ and access to a state ρ, we want to output a list of M classical values $v_1, \ldots, v_M$ such that $|v_i - \mathrm{Tr}(O_i\rho)| \leq \varepsilon$ for all $i = 1, \ldots, M$.

The Algorithm:

The algorithm works by iteratively improving a classical hypothesis state $\sigma^{(t)}$ until it is consistent with the true state ρ on all M observables simultaneously.

1. **Initialize:** Start with $t \leftarrow 1$, $i \leftarrow 1$, and an initial hypothesis: the maximally mixed state $\sigma^{(1)} = \mathbb{I}/2^n$. Initialize the Advanced Quantum Threshold Search (AQTS) algorithm from Corollary 1 by supplying the algorithm with $\mathcal{O}(\log^2 M/\varepsilon^2)$ samples of ρ.

2. **Check Consistency:** Use AQTS from Corollary 1 as a "consistency checker." We use our current hypothesis's predictions as the thresholds:

$$\text{Run AQTS on } (O_i, \theta_i = \mathrm{Tr}(O_i\sigma^{(t)}))$$

3. **Analyze Outcome:** There are two possible results from this check:

- **Case A: AQTS does NOT stop.** This is the success condition. By the Consistency condition (Corollary 1, point 2), this guarantees that our hypothesis is good for O_i:

$$|\mathrm{Tr}(O_i\rho) - \mathrm{Tr}(O_i\sigma^{(t)})| \leq \varepsilon.$$

- **Case B: AQTS STOPS.** This is our update trigger. By the Detection property (Corollary 1, point 1), this means

$$|\mathrm{Tr}(O_i\rho) - \mathrm{Tr}(O_i\sigma^{(t)})| > \varepsilon.$$

We have found a "large error." This is exactly the signal our Online Learning algorithm (Section 5) needs. We feed this error-witnessing observable O_i and its true value, $a_i = \mathrm{Tr}(O_i\rho)$, which we learn by consuming $\mathcal{O}(\log M/\varepsilon^2)$ samples of ρ into the Matrix Multiplicative Weight update rule. This produces a new, improved hypothesis $\sigma^{(t+1)}$.

When AQTS stops, it burns the $\mathcal{O}(\log^2 M/\varepsilon^2)$ samples of ρ. Hence, we need to reinitialize the AQTS algorithm by supplying it with $\mathcal{O}(\log^2 M/\varepsilon^2)$ samples of ρ. Update $t \leftarrow t+1$.

4. **Repeat:** Increment $i \leftarrow i+1$ and **go back to Step 2**.

6.2.1 Analysis and Total Sample Complexity

The provided algorithm is a *single-pass* procedure that iterates through the M observables one by one. Its total sample cost is determined by the number of times it must "restart" its quantum sample set, which only happens when an inconsistency is found (Case B). We can analyze the total cost by separating the two key components:

1. **The number of updates ($T_{\mathbf{stop}}$):** How many times can the algorithm enter Case B?
2. **The sample cost of each observable:** What is the cost of each observable?

1. Bound on Hypothesis Updates ($T_{\mathbf{stop}}$) The algorithm only updates its hypothesis ($\sigma^{(t)} \to \sigma^{(t+1)}$) when it enters **Case B**.

- A Case B event is triggered if and only if the Advanced Quantum Threshold Search (AQTS) stops, which means we have found a large error: $|\mathrm{Tr}(O_i\rho) - \mathrm{Tr}(O_i\sigma^{(t)})| > \varepsilon$.
- The total number of hypothesis updates, T_{stop}, is therefore equal to the total number of large errors our online learner makes over the entire run.
- The regret bound analysis for the Matrix Multiplicative Weight algorithm (Section 5, based on the bound $R_T \sim \mathcal{O}(n/\varepsilon)$) guarantees that this number is finite and bounded:

$$T_{\mathrm{stop}} = \#\{\text{large errors}\} \leq \mathcal{O}\left(\frac{n}{\varepsilon^2}\right)$$

2. Total Sample Complexity Analysis Now we can calculate the total sample cost N_{total} by summing the costs of the algorithm's entire execution (for $i = 1, \ldots, M$).

- **Initial Cost:** The algorithm begins (Step 1) by loading one "batch" of quantum data to initialize the AQTS.

$$N_{\mathrm{initial}} = \mathcal{O}\left(\frac{\log^2(M/\delta)}{\varepsilon^2}\right) \text{ samples}$$

- **Cost during the pass (Steps 2-4):** The algorithm takes M steps ($i = 1, \ldots, M$).
  - **Cost of "Pass" (Case A):** By the properties of AQTS, if the "pass" outcome occurs, the quantum state is not significantly disturbed. Therefore, we reuse the *same batch* of samples for the next check ($i + 1$). The sample cost is **0**.
  - **Cost of "Fail" (Case B):** This event happens T_{stop} times. According to the algorithm, this is an expensive failure:
    1. We consume $N_{\text{learn}} = \mathcal{O}(\log M/\varepsilon^2)$ new samples to learn the true value a_i.
    2. The AQTS check "burns" the current batch $N_{\text{AQTS}} = \mathcal{O}(\log^2 M/\varepsilon^2)$.
    3. We must re-initialize AQTS with a fresh batch, $N_{\text{re-init}} = \mathcal{O}(\log^2 M/\varepsilon^2)$.

