

Optimal Stabilizer Testing and Learning with Limited Quantum Memory

Srinivasan Arunachalam Louis Schatzki

1 Background

Learning and testing properties of unknown n -qubit quantum systems are fundamental problems in quantum information science. However, a major bottleneck is the difficulty of realizing coherent quantum memory. This motivates memory-limited streaming models for quantum data, where the algorithm sees multiple rounds of data qubits but is only allowed to preserve a small coherent workspace between measurements. In this spirit, a flurry of works over the past few years have studied quantum algorithms under streaming and memory constraints [2–14, 16, 18–21, 23–25, 27]. One of the main messages from this bulk of work is that restrictions to single-copy measurements often lead to large, potentially exponential, separations in sample complexity, suggesting that memory is not merely an implementation detail which can be removed for free, but rather a computational resource that can dramatically change the complexity of the problem. Although the tasks considered in these papers are very natural, the field still lacks a good understanding of the learning and testing complexity of the “simplest” family of states in a memory-limited setting, namely n -qubit *stabilizer states*. For single-copy measurements, recent work showed that $O(n)$ samples suffice and that $\Omega(\sqrt{n})$ samples are necessary [19]. Until our work, this gap remained open.

A seminal result of Gross, Nezami, and Walter [17] showed that stabilizer states can be tested using only 6 copies, in contrast to the $\Theta(n)$ sample complexity of learning a stabilizer state [1, 26]. Both these learning and testing results crucially rely upon *Bell sampling* and two-copy measurements. However, it is desirable to store as few coherent qubits as possible, motivating the central questions of our work:

1. Can we obtain constant-copy testers with limited memory?
2. Is stabilizer testing always easier than learning even with limited memory?


To tackle these questions, we consider a smooth interpolation between single-copy and two-copy measurements. In each round a fresh copy of an input state $|\psi\rangle$ is provided to the learning/testing algorithm. On top of this input, the algorithm is allowed to maintain k qubits of coherent memory, interpolating between single-copy measurements ($k = 0$) and two-copy measurements ($k = n$).

Our main result is to settle both of these questions. We prove that with k qubits of memory, the sample complexity of testing and learning stabilizer states is $\Theta(n - k)$ and $O(n^2/k)$ respectively, the latter being tight for non-adaptive learners. Hence, even with $0.99n$ qubits of coherent memory, $\Theta(n)$ copies are necessary and sufficient to test stabilizer states. Interestingly, with $0.99n$ qubits of memory, $\Theta(n)$ copies are also necessary and sufficient to *learn* an unknown stabilizer state. More generally, this exhibits a curious difference in scaling with limited memory. Take $k = 0$, i.e., single-copy measurements. Then, testing requires $\Theta(n)$ samples but the best known learning algorithm requires $O(n^2)$ samples. As k grows towards n , the sample complexity of learning decreases at a much faster rate than that of testing and, for $k = c \cdot n$ with $c < 1$ a constant, both become $\Theta(n)$. See Figure 1 for an illustration.


2 Testing stabilizer states

Theorem 2.1. (*Optimal testing bounds*) *Let $k \geq 0, \varepsilon > 0$. There is an adaptive protocol that uses k qubits of memory and $O((n - k)/\varepsilon)$ copies of an unknown $|\psi\rangle$ to distinguish between $\mathcal{F}_{\text{Stab}}(|\psi\rangle) = 1$ vs. $\mathcal{F}_{\text{Stab}}(|\psi\rangle) \leq 1 - \varepsilon$. Also every such tester needs $\Omega(n - k)$ copies.*

Upper bound. The central difficulty is that the tester has only k qubits of coherent memory, so it cannot directly Bell-sample all n qubits. Instead, partial Bell sampling reveals *prefix* of a Pauli label and leaves an *unknown suffix*. Our key contribution will then be a single-copy algorithm that takes a prefix



(a) With memory, the sample complexity for learning decreases at a faster rate than testing.



(b) Ratio between learning and testing sample complexity as a function of memory qubits. For $k = cn$, with $0 < c < 1$, $\Theta(n)$ samples is necessary and sufficient for both learning and testing.

