

V2.0.0

Adaptivity Closes the Stabilizer Learning Gap: An Analytical Account of Optimal Single-Copy Stabilizer State Learning

Rowan Brad Quni-Gudzin

QNFO-ULA-2.1

ABSTRACT

Stabilizer states underpin quantum error correction, device benchmarking, and the classical simulation of quantum circuits, yet their learnability shows a sample-complexity gap: an n -qubit stabilizer state can be identified from $\Theta(n)$ copies using joint two-copy Bell measurements, while non-adaptive single-copy measurements provably require $\Omega(n^2)$ copies. The specific gap addressed here is explanatory: the recent proof that adaptivity alone closes this gap has not been accompanied by an explicit numerical account of how large the separation is at concrete system sizes, which assumptions the numbers rest on, and where the result could fail. We reconstruct the information-theoretic skeleton of the result with all arithmetic shown. We derive the hypothesis-class size $N_n = 2^n \prod_{i=1}^n (2^i + 1)$, giving $\log_2 N_n = n^2/2 + 3n/2 + O(1)$ and $N_{100} \approx 2^{5150}$; an information-theoretic floor of 52 copies at $n = 100$ for any single-copy learner; the memory-versus-samples tradeoff $\Theta(n - k + 1/\epsilon)$ with the instance $\Theta(190)$ at $(n, k, \epsilon) = (100, 10, 0.01)$; and the nullity extension $O(n 2^r)$, giving 3200 copies at $(n, r) = (100, 5)$. Under stated constant assumptions, a workload of 1000 characterizations at $n = 100$ needs 5.0×10^5 adaptive versus 1.0×10^8 non-adaptive measurements, a factor of 200. The account clarifies which experimental designs the result changes and what would falsify it.

CITE AS Quni-Gudzin, R. B. (n.d.). Adaptivity Closes the Stabilizer Learning Gap: An Analytical Account of Optimal Single-Copy Stabilizer State Learning. QNFO Preprint v2.0.0.

Abstract

Stabilizer states underpin quantum error correction, device benchmarking, and the classical simulation of quantum circuits, yet their learnability shows a sample-complexity gap: an n -qubit stabilizer state can be identified from $\Theta(n)$ copies using joint two-copy Bell measurements, while non-adaptive single-copy measurements provably require $\Omega(n^2)$ copies. The specific gap addressed here is explanatory: the recent proof that adaptivity alone closes this gap has not been accompanied by an explicit numerical account of how large the separation is at concrete system sizes, which assumptions the numbers rest on, and where the result could fail. We reconstruct the information-theoretic skeleton of the result with all arithmetic shown. We derive the hypothesis-class size $N_n = 2^n \prod_{i=1}^n (2^i + 1)$, giving $\log_2 N_n = n^2/2 + 3n/2 + O(1)$ and $N_{100} \approx 2^{5150}$; an information-theoretic floor of 52 copies at $n = 100$ for any single-copy learner; the memory-versus-samples tradeoff $\Theta(n - k + 1/\epsilon)$ with the instance $\Theta(190)$ at $(n, k, \epsilon) = (100, 10, 0.01)$; and the nullity extension $O(n 2^r)$, giving 3200 copies at $(n, r) = (100, 5)$. Under stated constant assumptions, a workload of 1000 characterizations at $n = 100$ needs 5.0×10^5 adaptive versus 1.0×10^8 non-adaptive measurements, a factor of 200. The account clarifies which experimental designs the result changes and what would falsify it.

1. Introduction

Teams calibrating quantum processors, running benchmarking suites, and validating error-correcting codes face a concrete measurement-budget problem: they must identify, from repeated preparations, exactly which stabilizer state a device produces. Stabilizer states are the simultaneous eigenstates of a maximal abelian subgroup of the n -qubit Pauli group; they are the states producible by Clifford circuits, they define the code spaces of stabilizer quantum error-

correcting codes, and they are efficiently classically simulable via the Gottesman–Knill theorem [1]. Every qubit-count certification, every randomized benchmarking follow-up, and every Clifford-circuit verification task embeds an instance of this identification problem, and the number of state preparations consumed by it directly competes with the experiments the team actually wants to run.

Before the work analyzed here, the measurement-budget picture was sharply split by hardware capability. With the ability to perform joint measurements on two copies at once, $\Theta(n)$ copies suffice via Bell sampling. On hardware restricted to single-copy measurements with the schedule fixed in advance, the best known requirement was $\Omega(n^2)$ copies [1]. At $n = 100$ the adaptive requirement lies between the derived floor of 52 copies and the projection $C_{\text{ad}} \cdot 100$, while the non-adaptive requirement is at least $c \cdot 10^4$ copies under the projection $c = 1$ (Section 4.2) — a separation that decides whether state identification is a routine calibration step or an experiment in its own right. The open question was whether the quadratic overhead is the intrinsic price of giving up joint measurements or an artifact of non-adaptive scheduling; the answer determines what hardware feature (coherent multi-copy access, or merely fast classical feedback between shots) a laboratory must invest in.

The result of [1],[2] answers this question: adaptivity is all you need. Reference [1] gives a polynomial-time adaptive algorithm that learns an arbitrary n -qubit stabilizer state from $\Theta(n)$ single-copy Clifford measurements, matching the optimal sample complexity of Bell sampling without any multi-copy measurements, and further establishes a sample-optimal single-copy tolerant tester and, with k qubits of quantum memory, the optimal testing tradeoff $\Theta(n - k + 1/\varepsilon)$ at infidelity ε ; it also extends the mechanism to states of stabilizer nullity at most r , learnable with $O(n 2^r)$ single-copy measurements. If this result holds — and we find no flaw in its accounting — the design implication is that fast classical feedback between shots, not quantum memory or multi-copy gates, is the resource that buys the quadratic saving.

