Ph 220: Quantum Learning Theory
Lecture Note 6: Learning Noisy Quantum Devices

Hsin-Yuan Huang (Robert)

Caltech

1 Introduction

Question: Given access to a noisy quantum device capable of implementing many physical operations, how can we learn a complete description of all the physical operations it performs?

A quantum experiment performed on a noisy device can be represented as a sequence of operations:

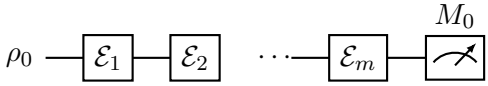

The components are:

- ρ_0: The initial state preparation, such as $\rho_0 \approx |0^n\rangle\langle 0^n|$.
- $\mathcal{E}_i$: A sequence of noisy quantum gates or circuits.
- M_0: The final measurement, described by a positive operator-valued measure (POVM), such as a noisy computational basis measurement $\{|b\rangle\langle b|\}_{b\in\{0,1\}^n}$.

Our objective is to learn a complete description of each component: the initial state ρ_0, the POVM M_0, and all quantum channels $\mathcal{E}_i$. Can this be achieved?

1.1 Pauli Representation

To formulate the learning problem precisely, we first establish a mathematical representation for all physical operations. The quantum benchmarking community commonly employs the Pauli representation. The set of n-qubit Pauli operators $P \in \{I, X, Y, Z\}^{\otimes n}$ forms an orthogonal basis for the space of $d \times d$ matrices, where $d = 2^n$. There are $d^2 = 4^n$ such operators. This allows us to represent states, channels, and measurements as vectors or matrices in a 4^n-dimensional space.

1.1.1 State Representation

A quantum state ρ (a $2^n \times 2^n$ density matrix) admits a unique expansion in the Pauli basis:

$$\rho = \frac{1}{2^n} \sum_P \mathrm{tr}(P\rho)\, P.$$

We define the *Pauli vector* v_ρ with components $(v_\rho)_P = \mathrm{tr}(P\rho)$. The component corresponding to the identity operator is always $(v_\rho)_I = \mathrm{tr}(I\rho) = 1$.

1.1.2 Measurement Representation

A POVM $\mathcal{M} = \{F_b\}$ can be represented analogously. Each POVM element F_b is a Hermitian operator and can be represented by a vector w_{F_b} with components $(w_{F_b})_P = \mathrm{tr}(PF_b)$. The Born rule probability $\mathrm{tr}(F_b\rho)$ becomes a simple inner product in Pauli space:

$$\mathrm{tr}(F_b\rho) = \mathrm{tr}\left(\left(\frac{1}{2^n}\sum_P \mathrm{tr}(PF_b)\,P\right)\rho\right) = \frac{1}{2^n}\sum_P \mathrm{tr}(PF_b)\ \mathrm{tr}(P\rho) = \frac{1}{2^n} w_{F_b}\cdot v_\rho.$$

This transforms the quantum mechanical trace operation into a standard vector inner product in the 4^n-dimensional Pauli space.

1.1.3 Channel Representation: The Pauli Transfer Matrix

Any quantum channel $\mathcal{E}$ can be represented by its *Pauli Transfer Matrix* (PTM), denoted $R_\mathcal{E}$. This is a $4^n \times 4^n$ real matrix describing how the channel maps input Pauli operators to linear combinations of output Pauli operators:

$$\mathcal{E}(P) = \sum_Q (R_\mathcal{E})_{Q,P}\, Q.$$

Learning the channel $\mathcal{E}$ is equivalent to learning all $4^n \times 4^n = 16^n$ entries of the matrix $R_\mathcal{E}$.

Goal (Reformulated): Our objective of learning the noisy quantum device can be rephrased precisely as follows. We seek to learn the Pauli vector v_{ρ_0} for the initial state ρ_0, the Pauli vectors w_{F_b} for all measurement outcomes b, and the PTM $R_{\mathcal{E}_i}$ for every quantum gate or circuit $\mathcal{E}_i$.