Figure 1: Sample complexity of learning and testing stabilizer states with limited quantum memory.

and checks for such a suffix. We do so via a novel connection to the *hidden shift problem*. [15, 22, 28–31]. Explicitly, Bell difference sampling on k qubits and computational difference sampling on the remaining $n - k$ qubit samples a marginal of the full Bell difference distribution. Say that the prefix $a \in \mathbb{F}_2^{2k}$ and $r \in \mathbb{F}_2^{n-k}$ is observed. For a stabilizer state, there must exist some $t^* \in \mathbb{F}_2^{n-k}$ such that $P_{a \times r, a_z t^*}$ is a stabilizer. For the observed prefix (a, r) , we define two Pauli families

$$F_{a,r}(t) := P_{a \times r, a_z t}, \quad G(t) := \mathbb{I}_k \otimes Z(t).$$

If $F_{a,r}(t^*)$ is a stabilizer, then $F_{a,r}(t^*) |\psi\rangle = \lambda |\psi\rangle$ for some $\lambda \in \{\pm 1\}$. Now compare any other completion t to this stabilizer prefix and the corresponding Pauli string $F_{a,r}(t)$. Their product is, up to a phase, a pure Z operator on the missing m qubits, $F_{a,r}(t) = i^{(t-t^*) \cdot r} G(t + t^*) F_{a,r}(t^*)$. Applying this identity to $|\psi\rangle$ gives $F_{a,r}(t) |\psi\rangle = \lambda i^{(t-t^*) \cdot r} G(t + t^*) |\psi\rangle$. Thus, if a stabilizer completion t^* exists, the indexed family $t \mapsto F_{a,r}(t) |\psi\rangle$ has the same action on $|\psi\rangle$ as the reference family $t \mapsto G(t) |\psi\rangle$, but shifted by t^* . This precisely is the *hidden-shift* structure that we exploit in our testing protocol, which is inspired by the usual Fourier sampling approach for the hidden subgroup problem. This step can be realized with *single-copy measurements* and returns linear constraints involving t^* . In particular, the data lies entirely in the graph of the affine function $s \mapsto t^* \cdot s \oplus \text{sign}(\lambda)$. But, if $|\psi\rangle$ is not a stabilizer state then no such t^* need exist and, if not, then the samples do not lie in the graph of an affine function with high probability. Hence, the tester obtains $O(n - k)$ samples and checks for affine consistency.

Lower bound. In order to obtain our lower bounds, we consider the ensemble of random *degree-2 phase states*, a subset of stabilizer states:

$$|\psi_A\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \mathbb{F}_2^n} (-1)^{x^\top A x} |x\rangle \quad (1)$$

and show that distinguishing these states from the maximally mixed state requires $\Omega(n - k)$ samples and that distinguishing a Haar random state from the maximally mixed state requires $2^{\Omega(n-k)}$ samples. Our lower bound is proven using the so-called likelihood ratios, which have been used often to prove quantum testing/learning lower bounds with memory constraints [6, 10, 12] and hinge upon proving that

$$\mathbb{E}_{\vec{x} \sim \mathcal{P}_{mm}} [|L(\vec{x}) - 1|] = o(1), \quad L(\vec{x}) = \frac{\mathbb{E}_A [\text{Tr}(E_{\vec{x}} \psi_A^{\otimes t})]}{\text{Tr}(E_{\vec{x}} I / 2^{nt})}, \quad (2)$$

when $t = o(n - k)$. To bound the (average) likelihood ratios, the first step is to understand $\mathbb{E}_A[\psi_A^{\otimes t}]$, which we show is close in trace distance to the subnormalized state

$$\sigma_P := \frac{1}{2^{nt}} \sum_{O \in \mathcal{O}_t} R(O) . \quad (3)$$

We then split $\mathcal{O}_t$ into two parts: $\mathcal{O}_{t-1}$ (all matrices that leave the first coordinate fixed) and its complement, $\mathcal{M}_t := \mathcal{O}_t \setminus \mathcal{O}_{t-1}$. The contribution of $\mathcal{O}_{t-1}$ will be handled recursively, so the main task is to prove that the $\mathcal{M}_t$ -contribution has negligible average bias. By triangle inequality, it suffices to bound

$$\mathbb{E}_{\vec{x} \sim \mathcal{P}_{mm}} \left[\frac{1}{\text{Tr}[E_{\vec{x}}]} \left| \sum_{O \in \mathcal{M}_t} \text{Tr}[R(O)E_{\vec{x}}] \right| \right] = \frac{1}{2^{nt}} \sum_{\vec{x}} \left| \sum_{O \in \mathcal{M}_t} \text{Tr}[R(O)E_{\vec{x}}] \right| . \quad (4)$$

We bound this term by considering branches of the learning tree conditioned upon the first round outcome x_1 . The first round induces a map $V_{x_1} : \mathcal{H} \rightarrow \mathcal{M}$ from the first input copy into the k -qubit memory space. The bias of this entire branch of the learning tree is compressed into the operator $\sum_{O \in \mathcal{M}_t} R(O)$ passed through the k -qubit memory via V_{x_1} , which we bound via combinatorics of binary orthogonal matrices. In particular, the combinatorial parameter $\text{rank}(\mathbb{I} + O)$, which is the codimension of the subspace of vectors fixed by O , appears when computing the kernels of certain linear maps resulting from partially contracting operators in $\mathcal{M}_t$ through the memory via V_{x_1} .