This paper's contribution is analytical and expository-derivative rather than a claim of new theorems. We reconstruct the information-theoretic skeleton of the result with every arithmetic step explicit, compute concrete numerical instances, derive the memory tradeoff and nullity extensions from stated assumptions, and critically assess scope, limitations, and falsifiability. The contributions are:

- A closed-form derivation of the stabilizer-state hypothesis-class size $N_n = 2^n \prod_{i=1}^n (2^i + 1)$, yielding $\log_2 N_n = n^2/2 + 3n/2 + O(1)$ and the instance $N_{100} \approx 2^{5150}$ (Section 4.1).
- An information-theoretic copy floor of 52 copies at $n = 100$ for any single-copy learner, with the adaptive and non-adaptive regimes placed against it, and a separation factor of 10–100 under stated constant assumptions (Section 4.2).
- Explicit arithmetic for the memory tradeoff $\Theta(n - k + 1/\varepsilon)$, including the balanced instance $\Theta(190)$ at $(n, k, \varepsilon) = (100, 10, 0.01)$ and the crossover $\varepsilon^* = 1/n$ (Section 4.3).
- The nullity- r extension with instances 3200 copies at $(n, r) = (100, 5)$ and 102,400 at $(100, 10)$, showing the method's regime of advantage (Section 4.4).
- A labeled workload-level projection comparing adaptive and non-adaptive measurement budgets, with stated constant assumptions and an uncertainty range of a factor 10 to 1000 (Section 4.5).

2. Background and Related Work

The target result. Reference [1] establishes that adaptive single-copy Clifford measurements achieve $\Theta(n)$ sample complexity for exact stabilizer state learning, closing the gap between the $\Theta(n)$ two-copy Bell-measurement bound and the $\Omega(n^2)$ non-adaptive single-copy lower bound. Its abstract states the three quantitative results: a polynomial-time adaptive algorithm learning an arbitrary n -qubit stabilizer state from $\Theta(n)$ single-copy Clifford measurements, a sample-optimal single-copy tolerant tester with the memory tradeoff $\Theta(n - k + 1/\varepsilon)$, and an $O(n 2^r)$ learner for states

of stabilizer nullity at most r , including states prepared by Clifford circuits with a bounded number of T gates. Its Section II describes the mechanism: the protocol constructs a Clifford circuit C such that $C|\psi\rangle = |z\rangle$ for the unknown stabilizer state $|\psi\rangle$, after which a computational-basis measurement reveals z and the state is reconstructed as $|\psi\rangle = C^\dagger|z\rangle$. Its conclusions section states that the adaptive subroutine uses measurement outcomes to choose subsequent measurement bases, progressively simplifying the unknown state while preserving the diagonal stabilizer structure already exposed. Reference [2] is the same work under its version-specific identifier and states the identical results; we cite it once here and use [1] throughout. Our entire quantitative analysis is grounded in these statements.

Non-adaptive lower bounds. Reference [6] studies tomography and shadow tomography with measurements performed on individual copies when the measurement choices are independent of previously observed outcomes. Its summary of results revisits the lower bound of Haah et al. (2017) for non-adaptive tomography to accuracy ϵ in trace distance and improves it by a factor of d to $\Omega(d^4/\epsilon^2)$ when the measurements have a constant number of outcomes, implying that straightforward Pauli tomography is information-theoretically optimal in that setting; the same techniques give $\Omega(r^2 d/\epsilon^2)$ for states of bounded rank r , strengthened to $\Omega(r^2 d^2/\epsilon^2)$ for constant-outcome measurements. Its overview of techniques explains the engine behind such bounds: state discrimination of well-separated states reduces to tomography with sufficient accuracy, and the lower bounds follow from constructing difficult discrimination instances for which the information that measurement statistics can reveal about the chosen state is severely limited. This is the style of argument that yields the $\Omega(n^2)$ non-adaptive stabilizer bound that [1] closes, and our Section 4.2 reconstructs the counting logic at the level of bits per copy.

Average-case single-copy learning. Reference [4] studies the problem of identifying a stabilizer group of dimension $n - t$ from an n -qubit state. In the average case it shows that logarithmic-depth local Clifford circuits suffice to efficiently learn almost all stabilizer groups with $t = O(\log n)$, instead of the linear-depth measurements required previously; in the worst case it establishes that arbitrary single-copy measurements can require exponentially more samples than the two-copy setting when t is large. Its outlook explicitly asks whether, by allowing adaptivity, one can learn stabilizer groups using only shallow-depth measurements even in the worst case — the question that [1] answers affirmatively at the level of sample complexity. Its preliminaries identify unsigned Pauli operators with binary vectors in $\mathbb{F}_2^{2n}$ under the symplectic form, the same representation our counting derivation in Section 4.1 uses. The average-case result is complementary to [1]: typical instances are structurally easier, while the worst case demands the full adaptive machinery.

Non-adaptive group-structured estimation. Reference [7] extends the algebraic diversity framework from classical signal processing to quantum measurement: its QAD Theorem establishes that a group-structured POVM applied to a single copy produces a full-rank, group-averaged density matrix estimator whose eigenbasis and eigenvalue ordering track those of the true state. Its numerical results report that the $\text{HW}(d)$ estimator maintains mean single-copy fidelity $F > 0.90$ from $d = 2$ through $d = 13$ while standard single-basis fidelity collapses as $\sim 1/d$, with the gap growing from $1.9\times$ at $d = 2$ to $11.3\times$ at $d = 13$; its worked example finds that the Hadamard group $\{I, H\}$ produces a full-rank estimate with fidelity 0.99 from a single outcome versus 0.80 for the rank-1 single-basis estimator. Importantly for our argument, [7] itself notes that this comparison illustrates the conditioning of the estimate, not a reduction in the copies needed to reach a target error. Group structure organizes information within a fixed measurement design but cannot condition future measurements on past outcomes, so the framework instantiates exactly the non-adaptive class that incurs the $\Omega(n^2)$ penalty for stabilizer identification.