1.2 Unlearnability

This learning objective faces a fundamental obstacle. Certain physically distinct descriptions of the device produce identical experimental outcomes, rendering them completely indistinguishable. Consider a simple scenario where all quantum gates $\mathcal{E}_i$ are perfect (we can implement any unitary quantum circuit without error or noise). Even in this idealized setting, can we uniquely learn ρ_0 and $M_0 = \{F_b\}$? The answer is no.

To see why, let $D_\lambda(\rho)$ denote the depolarizing channel:

$$D_\lambda(\rho) = (1-\lambda)\rho + \lambda\frac{I}{2^n}.$$

Now consider two possible worlds describing the device:

World A: Noisy State Preparation

- Initial state: $\rho_0 = D_\lambda(|0^n\rangle\langle 0^n|)$
- Measurement: $M_0 = \{F_b\} = \{|b\rangle\langle b|\}_{b\in\{0,1\}^n}$ (perfect)

World B: Noisy Measurement

- Initial state: $\rho_0 = |0^n\rangle\langle 0^n|$ (perfect)
- Measurement: $M_0 = \{F_b\} = \{D_\lambda(|b\rangle\langle b|)\}_b$

For any experiment we perform (implementing an arbitrary unitary U), the probability of outcome b is $P(b) = \mathrm{tr}(F_b \cdot U \cdot \rho_0 \cdot U^\dagger)$.

In World A, the probability of outcome b after applying unitary U is:

$$\begin{aligned} P_A(b) &= \mathrm{tr}\left((|b\rangle\langle b|) \cdot U \cdot \left((1-\lambda)|0^n\rangle\langle 0^n| + \lambda\frac{I}{2^n}\right) \cdot U^\dagger\right) \\ &= (1-\lambda)\ \mathrm{tr}\left(|b\rangle\langle b| \cdot U|0^n\rangle\langle 0^n|U^\dagger\right) + \lambda\ \mathrm{tr}\left(|b\rangle\langle b| \cdot U\frac{I}{2^n}U^\dagger\right) \\ &= (1-\lambda)\,|\langle b|U|0^n\rangle|^2 + \frac{\lambda}{2^n}\ \mathrm{tr}(|b\rangle\langle b|) \\ &= (1-\lambda)\,|\langle b|U|0^n\rangle|^2 + \frac{\lambda}{2^n}. \end{aligned}$$

In World B, the probability of outcome b after applying unitary U is:

$$\begin{aligned} P_B(b) &= \mathrm{tr}\left(D_\lambda(|b\rangle\langle b|) \cdot U|0^n\rangle\langle 0^n|U^\dagger\right) \\ &= \mathrm{tr}\left(\left((1-\lambda)|b\rangle\langle b| + \lambda\frac{I}{2^n}\right) \cdot U|0^n\rangle\langle 0^n|U^\dagger\right) \\ &= (1-\lambda)\ \mathrm{tr}\left(|b\rangle\langle b| \cdot U|0^n\rangle\langle 0^n|U^\dagger\right) + \lambda\ \mathrm{tr}\left(\frac{I}{2^n} \cdot U|0^n\rangle\langle 0^n|U^\dagger\right) \\ &= (1-\lambda)\,|\langle b|U|0^n\rangle|^2 + \frac{\lambda}{2^n}\ \mathrm{tr}(U|0^n\rangle\langle 0^n|U^\dagger) \\ &= (1-\lambda)\,|\langle b|U|0^n\rangle|^2 + \frac{\lambda}{2^n}. \end{aligned}$$

Since $P_A(b) = P_B(b)$ for *any* unitary U we implement, these two worlds are completely indistinguishable from experimental data. This demonstrates that the learning problem, as originally stated, is **unlearnable**. We cannot uniquely determine whether the noise parameter λ resides in the state preparation or in the measurement. This ambiguity is known as **gauge freedom** in quantum characterization and benchmarking.