    The total cost *per failure* is

$$N_{\text{per\_stop}} = \mathcal{O}\left(\frac{\log^2(M/\delta)}{\varepsilon^2}\right).$$

6.2.2 Final Result

We now substitute our bound for T_{stop} into the total cost equation.

$$N_{\text{total}} = \mathcal{O}\left(\frac{n}{\varepsilon^2}\right) \times \mathcal{O}\left(\frac{\log^2(M/\delta)}{\varepsilon^2}\right) = \mathcal{O}\left(\frac{n \log^2(M)}{\varepsilon^2}\right).$$

This final result is the culmination of our tools. We have designed an algorithm that can predict M arbitrary expectation values of an n-qubit state ρ with a total sample complexity that scales:

- **Linearly in** n (the number of qubits).
- **Polylogarithmically in** M (the number of observables).
- Polynomially in $1/\varepsilon$ (the error).

7 Catalytic Single-Copy Tomography

All the methods we have discussed so far assume we have *many copies* of the state ρ that are consumed during the measurement process. We now consider a fundamentally different paradigm.

Question: Suppose we wish to predict M observables $O_1, \ldots, O_M$ (with eigenvalues in $[0, 1]$) for a state $|\psi_g\rangle$ that is the gapped ground state of a known Hamiltonian H. How many copies of $|\psi_g\rangle$ do we need to predict all M observables?

Answer: Remarkably, a single copy suffices.

7.1 Single-Copy Tomography

This idea, known as single-copy tomography, relies on the fact that the single copy of $|\psi_g\rangle$ is not consumed. Instead, the ground state acts as a **catalyst**, enabling us to perform many measurements without destroying the state.

7.1.1 Algorithmic Idea

The core idea is to gently couple our single copy of $|\psi_g\rangle$ to an ancilla qubit. We wish to measure the expectation value $\langle\psi_g|O|\psi_g\rangle$ for some observable O. A simple algorithm proceeds as follows:

1. Prepare an ancilla qubit in the state $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$.

2. Evolve the joint system (the state $|\psi_g\rangle$ and the ancilla $|+\rangle$) for a short time Δt under the interaction Hamiltonian
$$H_{\text{int}} = O \otimes Z. \tag{66}$$
The state of the joint system evolves to
$$|\psi(\Delta t)\rangle = e^{-i(O\otimes Z)\Delta t}(|\psi_g\rangle \otimes |+\rangle). \tag{67}$$

3. Perform quantum phase estimation using the Hamiltonian H as a measurement on the n-qubit system. The measurement outcome will be the ground state energy of H with high probability. If the measurement outcome is the ground state energy, the post-measurement n-qubit state is restored to $|\psi_g\rangle$.

7.1.2 The Quantum Zeno Effect as Catalyst

How does this work? The key lies in the fact that $|\psi_g\rangle$ is a gapped ground state.

The interaction Hamiltonian H_{int} applies a perturbation (namely O) to the system. After a small time Δt, the quantum phase estimation measurement projects the system back onto the ground state subspace. This is a manifestation of the quantum Zeno effect: frequent measurements suppress transitions out of the measured subspace.

Crucially, although the system state returns to $|\psi_g\rangle$, the effect of the evolution does not vanish. Instead, the phase information from the interaction is imprinted onto the ancilla qubit. The total evolution can be approximated (to first order in Δt) as
$$e^{-i(O\otimes Z)\Delta t}(|\psi_g\rangle \otimes |+\rangle) \approx |\psi_g\rangle \otimes \left(e^{-i\langle\psi_g|O|\psi_g\rangle Z\Delta t}\,|+\rangle\right) + \mathcal{O}(\Delta t^2). \tag{68}$$

7.1.3 Extracting the Expectation Value

Examining the resulting state after the phase estimation measurement:

- The system state $|\psi_g\rangle$ is unchanged, acting as a catalyst.
- The ancilla qubit has undergone a rotation
$$|\psi_{\text{anc}}\rangle \approx e^{-i\theta Z}\,|+\rangle\,, \quad \text{where} \quad \theta = \langle\psi_g|O|\psi_g\rangle\Delta t. \tag{69}$$

The ancilla has undergone a Z-rotation by an angle θ that is directly proportional to the expectation value we seek. Thus, the tomography problem reduces to a quantum metrology problem. We can employ standard quantum sensing techniques (covered in Lecture 1) on the ancilla qubit to determine θ with high precision, from which we extract $\langle\psi_g|O|\psi_g\rangle$.

Because the state $|\psi_g\rangle$ is preserved, we can repeat this entire procedure for O_1, then O_2, and so on, for all M observables, using only our single catalytic copy of $|\psi_g\rangle$.