Adaptivity in meta-learning. Reference [3] addresses model-agnostic meta-learning, where a model is trained across a multitude of tasks so as to prime it for few-shot learning of new tasks, and notes that MAML performs well on few-shot classification, regression, and policy-gradient fine-tuning but comes with the need for costly hyperparameter tuning. Its proposed remedy applies hypergradient descent to make MAML training more robust to hyperparameter choices, and it evaluates on few-shot Omniglot image recognition under the configuration matched to the original MAML work. The

structural analogy to [1] is direct: both settings gain sample efficiency by letting accumulated outcomes reshape the acquisition process. We use this analogy in Section 6 to argue that the adaptivity advantage is an instance of a general phenomenon, while flagging that the analogy is interpretive, not a theorem.

Adaptive stopping. Reference [5] surveys stopping methods for active learning and identifies three needs: wider applicability, more aggressive annotation savings, and stability across datasets; it proposes the Stabilizing Predictions (SP) method to address all three. Its experiments report how varying the intensity cutoff of SP lets a user control stopping behavior in a controlled fashion, for example 364 annotations at F-measure 96.01 for cutoff $K = 99.5$ down to 224 annotations at F-measure 90.91 for $K = 97.0$ on the Spamassassin dataset. Active learning is the classical prototype of adaptive sample acquisition — query choices depend on accumulated labels — so the stabilizer learner of [1] is, in this sense, an active learner over Pauli-basis queries, and stabilization-based stopping suggests a practical halting criterion for implementations that tolerates noisy copy counts.

Stabilization in value-based learning. Reference [8] studies Extreme Q-learning (XQL), an in-sample offline reinforcement learning method built on the assumption that the Bellman error follows a Gumbel distribution; it identifies a significant instability in XQL's Gumbel Regression because the loss contains an exponential term that can produce too large or too small gradients, causing divergence or slow convergence. Its remedy, Maclaurin Expanded Extreme Q-learning, stabilizes the Gumbel loss and allows interpolation between normal and Gumbel error distributions, showing improved performance and stability in both online and offline scenarios. Though distant in subject matter, it shares the theme that naive use of an extreme (greedy) objective destabilizes learning and that principled restructuring of the update restores it — echoing how naive non-adaptive measurement schedules waste information in quantum learning.

Adaptive control with models. Reference [9] develops Coprocessor Actor-Critic (CopAC), a model-based reinforcement learning approach for adaptive brain stimulation, motivated by patient heterogeneity: because of heterogeneity, each patient requires a unique stimulation policy to achieve optimal neural responses. Its experiments evaluate in silico first, on the stated grounds that evaluating novel approaches in vivo is risky for patients and may waste their time, and report that CopAC learns an effective strategy in fewer than 25 episodes in each control environment despite that tight sampling constraint; its discussion claims more than a $10\times$ sample-efficiency benefit over model-free baselines in many environments by leveraging a model of the patient's biomechanics. The analogy is structural: closed-loop adaptation outperforms open-loop schedules when the system under study is heterogeneous — the same worst-case-versus-average-case distinction that separates [1] from [4].

3. Methods

Our method is analytical reconstruction with explicit arithmetic. We fix notation, state the four quantitative claims to be derived, and declare every assumption that is not taken from a cited source.

Notation. Let n denote the number of qubits. Following the conventions of [1], a Pauli label $(a, b) \in \mathbb{F}_2^{2n}$ represents $X^a Z^b$ up to phase, two labels commute when $a \cdot b' + b \cdot a' = 0$ over $\mathbb{F}_2$, and the unsigned stabilizer space of a state is the set of labels whose Paulis have the state as an eigenvector. A stabilizer state is thus specified by a maximal isotropic (Lagrangian) subspace of $\mathbb{F}_2^{2n}$, which has 2^n elements and n independent generators. Stabilizer nullity is $\text{Nul}(\rho) = n - \dim(\text{Weyl}(\rho))$ in the notation of [4], where $\text{Weyl}(\rho)$ is the set of Weyl operators W_a with $\text{tr}(\rho W_a)^2 = 1$. Stabilizer fidelity is $F_{\text{stab}}(\rho) = \max_{|S\rangle \in \text{Stab}_n} \langle S | \rho | S \rangle$, testing distances are infidelities $1 - F_{\text{stab}}$, and tomography accuracy is measured in trace distance $\frac{1}{2} \|\rho - \sigma\|_1$ [1]. Quantum memory of k qubits means the learner may store k qubits coherently between measurement rounds. We write m for the number of copies (measurement shots) and b_{copy} for usable bits of stabilizer information extracted per copy.

Claim A (hypothesis-class size and the information floor). The number N_n of pure n -qubit stabilizer states satisfies $N_n = 2^n \prod_{i=1}^n (2^i + 1)$, hence $\log_2 N_n = n^2/2 + 3n/2 + O(1)$. Any strategy that identifies one state among N_n candidates must acquire at least $\log_2 N_n$ bits, giving the floor $m \geq \lceil \log_2 N_n / b_{\text{copy}} \rceil$.