Remark 1 (Physical Realities and Gauge Freedoms)**.** *These two worlds correspond to distinct physical realities. If the noise is in state preparation, we should focus on improving the state preparation process. If the noise is in measurement, we should instead improve the measurement apparatus. The unlearnability arises from our restricted experimental access to the device components, not because the two scenarios are physically identical. If we had access to additional physical operations (such as multiple distinct state preparations or measurements), we could leverage the "purest" available operations to partially disentangle these ambiguities. While some gauge freedom remains inherent, it diminishes as the device offers more diverse and higher-fidelity operations.*

1.3 Gauges

This observation motivates a formal treatment of the ambiguity.

Definition 1 (Gauge)**.** *A gauge is the equivalence class of all possible physical descriptions of a noisy quantum device (characterized by different ρ_0, M_0, and $\{\mathcal{E}_i\}$) that produce identical experimental outcomes for all possible experiments.*

The unlearnability problem can now be precisely formulated: what are all the gauges? If we can characterize the gauge freedom completely, we can impose a convention to "fix" the gauge, thereby

making the remaining description unique and learnable. Our strategy for addressing gauge freedom proceeds in three steps.

Strategy for Handling Gauge Freedom:

1. **Identify gauge transformations:** Characterize the family of maps that transform one device description to another experimentally indistinguishable description.
2. **Fix a gauge:** Impose a specific, convenient convention that uniquely selects one description from each equivalence class identified in Step 1.
3. **Learn everything:** Demonstrate that once the gauge is fixed, all remaining device parameters become uniquely learnable from experimental data.

2 Gauge Identification, Gauge Fixing, and Learning

2.1 Identifying Gauge Transformations

Let $\mathcal{L}$ be any invertible linear map (a superoperator) on the space of $d \times d$ matrices, where $d = 2^n$. We can construct a new device description (which we denote with primes) by applying the following transformations:

- State preparation: $\rho_0 \mapsto \rho_0' = \mathcal{L}(\rho_0)$
- Measurement operators: $M_0 = \{F_b\} \mapsto M_0' = \{F_b'\} = \{\mathcal{L}^{-\dagger}(F_b)\}_b$
- Quantum channels: $\mathcal{E}_i \mapsto \mathcal{E}_i' = \mathcal{L} \circ \mathcal{E}_i \circ \mathcal{L}^{-1}$

Here $\mathcal{L}^\dagger$ denotes the adjoint of $\mathcal{L}$ with respect to the Hilbert-Schmidt inner product, defined by $\mathrm{tr}(A \cdot \mathcal{L}(B)) = \mathrm{tr}(\mathcal{L}^\dagger(A) \cdot B)$ for all operators A and B, and $\mathcal{L}^{-\dagger} = (\mathcal{L}^{-1})^\dagger = (\mathcal{L}^\dagger)^{-1}$.

This transformation preserves all experimental outcomes. To see this, consider any experiment $\rho_0 \to \mathcal{E}_1 \to \ldots \to \mathcal{E}_m \to M_0$. The probability of outcome b after the gate sequence $\mathcal{C} = \mathcal{E}_m \circ \cdots \circ \mathcal{E}_1$ is $P(b) = \mathrm{tr}(F_b \cdot \mathcal{C}(\rho_0))$. In the transformed (primed) gauge, the probability $P'(b)$ is:

$$
\begin{aligned}
P'(b) &= \mathrm{tr}\left(F_b' \cdot \mathcal{C}'(\rho_0')\right) \\
&= \mathrm{tr}\left(\mathcal{L}^{-\dagger}(F_b) \cdot (\mathcal{L} \circ \mathcal{E}_m \circ \mathcal{L}^{-1}) \circ \cdots \circ (\mathcal{L} \circ \mathcal{E}_1 \circ \mathcal{L}^{-1}) \circ \mathcal{L}(\rho_0)\right) \\
&= \mathrm{tr}\left(\mathcal{L}^{-\dagger}(F_b) \cdot \mathcal{L} \circ \mathcal{E}_m \circ \cdots \circ \mathcal{E}_1 \circ (\mathcal{L}^{-1} \circ \mathcal{L})(\rho_0)\right) \quad \text{(telescoping)} \\
&= \mathrm{tr}\left(\mathcal{L}^{-\dagger}(F_b) \cdot \mathcal{L}\left((\mathcal{E}_m \circ \cdots \circ \mathcal{E}_1)(\rho_0)\right)\right) \\
&= \mathrm{tr}\left(F_b \cdot (\mathcal{E}_m \circ \cdots \circ \mathcal{E}_1)(\rho_0)\right) = P(b), \quad \text{(by the adjoint property)}
\end{aligned}
$$

where in the last step we used the definition of the adjoint operator.

Since all experimental predictions remain identical under this transformation, any invertible linear map $\mathcal{L}$ generates a gauge transformation.

2.2 Fixing a Gauge

To resolve the ambiguity, we must select one specific description from each gauge equivalence class. We accomplish this through a sequence of calibration experiments.

1. **Define ideal observables:** Consider the complete set of 4^n ideal Pauli observables $P \in \{I, X, Y, Z\}^{\otimes n}$.

2. **Define calibration circuits:** For each Pauli observable P, define a corresponding depth-1 calibration circuit $\mathcal{C}_P$. For a single qubit, the correspondence is:

   - $I \to \mathcal{C}_I = I$ (identity circuit)
   - $Z \to \mathcal{C}_Z = I$ (identity circuit)
   - $X \to \mathcal{C}_X = H$ (Hadamard gate)
   - $Y \to \mathcal{C}_Y = SH$ (Phase gate followed by Hadamard gate)

   For n qubits, $\mathcal{C}_P$ is the tensor product of the corresponding single-qubit circuits. These circuits represent the ideal operations one would implement to measure the expectation value $\mathrm{tr}(P\rho)$ of a Pauli observable P.

3. **Calibrate effective observables:** We now experimentally define a set of "effective observables" $\tilde{P}$ based on what the noisy device actually implements. We execute the noisy circuit $\mathcal{E}_{\mathcal{C}_P}$ and measure with the noisy measurement $M_0 = \{F_b\}$.

   The effective observable $\tilde{P}$ is reconstructed from the experimental probabilities using the Heisenberg picture:

   $$\tilde{P} = \sum_{b \in \{0,1\}^n} (-1)^{f_P(b)} \, \mathcal{E}^\dagger_{\mathcal{C}_P}(F_b),$$

   where $f_P(b)$ denotes the parity function corresponding to measuring the Pauli P in the computational basis. Specifically, $f_P(b)$ counts the number of qubits where both P and b have a nontrivial value (not identity or zero), modulo 2. This encodes how one measures a Pauli observable P by performing computational basis measurement after the circuit $\mathcal{C}_P$.

4. **Construct the gauge transformation:** If the noise level is sufficiently low, the set of 4^n effective observables $\{\tilde{P}\}$ will be linearly independent and therefore form a complete basis for the space of observables, just as the ideal Paulis $\{P\}$ do.

   Consequently, there exists a unique invertible linear map $\mathcal{L}^\dagger$ (the adjoint of the gauge transformation) that relates the ideal Pauli basis to our experimentally determined basis:

   $$\mathcal{L}^\dagger(P) = \tilde{P}, \quad \forall P \in \{I, X, Y, Z\}^{\otimes n}.$$

   This uniquely determines the gauge transformation $\mathcal{L}$.

5. **Fix the gauge:** We now fix the gauge by applying the gauge transformation $\mathcal{L}$ to our entire device description. We denote quantities in this new, gauge-fixed description with a superscript "(new)".