3 Learning stabilizer states

Theorem 3.1. (*Optimal learning bounds*) *Let $k \geq 1$. There is a non-adaptive protocol that uses k qubits of memory and $O(n^2/k)$ copies of an unknown $|\psi\rangle \in \text{Stab}$ to learn $|\psi\rangle$. Also, every non-adaptive protocol needs $\Omega(n^2/k)$ copies.*

Upper bound. Say that the state is in *full-support form*, by which we mean that the unsigned stabilizer group is $M = \{(u, Bu) : u \in \mathbb{F}_2^n\}$ for a symmetric matrix B . Let S be a block of qubits of size at most k . Bell-measuring this block and computational basis-measuring the remaining qubits yields samples of the form $(a, \Pi_S Ba)$, where $a \in \mathbb{F}_2^n$ is uniform and Π_S indicates restricting to the rows corresponding to S . Taking $O(n)$ samples suffices to recover this block of B . Partitioning the full n -qubit state into n/k blocks of size roughly k then suffices to learn the state using $O(n^2/k)$ copies. Now, the state may not actually be in full-support form, but a random Clifford unitary puts it into such a form with constant probability. Hence, repeating this procedure with random Cliffords several times suffices to learn stabilizer group M_ψ non-adaptively. To recover signs, we use separate random stabilizer-basis measurements. For each random Clifford basis, the measurement reveals eigenvalues of a known random commuting Pauli subspace L_s . This determines the signs for the Pauli strings in the intersection $L_s \cap M_\psi$, which is nontrivial with constant probability. With $O(n)$ random Clifford basis measurements, the intersections span M_ψ with high probability, which then fully determines the signs of all Paulis in M_ψ . This sign-recovery stage costs only $O(n)$ copies, so the total remains $O(n^2/k)$.

Lower Bound. It again suffices to consider the smaller ensemble of real degree-2 phase states

$$|\psi_A\rangle = 2^{-n/2} \sum_{x \in \mathbb{F}_2^n} (-1)^{x^T A x} |x\rangle ,$$

where A is a uniformly random upper-triangular matrix over $\mathbb{F}_2$. The parameter A contains $\Theta(n^2)$ bits. A one-copy accessible-information calculation shows that any measurement on a single fresh copy reveals only $O(1)$ bits about A . If the protocol also carries a k -qubit coherent memory, then one round can increase the information about A by at most $O(k+1)$ bits (this is where non-adaptivity is used). Thus, after T copies, the final transcript contains at most $O(T(k+1))$ bits of information about A . On the other hand, any learner that identifies the state with constant success probability must identify A with constant success probability, and hence must obtain $\Omega(n^2)$ bits of information by Fano's inequality. Therefore