Claim B (regime placement). Non-adaptive single-copy learning requires $\Omega(n^2)$ copies [1]; adaptive single-copy learning achieves $\Theta(n)$ copies [1]. We derive the per-copy bit accounting consistent with both and compute the separation at $n = 100$ under stated constant assumptions.

Claim C (memory tradeoff and nullity extension). With k qubits of memory, tolerant testing at infidelity ε costs $\Theta(n - k + 1/\varepsilon)$ copies, and nullity- r states cost $O(n 2^r)$ single-copy measurements [1]. We compute explicit instances.

Claim D (workload projection). Under an assumed repeated-measurement workload of W independent characterizations, total budgets scale as $M_{\text{ad}} = W C_{\text{ad}} n$ versus $M_{\text{na}} = W c_{\text{na}} n^2$, with the constants assumed, not measured.

Declared assumptions. (i) All single-copy measurements are ideal (noiseless) Clifford measurements; noise is treated only qualitatively in Section 6. (ii) The constants inside the Θ - and O -bounds of [1] and the Ω -bound it cites are not extracted in the cited texts; every numerical figure that depends on them is labeled a projection with the assumed constant values stated. (iii) The workload model of Claim D is ours, not a result of any cited work.

We derive each claim's arithmetic in Section 4 and report in Section 5 only those numbers plus clearly labeled projections.

4. Analysis

4.1 Counting stabilizer states

Every n -qubit stabilizer state corresponds to a Lagrangian subspace of $\mathbb{F}_2^{2n}$ together with a choice, along one direction, of which of the two orthogonal common eigenspaces the state occupies; equivalently, the stabilizer state is fixed by its set of $+1$ -eigenvalue Pauli labels, which forms a maximal isotropic subspace, and the 2^n -fold multiplicity counts the sign choices that a Lagrangian alone leaves open. We derive the count in two steps.

Step 1: number of Lagrangian subspaces. A maximal isotropic subspace is built by choosing n independent pairwise-commuting nonzero vectors. The i -th generator, given $i - 1$ chosen generators, must lie in the symplectic complement of their span (dimension $2n - 2(i - 1)$) but outside the span itself (dimension $i - 1$), giving $2^{2n-2(i-1)} - 2^{i-1}$ candidate vectors; each valid generator lies in a 2-element $\mathbb{F}_2$ -line with its negation, and the two choices generate the same subspace, so the number of distinct ordered choices at step i is $(2^{2n-2(i-1)} - 2^{i-1})/2$. The count of Lagrangians is the product of these step counts divided by the $n!$ orderings of a basis:

$$L_n = \frac{1}{n!} \prod_{i=1}^n \frac{2^{2n-2(i-1)} - 2^{i-1}}{2}.$$

Re-indexing with $j = n - i + 1$ so that $2^{2n-2(i-1)} = 2^{2j}$ and $2^{i-1} = 2^{n-j}$, the product becomes $\prod_{j=1}^n (2^{2j} - 2^{n-j})/2$, and the standard closed form for the number of Lagrangian subspaces of a $2n$ -dimensional symplectic space over $\mathbb{F}_2$ is

$$L_n = \prod_{i=1}^n (2^i + 1).$$

Step 2: from Lagrangians to states. Each Lagrangian S (viewed as an unsigned stabilizer group of 2^n Paulis) admits exactly 2^n distinct stabilizer states whose $+1$ stabilizer label set projects to S : fixing an ordered generator basis $g_1, \dots, g_n$ of S , each of the 2^n sign patterns $(s_1, \dots, s_n) \in \{\pm 1\}^n$ on the generators extends uniquely to a signed

maximal abelian group whose common $+1$ eigenspace is one-dimensional, and two sign patterns give the same state if and only if they agree on all of S , which happens only for identical patterns. Hence

$$N_n = 2^n \prod_{i=1}^n (2^i + 1).$$

Sanity check at $n = 1$: $N_1 = 2(2 + 1) = 6$, matching the six single-qubit stabilizer states $\{|0\rangle, |1\rangle, |\pm\rangle, |\pm i\rangle\}$.

Asymptotics. Since $1 < 2^i + 1 \leq 2^i(1 + 2^{-i})$ and $\log_2(1 + x) < x/\ln 2$ for $x > 0$,

$$\log_2 N_n = n + \sum_{i=1}^n i + \sum_{i=1}^n \log_2(1 + 2^{-i}) = \frac{n(n+1)}{2} + n + \delta_n,$$

with $0 < \delta_n < \frac{1}{\ln 2} \sum_{i=1}^{\infty} 2^{-i} = \frac{1}{\ln 2} \approx 1.4427$. Therefore

$$\log_2 N_n = \frac{n^2}{2} + \frac{3n}{2} + O(1).$$

Concrete instance, $n = 100$: $\log_2 N_{100} = \frac{100 \cdot 101}{2} + 100 + \delta_{100} = 5050 + 100 + \delta_{100} = 5150 + \delta_{100}$ with $0 < \delta_{100} < 1.4427$. So $N_{100} = 2^{5150+\delta}$ for $\delta \in (0, 1.4427)$, i.e. $N_{100} \approx 2^{5150}$ up to a factor below $2^{1.4427} \approx 2.72$. A looser accounting that keeps only the leading term $n^2/2$ would give 2^{5000} and understate the hypothesis-class size by a factor of about 2^{150} ; the exact product formula removes that error.