   What happens to our effective observables $\tilde{P}$ in this new gauge? The transformation law for observables is $\tilde{P} \mapsto \tilde{P}^{(\mathrm{new})} = \mathcal{L}^{-\dagger}(\tilde{P})$.

   Substituting our definition of $\mathcal{L}^\dagger$, we obtain:

   $$\tilde{P}^{(\mathrm{new})} = \mathcal{L}^{-\dagger}(\tilde{P}) = \mathcal{L}^{-\dagger}(\mathcal{L}^\dagger(P)) = P.$$

Under this gauge-fixed description, our experimentally measured observables $\tilde{P}^{(\text{new})}$ are, by construction, identical to the ideal Pauli observables P.

In the gauge-fixed description, we have perfect Pauli observables by construction. This is the essence of gauge fixing: we have chosen our description such that a specific set of operations (in this case, the effective observables) takes on a simple, canonical form.

2.3 Learning Everything

Having fixed the gauge, we can now learn all remaining components of the device uniquely.

1. **Learning the initial state $\rho_0^{(\text{new})}$:** The initial state in the gauge-fixed description can be fully characterized through standard quantum state tomography. We experimentally measure all 4^n expectation values:

$$\text{tr}\left(P \cdot \rho_0^{(\text{new})}\right) = \text{tr}\left(\tilde{P}^{(\text{new})} \cdot \rho_0^{(\text{new})}\right),$$

which is possible because we have fixed our gauge such that the experimental observables $\tilde{P}^{(\text{new})}$ coincide with the ideal Paulis P. This standard tomography procedure reconstructs the complete Pauli vector $v_{\rho_0^{(\text{new})}}$.

2. **Learning a state basis:** Consider an additional set of 4^n depth-1 circuits corresponding to preparations $D \in \{I, X, H, SH\}^{\otimes n}$. We execute these (noisy) circuits $\mathcal{E}_D^{(\text{new})}$ starting from the now-known initial state $\rho_0^{(\text{new})}$ to generate a family of 4^n output states:

$$\rho_D^{(\text{new})} = \mathcal{E}_D^{(\text{new})}(\rho_0^{(\text{new})}).$$

We can learn all of these states $\rho_D^{(\text{new})}$ (that is, determine all their Pauli vectors $v_{\rho_D^{(\text{new})}}$) using the tomography procedure from step 1. If the noise level is sufficiently low, these 4^n states will be linearly independent and therefore form a complete basis for the space of density matrices.

3. **Learning an arbitrary channel $\mathcal{E}^{(\text{new})}$:** Consider any unknown noisy gate $\mathcal{E}^{(\text{new})}$. We can determine its complete Pauli Transfer Matrix $R_{\mathcal{E}^{(\text{new})}}$ as follows.

We experimentally measure the $4^n \times 4^n$ matrix of expectation values:

$$E_{P,D}^{(\text{new})} = \text{tr}\left(\tilde{P}^{(\text{new})} \cdot \mathcal{E}^{(\text{new})}(\rho_D^{(\text{new})})\right) = \text{tr}\left(P \cdot \mathcal{E}^{(\text{new})}(\rho_D^{(\text{new})})\right).$$

We use our perfect Pauli measurements (from gauge fixing) on the states generated by applying $\mathcal{E}^{(\text{new})}$ to our basis of states $\{\rho_D^{(\text{new})}\}$.

The measured matrix $E_{P,D}^{(\text{new})}$ is closely related to the PTM. The PTM $R_{\mathcal{E}^{(\text{new})}}$ is defined in the Pauli basis $\{Q\}$ as $R_{P,Q}$, whereas the matrix we measured, $E_{P,D}^{(\text{new})}$, uses the Pauli basis for outputs (index P) and the D-state basis for inputs (index D).