$$T(k+1) = \Omega(n^2) \implies T \geq \Omega(n^2/k) .$$

References

- [1] Scott Aaronson and Daniel Gottesman. Identifying stabilizer states, 2009. <https://pirsa.org/08080052>. 1
- [2] Dorit Aharonov, Jordan Cotler, and Xiao-Liang Qi. Quantum algorithmic measurement. *Nature communications*, 13(1):887, 2022. 1
- [3] Srinivasan Arunachalam, Vojtech Havlicek, and Louis Schatzki. On the role of entanglement and statistics in learning. *Advances in Neural Information Processing Systems*, 36:55064–55076, 2023.
- [4] Srinivasan Arunachalam and Louis Schatzki. Generalized inner product estimation with limited quantum communication. In *42nd International Symposium on Theoretical Aspects of Computer Science (STACS 2025)*, pages 11–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2025.
- [5] Jacob Beckey, Luke Coffman, Ariel Shlosberg, Louis Schatzki, and Felix Leditzky. Product testing with single-copy measurements. *arXiv preprint arXiv:2510.07820*, 2025.
- [6] Sebastien Bubeck, Sitan Chen, and Jerry Li. Entanglement is necessary for optimal quantum property testing. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 692–703. IEEE, 2020. 2
- [7] Matthias C Caro. Learning quantum processes and hamiltonians via the pauli transfer matrix. *ACM Transactions on Quantum Computing*, 5(2):1–53, 2024.
- [8] Senrui Chen, Changhun Oh, Sisi Zhou, Hsin-Yuan Huang, and Liang Jiang. Tight bounds on pauli channel learning without entanglement. *Physical Review Letters*, 132(18):180805, 2024.
- [9] Senrui Chen, Sisi Zhou, Alireza Seif, and Liang Jiang. Quantum advantages for pauli channel estimation. *Physical Review A*, 105(3):032435, 2022.
- [10] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. Exponential separations between learning with and without quantum memory. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 574–585. IEEE, 2022. 2
- [11] Sitan Chen and Weiyuan Gong. Efficient pauli channel estimation with logarithmic quantum memory. *PRX Quantum*, 6(2):020323, 2025.
- [12] Sitan Chen, Weiyuan Gong, and Qi Ye. Optimal tradeoffs for estimating Pauli observables. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1086–1105, 2024. 2
- [13] Sitan Chen, Jerry Li, Brice Huang, and Allen Liu. Tight bounds for quantum state certification with incoherent measurements. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1205–1213. IEEE, 2022.
- [14] Sitan Chen, Jerry Li, and Allen Liu. An optimal tradeoff between entanglement and copy complexity for state tomography. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1331–1342, 2024. 1
- [15] Andrew M Childs and Wim Van Dam. Quantum algorithms for algebraic problems. *Reviews of Modern Physics*, 82(1):1–52, 2010. 2
- [16] Weiyuan Gong, Jonas Haferkamp, Qi Ye, and Zhihan Zhang. On the sample complexity of purity and inner product estimation. *arXiv preprint arXiv:2410.12712*, 2024. 1

- [17] David Gross, Sepehr Nezami, and Michael Walter. Schur–Weyl duality for the Clifford group with applications: Property testing, a robust Hudson theorem, and de Finetti representations. *Communications in Mathematical Physics*, 385(3):1325–1393, 2021. [1](#)
- [18] Marcel Hinsche, Zongbo Bao, Philippe van Dordrecht, Jens Eisert, Jop Briët, and Jonas Helsen. Clifford testing: Algorithms and lower bounds. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, pages 869–873, 2026. [1](#)
- [19] Marcel Hinsche and Jonas Helsen. Single-copy stabilizer testing. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 439–450, 2025. [1](#)
- [20] Hsin-Yuan Huang, Michael Broughton, Jordan Cotler, Sitan Chen, Jerry Li, Masoud Mohseni, Hartmut Neven, Ryan Babbush, Richard Kueng, John Preskill, et al. Quantum advantage in learning from experiments. *Science*, 376(6598):1182–1186, 2022.
- [21] Robbie King, David Gosset, Robin Kothari, and Ryan Babbush. Triply efficient shadow tomography. *PRX Quantum*, 6:010336, Feb 2025. [1](#)
- [22] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM Journal on Computing*, 35(1):170–188, 2005. [2](#)
- [23] Qipeng Liu, Ran Raz, and Wei Zhan. Memory-sample lower bounds for learning with classical-quantum hybrid memory. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1097–1110, 2023. [1](#)
- [24] Yuhan Liu and Jayadev Acharya. The role of randomness in quantum state certification with unentangled measurements. In *The Thirty Seventh Annual Conference on Learning Theory*, pages 3523–3555. PMLR, 2024.
- [25] Angus Lowe and Ashwin Nayak. Lower bounds for learning quantum states with single-copy measurements. *ACM Transactions on Computation Theory*, 17(1):1–42, 2025. [1](#)
- [26] Ashley Montanaro. Learning stabilizer states by Bell sampling. *arXiv:1707.04012*, 2017. [1](#)
- [27] Aadil Oufkir. Sample-optimal quantum process tomography with non-adaptive incoherent measurements. In *2023 IEEE International Symposium on Information Theory (ISIT)*, pages 1919–1924. IEEE, 2023. [1](#)
- [28] Oded Regev. Quantum computation and lattice problems. *SIAM Journal on Computing*, 33(3):738–760, 2004. [2](#)
- [29] Martin Rötteler. Quantum algorithms for highly non-linear boolean functions. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete algorithms*, pages 448–457. SIAM, 2010.
- [30] Daniel R Simon. On the power of quantum computation. *SIAM journal on computing*, 26(5):1474–1483, 1997.
- [31] Wim Van Dam, Sean Hallgren, and Lawrence Ip. Quantum algorithms for some hidden shift problems. *SIAM Journal on Computing*, 36(3):763–778, 2006. [2](#)