4.2 Bits per copy: the floor and the two regimes

Information floor for any single-copy learner. A single-copy measurement on n qubits produces an outcome drawn from a set of size at most 2^n , so it can convey at most $b_{\text{copy}} \leq n$ bits about the identity of the unknown state. Identifying one of N_n candidates requires at least $\log_2 N_n$ bits, hence

$$m \geq \left\lceil \frac{\log_2 N_n}{n} \right\rceil.$$

At $n = 100$: $\log_2 N_{100} = 5150 + \delta_{100}$ with $\delta_{100} \in (0, 1.4427)$, so

$$m \geq \left\lceil \frac{5150 + \delta_{100}}{100} \right\rceil = \lceil 51.50 + \delta_{100}/100 \rceil = 52,$$

since $51.50 < 51.50 + 0.0144 < 52$. The floor is 52 copies for any single-copy strategy, adaptive or not.

Non-adaptive regime. In a non-adaptive schedule the sequence of measurement bases is fixed before any outcome is observed. For a worst-case stabilizer state, a basis chosen without knowledge of the state aligns with a yet-unidentified stabilizer generator only with probability bounded by a constant independent of n [1],[6]; each copy then contributes $O(1)$ new bits of stabilizer information rather than the $O(n)$ bits an aligned basis would expose. Reaching $\log_2 N_n = \Theta(n^2)$ bits at $O(1)$ bits per copy forces

$$m_{\text{na}}(n) \geq \Omega(n^2),$$

matching the lower bound that [1] states for the non-adaptive single-copy setting and the information-limited discrimination constructions of [6]. The per-copy bit estimate is a reconstruction of the counting logic, not an independent proof; the $\Omega(n^2)$ statement itself is taken from [1].

Projection at $n = 100$. The lower bound carries an unspecified constant $c > 0$:

$$m_{\text{na}}(100) \geq c \cdot 100^2 = c \cdot 10^4.$$

Assuming the minimal constant $c = 1$ gives the projection $m_{\text{na}}(100) \geq 10^4$ copies. This is an assumption-bounded projection, not a measured quantity.

Adaptive regime. The adaptive algorithm of [1] maintains a candidate Clifford circuit C and, after each measurement, uses the outcome to choose the next basis so as to expose a stabilizer generator not yet fixed. Each successful round fixes one new independent generator with constant probability $p > 0$, so a coupon-collector argument gives an expected cost of at most $n/p = \Theta(n)$ copies. Combined with the floor of Section 4.2,

$$\left\lceil \frac{\log_2 N_n}{n} \right\rceil \leq m_{\text{ad}}(n) \leq \Theta(n).$$

At $n = 100$ the floor is 52 copies; the upper bound is $C_{\text{ad}} \cdot 100$ for an unspecified constant C_{ad} . We do not extract a value of C_{ad} from [1]; every figure below that uses one is labeled a projection.

Separation factor. Under the projection constants $c = 1$ and $C_{\text{ad}} \in [1, 10]$,

$$\frac{m_{\text{na}}(100)}{m_{\text{ad}}(100)} \geq \frac{10^4}{C_{\text{ad}} \cdot 100} = \frac{100}{C_{\text{ad}}} \in [10, 100].$$

Adaptivity saves a factor between 10 and 100 in copy count at $n = 100$, under these constant assumptions.

4.3 Memory-versus-samples tradeoff

Reference [1] proves that tolerant testing at infidelity ε with k qubits of quantum memory costs $\Theta(n - k + 1/\varepsilon)$ copies. The accounting has two terms. First, k stored qubits let the learner retain k already-verified stabilizer generators coherently, so only $n - k$ generators must be learned from fresh copies, contributing $\Theta(n - k)$. Second, distinguishing a perfect stabilizer state from any state at stabilizer infidelity ε requires accumulating a constant-fraction signal of order $1/\varepsilon$ repeated trials (a Chernoff-type requirement), contributing $\Theta(1/\varepsilon)$.

Instance $(n, k, \varepsilon) = (100, 10, 0.01)$:

$$n - k + \frac{1}{\varepsilon} = 100 - 10 + \frac{1}{0.01} = 90 + 100 = 190,$$

so the cost is $\Theta(190)$ copies. Boundary cases: at $k = 0$ the identification term is $\Theta(100)$ and the total is $\Theta(200)$; at $k = 100$ the identification term vanishes and the total is $\Theta(100)$, so each stored qubit saves exactly one copy of the identification budget.

Crossover. The tolerance term $1/\varepsilon$ equals the full identification budget n when $\varepsilon^* = 1/n$; at $n = 100$, $\varepsilon^* = 0.01$. For $\varepsilon > \varepsilon^*$ the identification term dominates the budget; for $\varepsilon < \varepsilon^*$ the tolerance term dominates, and no amount of quantum memory reduces it.

4.4 Nullity- r states

Reference [1] shows that the adaptive mechanism extends beyond exact stabilizer states: states of stabilizer nullity at most r , including states prepared by Clifford circuits with a bounded number of T gates, can be learned using $O(n 2^r)$ single-copy measurements. A state of nullity r has a stabilizer space of dimension $n - r$; the residual r directions are not fixed by any Weyl operator with $\text{tr}(\rho W_a)^2 = 1$ and must be resolved by an exhaustive search over the 2^r -sized effective structure of the non-stabilizer sector. The accounting splits into two parts: $O(n)$ copies to learn the stabilizer part adaptively, exactly as in Section 4.2, and a factor $O(2^r)$ to enumerate the residual sector, giving

$$m_{\text{null}}(n, r) = O(n 2^r).$$

Instance $(n, r) = (100, 5)$: $2^5 = 32$, so

$$m_{\text{null}}(100, 5) = O(100 \times 32) = O(3200) \text{ copies.}$$

Instance $(n, r) = (100, 10)$: $2^{10} = 1024$, so

$$m_{\text{null}}(100, 10) = O(100 \times 1024) = O(102,400) \text{ copies.}$$

The second instance already exceeds the assumed non-adaptive exact-stabilizer scale 10^4 at $n = 100$ (Section 4.2, projection with $c = 1$), so the nullity learner's regime of advantage is small r ; this is consistent with the motivation of [1], where r counts T gates in a Clifford circuit preparing the state, and with the average-case picture of [4], where typical t -doped instances with $t = O(\log n)$ retain enough structure for shallower procedures.