Since the states $\{\rho_D^{(\text{new})}\}$ form a complete basis, there exists an invertible change-of-basis matrix K relating them to the Pauli basis: $\rho_D^{(\text{new})} = \sum_Q K_{Q,D}\, Q$. We can determine K explicitly because we learned all the $\rho_D^{(\text{new})}$ states in step 2.

We can now solve for the PTM through a change of basis:

$$R_{\mathcal{E}^{(\text{new})}}^{P,Q} = \left(E^{(\text{new})} \cdot K^{-1}\right)_{P,Q}.$$

This determines the PTM for all quantum channels.

By fixing the gauge, we have established a self-consistent framework in which the state preparation ρ_0, the measurement operators M_0, and all gate channels $\mathcal{E}_i$ become uniquely defined and experimentally learnable. This completes our demonstration that gauge fixing enables complete characterization of noisy quantum devices.

Remark 2 (Gate Set Tomography)**.** *The procedure described above (fixing a gauge and then performing full tomography on every physical operation) is commonly known as* ***Gate Set Tomography (GST)****. It is generally very inefficient due to exponential scaling with system size n in multiple steps. Additional structural assumptions are required to make the learning efficient.*

3 Efficient Learning via Noise Twirling

Question: How can we achieve more efficient learning of noisy quantum devices?

Answer: We can circumvent the exponential scaling of GST by assuming **gate-independent noise** and applying a technique called **noise twirling**.

3.1 The Gate-Independent Noise Model

Instead of assuming every gate $\mathcal{E}_i$ has a unique, arbitrary noise channel, we assume the noise possesses additional structure. We model the noisy device as follows:

- **State Preparation:** The noisy preparation consists of an ideal preparation followed by a noise channel $\mathcal{N}_1$:
$$\rho_0 = \mathcal{N}_1(|0^n\rangle\langle 0^n|).$$
- **Measurement:** The noisy measurement consists of a noise channel $\mathcal{N}_2$ followed by ideal measurement:
$$M_0 = \{\mathcal{N}_2^\dagger(|b\rangle\langle b|)\}_{b\in\{0,1\}^n}.$$
- **Gates:** For any ideal Clifford gate $C \in \mathrm{Cl}_n$, the noisy implementation $\mathcal{E}_C$ is modeled as the ideal unitary followed by a noise channel $\mathcal{N}$:
$$\mathcal{E}_C(\rho) = \mathcal{N}(C\rho C^\dagger).$$

Schematically, a circuit implementation looks like this:

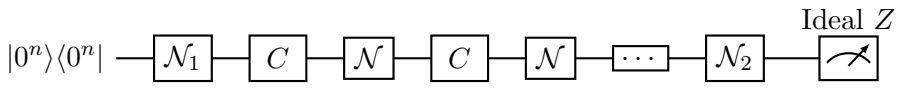

The crucial assumption is that the same noise channel $\mathcal{N}$ applies after every gate, independent of which gate was implemented. This gate-independent noise model captures the dominant noise sources in many physical implementations.

3.2 Noise Twirling

To simplify the learning of these noise channels, we employ **noise twirling** (also known as randomized compiling). The core mathematical insight is that conjugating an arbitrary channel by random Pauli operators symmetrizes the noise, transforming it into a **Pauli channel** (a channel whose PTM is diagonal in the Pauli basis). We apply this technique in two contexts.

3.2.1 Gate Twirling

This is the standard form of randomized compiling used to simplify the noise $\mathcal{N}$ occurring after a Clifford gate C. We twirl the noisy gate by inserting a random Pauli P before the gate and a compiling correction Pauli P^c after the gate such that the logical operation remains unchanged (that is, $P^cCP = C$).

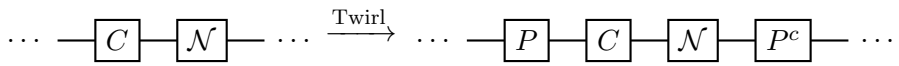

Physically, the twirling is implemented by executing the Clifford $P \to C$ instead of the original C for a randomly chosen Pauli P. The correction P^c must be combined with the physical operation that immediately follows.

Mathematically, we average the channel over all $P \in \{I, X, Y, Z\}^{\otimes n}$. The effective channel $\mathcal{E}_{\text{eff}}$ becomes:

$$\mathcal{E}_{\text{eff}}(\rho) = \mathbb{E}_P \left[P^c \, \mathcal{N}(CP\rho PC^\dagger) \, P^c \right].$$

Since $P^c = CPC^\dagger$, this can be rewritten as:

$$\mathcal{E}_{\text{eff}}(\rho) = C \left(\mathbb{E}_P \left[P \, (C^\dagger \mathcal{N} C)(P\rho P) \, P \right] \right) C^\dagger.$$

The effective noise channel becomes:

$$\mathcal{N}_{\text{twirled}}(\rho) = \mathbb{E}_P \left[P \tilde{\mathcal{N}}(P\rho P) \, P \right],$$

where $\tilde{\mathcal{N}} = C^\dagger \circ \mathcal{N} \circ C$ is the noise channel $\mathcal{N}$ rotated by the ideal Clifford circuit C. Note that $\tilde{\mathcal{N}}$ remains close to the identity channel.

Proof of Diagonalization: Let $\tilde{\mathcal{N}} = C^\dagger \circ \mathcal{N} \circ C$ be the rotated noise channel. We show that the twirled channel,

$$\mathcal{N}_{\text{twirled}}(\rho) = \mathbb{E}_P \left[P \tilde{\mathcal{N}}(P\rho P) \, P \right],$$

is a Pauli channel (its PTM is diagonal).

We represent the action of $\tilde{\mathcal{N}}$ in the Pauli basis. For any Pauli operator A, we can expand the output as:

$$\tilde{\mathcal{N}}(A) = \sum_B \chi_{BA} \, B,$$

where χ_{BA} are the PTM coefficients.

Now, consider the action of the twirled channel on a specific Pauli basis element A:

$$\begin{aligned}
\mathcal{N}_{\text{twirled}}(A) &= \mathbb{E}_P \left[P \tilde{\mathcal{N}}(PAP) \, P \right] \\
&= \mathbb{E}_P \left[P \tilde{\mathcal{N}} \left((-1)^{\{P,A\}} A \right) \, P \right] \quad (\text{using } PAP = (-1)^{\{P,A\}} A) \\
&= \mathbb{E}_P \left[(-1)^{\{P,A\}} P \left(\sum_B \chi_{BA} \, B \right) P \right] \\
&= \sum_B \chi_{BA} \left(\mathbb{E}_P \left[(-1)^{\{P,A\}} PBP \right] \right),
\end{aligned}$$

where $\{P, A\} = 1$ if P and A anticommute and 0 if they commute.

Using the commutation relation $PBP = (-1)^{\{P,B\}}B$, the term inside the expectation becomes:

$$\mathbb{E}_P\left[(-1)^{\{P,A\}+\{P,B\}}B\right] = B \cdot \mathbb{E}_P\left[(-1)^{\{P,A\}+\{P,B\}}\right].$$

The expectation value $\mathbb{E}_P[(-1)^{\{P,A\}+\{P,B\}}]$ equals 1 if $A = B$ (since the exponent is always 0 (mod 2)) and equals 0 if $A \neq B$ (due to the orthogonality of characters). Therefore, only the term where $B = A$ survives:

$$\mathcal{N}_{\text{twirled}}(A) = \chi_{AA}\, A = \lambda_A\, A.$$

This confirms that the effective noise channel is diagonal in the Pauli basis, with eigenvalues λ_A. Thus, the twirled channel is a Pauli channel.