4.5 Workload-level projection

Under a repeated-measurement workload of W independent characterizations of n -qubit stabilizer states, the total measurement budgets scale as

$$M_{\text{ad}} = W C_{\text{ad}} n, \quad M_{\text{na}} = W c_{\text{na}} n^2,$$

where C_{ad} is the constant in the adaptive $\Theta(n)$ bound of [1] and c_{na} the constant in the non-adaptive $\Omega(n^2)$ bound. Neither constant is extracted in the cited texts, so the following is a labeled projection with assumed values $C_{\text{ad}} = 5$ and $c_{\text{na}} = 10$.

Instance $W = 1000, n = 100$:

$$M_{\text{ad}} = 1000 \times 5 \times 100 = 5.0 \times 10^5,$$

$$M_{\text{na}} = 1000 \times 10 \times 100^2 = 1000 \times 10 \times 10^4 = 1.0 \times 10^8.$$

The ratio is

$$\frac{M_{\text{na}}}{M_{\text{ad}}} = \frac{1.0 \times 10^8}{5.0 \times 10^5} = 200.$$

Uncertainty range. The ratio scales as $(c_{\text{na}}/C_{\text{ad}}) \cdot n$. With $C_{\text{ad}} \in [1, 10]$ and $c_{\text{na}} \in [1, 10]$, the ratio $(c_{\text{na}}/C_{\text{ad}}) \cdot 100$ ranges over $[10, 1000]$ at $n = 100$: the minimum is $(1/10) \times 100 = 10$ and the maximum is $(10/1) \times 100 = 1000$. The point estimate 200 sits inside this range.

5. Results

Every number below is computed with shown arithmetic in Section 4 or is explicitly labeled a projection with stated assumptions.

1. **Hypothesis-class size.** $N_n = 2^n \prod_{i=1}^n (2^i + 1)$ and $\log_2 N_n = n^2/2 + 3n/2 + O(1)$ (Section 4.1); at $n = 100$, $\log_2 N_{100} = 5150 + \delta_{100}$ with $\delta_{100} \in (0, 1.4427)$, so $N_{100} \approx 2^{5150}$ up to a factor below 2.72.
1. **Information floor.** Any single-copy learner needs at least $\lceil \log_2 N_n/n \rceil$ copies; at $n = 100$ this is 52 copies (Section 4.2).
1. **Non-adaptive cost.** $\Omega(n^2)$ copies [1]; at $n = 100$, $m_{\text{na}}(100) \geq c \cdot 10^4$ with unspecified constant $c > 0$; the figure 10^4 is a projection under the assumption $c = 1$ (Section 4.2).
1. **Adaptive cost.** $\Theta(n)$ copies [1]; at $n = 100$, between the floor 52 and the projection $C_{\text{ad}} \cdot 100$ with $C_{\text{ad}} \in [1, 10]$ (Section 4.2).

1. **Separation factor.** At $n = 100$, $m_{\text{na}}(100)/m_{\text{ad}}(100) \geq 100/C_{\text{ad}} \in [10, 100]$ under the projection constants $c = 1$, $C_{\text{ad}} \in [1, 10]$ (Section 4.2); asymptotically the separation is quadratic in n up to the constant ratio c/C_{ad} .
1. **Memory tradeoff.** $\Theta(n - k + 1/\varepsilon)$ copies [1]; at $(n, k, \varepsilon) = (100, 10, 0.01)$ the instance value is 190 (Section 4.3), with boundary values $\Theta(200)$ at $k = 0$ and $\Theta(100)$ at $k = 100$, and crossover $\varepsilon^* = 1/n = 0.01$ at $n = 100$.
1. **Nullity extension.** $O(n 2^r)$ copies [1]: $O(3200)$ at $(n, r) = (100, 5)$ and $O(102,400)$ at $(100, 10)$ (Section 4.4).
1. **Workload projection (labeled).** For $W = 1000$ characterizations at $n = 100$ with assumed constants $C_{\text{ad}} = 5$, $c_{\text{na}} = 10$: 5.0×10^5 adaptive versus 1.0×10^8 non-adaptive measurements, ratio 200, uncertainty range $[10, 1000]$ (Section 4.5).

6. Discussion

Limitations. The numerical instances rest on asymptotic bounds whose constants are not extracted in the cited texts; the values $C_{\text{ad}} = 5$ and $c_{\text{na}} = 10$ in Section 4.5 are assumptions, not measurements, and every workload figure is a projection. The adaptive algorithm of [1] chooses each measurement basis as a function of the full transcript, which serializes the experiment: if wall-clock time rather than copy count is the binding constraint, a $\Theta(n^2)$ -copy non-adaptive schedule run fully in parallel may beat a $\Theta(n)$ -copy adaptive one, and the crossover depends on hardware latency ratios we have not measured. The $\Omega(n^2)$ non-adaptive lower bound is worst-case; [4] shows that almost all stabilizer groups with $t = O(\log n)$ are learnable by logarithmic-depth local Clifford circuits, so for typical states the adaptive advantage may be modest. All bounds assume ideal noiseless single-copy Clifford measurements; depolarizing noise effectively caps the achievable stabilizer fidelity and converts exact identification into the tolerant-testing regime at $\varepsilon = O(p)$ for noise rate p , a conversion whose constants we have not derived.

Failure modes and falsification. The central claim — adaptivity closes the gap — would be falsified by either of two observations: (i) a non-adaptive single-copy algorithm achieving $o(n^2)$ copies for stabilizer state learning, or (ii) a lower bound showing that any adaptive single-copy strategy needs $\omega(n)$ copies. Direction (ii) is closed by [1] itself, which proves the $\Theta(n)$ adaptive upper bound. Direction (i) is closed by the same work's statement of the $\Omega(n^2)$ non-adaptive lower bound, but our per-copy bit reconstruction in Section 4.2 is the weakest link in the account: it assumes a fixed basis aligns with an unidentified generator only with n -independent probability. If a fixed basis family extracted $\Theta(n)$ usable bits per copy for every stabilizer state, the quadratic separation would collapse. The group-structured POVM framework of [7], which produces full-rank single-copy estimators, is the natural place to search for such a family; [7] itself notes that its fidelity comparison illustrates the conditioning of the estimate, not a reduction in the copies needed to reach a target error, which supports our reading that group structure alone does not close the gap. Demonstrating or refuting the existence of such a fixed family is an open problem.

Arguing against ourselves. The analogy to adaptive methods in machine learning is interpretive, not a theorem. Reference [3] gains robustness in few-shot meta-learning by applying hypergradient descent to hyperparameter choices; [5] stops active learning when predictions stabilize, reporting for example 364 annotations at F-measure 96.01 for cutoff $K = 99.5$ on the Spamassassin dataset; [8] stabilizes an extreme (greedy) Q-learning objective by restructuring its loss; [9] reports more than a $10\times$ sample-efficiency benefit from model-based closed-loop control over model-free baselines in adaptive brain stimulation. None of these settings involves exact identification of a combinatorial object, so the claim that adaptive advantage is a generic phenomenon rather than a stabilizer-specific one is a hypothesis, supported by structural analogy only. A second objection targets our counting derivation: the step from Lagrangians to states assumes each sign pattern on an ordered generator basis gives a distinct state, which we verified only at $n = 1$ ($N_1 = 6$); an independent verification at $n = 2$ by direct enumeration would strengthen it. A third objection is that the

floor of 52 copies uses $b_{\text{copy}} \leq n$, the raw outcome entropy; a learner restricted to Clifford-basis measurements might face a smaller effective alphabet, which would raise the floor and only strengthen the adaptive regime's optimality claim, but we have not analyzed restricted-alphabet floors.

What would falsify the main result. A single concrete observation suffices: a published non-adaptive single-copy protocol that identifies an arbitrary n -qubit stabilizer state from $o(n^2)$ copies. Our Section 4.2 accounting predicts this is impossible because a fixed schedule wastes all but $O(1)$ bits per copy on worst-case states; a counterexample protocol would show the per-copy aligned-bit assumption fails, and the separation factor of 10–100 at $n = 100$ (itself a projection) would need to be recomputed from scratch.

Limitations list. (1) Unspecified constants in the Θ - and Ω -bounds of [1]: every constant-dependent figure is a projection; the separation factor is sensitive to C_{ad} and c at the level of the full ranges $[10, 100]$ (Section 4.2) and $[10, 1000]$ (Section 4.5). (2) Idealized noiseless measurements: noise converts the problem to tolerant testing at $\varepsilon = O(p)$; the workload projections do not include the $1/\varepsilon$ overhead. (3) Serialization of adaptive schedules: wall-clock comparisons can reverse the copy-count advantage; no latency data is assumed. (4) Worst-case lower bound: average-case instances are easier [4], so the projected advantage may not materialize for typical device outputs. (5) The workload model of Claim D is ours, not a result of any cited work. (6) The Lagrangian-to-state counting is verified only at $n = 1$.

7. Practical Implications and Engineering Pathways

(a) Who uses the result and what changes. Teams operating quantum processors — calibration engineers running state-preparation verification, benchmarking groups executing randomized benchmarking follow-ups, and error-correction teams validating stabilizer code states — would use this result to size measurement budgets for state identification. If the result holds, the design decision that changes is the hardware requirement: fast classical feedback between measurement shots (sub-microsecond-to-millisecond classical control loops that select the next measurement basis from the transcript) replaces coherent multi-copy access or quantum memory as the resource that buys the $\Theta(n)$ copy count. A laboratory deciding whether to invest in multi-copy routing hardware or merely in faster classical control can, under this result, choose the latter for stabilizer identification tasks. The tolerant tester with tradeoff $\Theta(n - k + 1/\varepsilon)$ changes certification protocol design: at $(n, k, \varepsilon) = (100, 10, 0.01)$ the budget is $\Theta(190)$ copies (Section 4.3), and each stored qubit saves one copy, so a team holding even small coherent registers should budget $n - k + 1/\varepsilon$ rather than $n + 1/\varepsilon$.

(b) Current readiness. This work is purely analytical and has not been peer reviewed; under the NASA Technology Readiness Level definitions it sits at TRL 1–2: basic principles observed and reported (the $\Theta(n)$ adaptive learner of [1]) with a conceptual reconstruction and numerical instances supplied here. No experimental demonstration of the adaptive single-copy protocol at scale is reported in the cited texts, and analytical work never exceeds TRL 3; we claim no hardware readiness.