3.2.2 Measurement Twirling

In addition to gate errors, we must address readout errors. In our gate-independent noise model, we cannot treat measurement noise $\mathcal{N}_2$ in isolation. This is because performing the measurement typically requires a preceding control operation (or "idling," which acts as an identity gate). In our model, any operation incurs the gate noise $\mathcal{N}$.

Therefore, the effective noise acting before the ideal computational basis measurement is the composition:

$$\Lambda_{\text{readout}} = \mathcal{N}_2 \circ \mathcal{N}.$$

To simplify this combined noise channel, we apply **measurement twirling**. We insert a random Pauli gate P immediately before the measurement.

Protocol:

1. Apply a random Pauli gate $P \in \{I, X, Y, Z\}^{\otimes n}$ immediately before the measurement.
2. According to our noise model, the physical implementation of P consists of the ideal P followed by the gate noise $\mathcal{N}$.
3. The measurement process immediately follows, adding the measurement noise $\mathcal{N}_2$.
4. The total noise $\Lambda_{\text{readout}} = \mathcal{N}_2 \circ \mathcal{N}$ acts on the state *after* the ideal P has been applied.
5. We measure in the computational basis.
6. We perform a classical correction: if P anticommutes with Z on qubit i (that is, P maps between $|0\rangle$ and $|1\rangle$ on that qubit), we flip the classical measurement outcome for that qubit.

The effective POVM element for a specific outcome b is obtained by averaging over the Pauli group. Note that the classical flip is equivalent to applying P to the measurement operator in the Heisenberg picture for Z-basis measurements. This means we are effectively performing:

$$P \to \mathcal{N} \to \mathcal{N}_2 \to P \to \{|b\rangle\langle b|\}_{b\in\{0,1\}^n}.$$

Through a calculation analogous to gate twirling, the arbitrary combined readout noise $\Lambda_{\text{readout}} = \mathcal{N}_2 \circ \mathcal{N}$ is transformed into a Pauli channel. Furthermore, the action of a Pauli channel before the Z-basis measurement is equivalent to applying a correlated distribution of bitflips. The effective measurement behaves like an ideal measurement preceded by a simple symmetric error channel (characterized by bitflip probabilities), parametrized by the diagonal elements of the noise PTM. This drastically reduces the complexity of learning the readout noise.

3.3 The Simplified Learning Landscape

After applying both gate twirling and measurement twirling, the description of our noisy device undergoes a dramatic simplification. We have effectively transformed a device with arbitrary, complex, coherent errors into a device characterized solely by **Pauli noise**. In this "twirled world," the effective circuit execution looks like this:

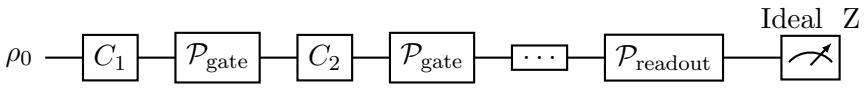

where:

- C_1, C_2 are the ideal Clifford unitaries we intended to implement.
- $\mathcal{P}_{\text{gate}}$ is the Pauli channel resulting from twirling the gate noise $\mathcal{N}$. It is diagonal in the Pauli basis: $\mathcal{P}_{\text{gate}}(Q) = \lambda_Q^{\text{gate}} Q$.
- $\mathcal{P}_{\text{readout}}$ is the Pauli channel resulting from twirling the combined readout noise $\mathcal{N}_2 \circ \mathcal{N}$. It is also diagonal: $\mathcal{P}_{\text{readout}}(Q) = \lambda_Q^{\text{readout}} Q$.

Because a Pauli channel is equivalent to a classical probability distribution over Pauli operators, the problem of learning quantum channels is reduced to the problem of learning classical probability distributions. This can be accomplished much more efficiently, especially if we impose locality assumptions about the classical probability distribution (such as assuming correlations are limited to bounded-size regions). This reduction from quantum to classical learning represents a fundamental simplification that makes efficient device characterization tractable.