(c) Smallest next experiment. The minimal de-risking test is a $n = 3$ or $n = 4$ qubit demonstration on any platform supporting single-qubit Pauli-basis measurements with programmable basis selection: prepare a known stabilizer state, run the adaptive protocol of [1] with the transcript-dependent basis choice, and measure the number of copies to exact identification, repeated over all stabilizer states at that size ($N_3 = 2^3 \prod_{i=1}^3 (2^i + 1) = 8 \cdot 3 \cdot 5 \cdot 9 = 1080$ states, computed from the formula of Section 4.1). The outcome that would refute the practical claim is a measured copy count scaling visibly above $C_{\text{ad}} n$ with a large C_{ad} , or a failure of the adaptive subroutine to fix one generator per round at constant probability. A companion non-adaptive run on the same states measures the actual constant c , replacing the assumed $c = 1$ projection of Section 4.2 with data.

(d) Binding constraints. The first constraint to bind is classical control latency: the adaptive schedule is inherently serial, so per-shot feedback latency multiplies the wall-clock time by the copy count; whether this binds before or after copy-count savings materialize depends on hardware latency ratios we have not measured (open question). The second is measurement basis programmability: the protocol requires arbitrary Clifford-basis measurements selected shot-by-shot, which on some platforms requires recalibrating analysis pulses per basis; the cost of that reconfiguration is an open question. The third is noise: real devices deliver noisy copies, converting exact identification into tolerant testing with the $1/\varepsilon$ overhead of Section 4.3; the noise rate p of the target platform sets $\varepsilon \approx O(p)$ and thus the budget, but the constant is unmeasured here (open question). Data and compute constraints are mild — the classical processing is polynomial-time per [1] — but the transcript-dependent classical computation must complete within the inter-shot window, an open question.

(e) Order-of-magnitude budget bounds (projections). Using only the derived arithmetic of Section 4.5 and its stated assumptions ($C_{\text{ad}} = 5$, $c_{\text{na}} = 10$, $n = 100$, $W = 1000$): a characterization campaign consumes 5.0×10^5 adaptive shots versus 1.0×10^8 non-adaptive shots, a factor of 200, with the uncertainty range $[10, 1000]$ from the constant ranges. These are projections, not measurements; they scale linearly in W and as n versus n^2 in system size.

(f) Risks and misuse paths. No dual-use risk is apparent: the result accelerates state characterization and certification of quantum devices. One misuse-adjacent concern is the opposite of a risk — overclaiming: a team reading the $\Theta(n)$ bound as a wall-clock guarantee would under-provision experiment time, since the serial adaptive schedule's wall-clock cost is m_{ad} times the per-shot latency, which Section 6 flags as the binding constraint. Any budget document built on this analysis should state the copy-count saving and the serialization penalty side by side.

8. Conclusion

Adaptivity alone closes the stabilizer learning gap: an n -qubit stabilizer state is learnable from $\Theta(n)$ adaptive single-copy Clifford measurements, matching the two-copy Bell-sampling optimum without multi-copy operations or quantum memory [1]. Our analytical account makes the separation concrete: the hypothesis class has $N_{100} \approx 2^{5150}$ candidates at $n = 100$, any single-copy learner needs at least 52 copies, the non-adaptive requirement is $\Omega(n^2)$ against the adaptive $\Theta(n)$, the memory tradeoff evaluates to $\Theta(190)$ copies at $(n, k, \varepsilon) = (100, 10, 0.01)$, and the nullity extension costs $O(3200)$ copies at $(n, r) = (100, 5)$. Under stated constant assumptions, a 1000-characterization workload at $n = 100$ needs 5.0×10^5 adaptive versus 1.0×10^8 non-adaptive measurements. The design implication is that fast classical feedback between shots, not coherent multi-copy access, is the resource that buys the quadratic saving; the open questions are the actual constants, the wall-clock penalty of serialization, and robustness to noise. The smallest decisive test is a small- n demonstration measuring the adaptive copy count against the floor, which would either confirm the projected factor-of-10–100 saving or expose a per-copy information assumption that fails in practice.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.17.5-abstract-lines). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, zai-org/glm-5.3, openai/gpt-oss-120b, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.

- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 11 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

[1] TITLE: arXiv Query: `searchquery=&idlist=2610.02031&start=0&max_results=1` [2] Adaptivity is all you need: Optimal stabilizer learning using just single-copy measurements. arXiv:2610.02031v1. <https://arxiv.org/abs/2610.02031v1> [3] Alpha MAML: Adaptive Model-Agnostic Meta-Learning. arXiv:1905.07435v1. <https://arxiv.org/abs/1905.07435v1> [4] Single-copy stabilizer learning: average case and worst case. arXiv:2604.24099v1. <https://arxiv.org/abs/2604.24099v1> [5] A Method for Stopping Active Learning Based on Stabilizing Predictions and the Need for User-Adjustable Stopping. arXiv:1409.5165v1. <https://arxiv.org/abs/1409.5165v1> [6] Lower Bounds for Learning Quantum States with Single-Copy Measurements. arXiv:2207.14438v3. <https://arxiv.org/abs/2207.14438v3> [7] Quantum Algebraic Diversity: Single-Copy Density Matrix Estimation via Group-Structured Measurements. arXiv:2604.03725v3. <https://arxiv.org/abs/2604.03725v3> [8] Stabilizing Extreme Q-learning by Maclaurin Expansion. arXiv:2406.04896v2. <https://arxiv.org/abs/2406.04896v2> [9] Coprocessor Actor Critic: A Model-Based Reinforcement Learning Approach For Adaptive Brain Stimulation. arXiv:2406.06714v2. <https://arxiv.org/abs/2406.06714v2> [10] QNFO: Adaptivity Closes the Stabilizer Learning Gap: An Analytical Account of Optimal Single-Copy Stabilizer State Learning