

Cloning is as Hard as Learning for Stabilizer States

(arXiv:2604.15269)

QCOW-II

Nikhil Bansal

April 23, 2026



My Collaborators



Matthias C. Caro
(University of Warwick)



Gaurav Mahajan
(Yale University)

Outline

1. Introduction/Motivation
2. Main Result
3. Proof Overview
4. Open Questions

Motivation

A single quantum cannot be cloned

No Cloning Theorem (Informal)

There doesn't exist a unitary map U , such that

$$U(|\psi\rangle \otimes |0\rangle) = |\psi\rangle \otimes |\psi\rangle \quad (1)$$

holds simultaneously for non-orthogonal states $|\psi\rangle$.

William K Wootters and Wojciech H Zurek (1982). "A single quantum cannot be cloned". In: *Nature* 299.5886, pp. 802–803. doi: 10.1038/299802a0.

D. Dieks (1982). "Communication by EPR devices". In: *Physics Letters A* 92.6, pp. 271–272. issn: 0375-9601. doi: [https://doi.org/10.1016/0375-9601\(82\)90084-6](https://doi.org/10.1016/0375-9601(82)90084-6).

Approximate Cloning




Figure: Approximate Cloning map Λ for d -dimensional states $|\psi\rangle \in \mathcal{H}$

- Figure of Merit:

$$\mathcal{F}(\Lambda, t, t+m) = \inf_{|\psi\rangle \in \mathcal{H}} \text{Tr} \left(|\psi\rangle\langle\psi|^{\otimes t+m} \Lambda(|\psi\rangle\langle\psi|^{\otimes t}) \right). \quad (2)$$

Optimal approximate cloner for general pure states have been fully characterised!

Optimal Cloning of Pure States

R.F. Werner^{1,2}

Abstract. We construct the unique optimal quantum device for turning a finite number of d -level quantum systems in the same unknown pure state σ into M systems of the same kind, in an approximation of the M -fold tensor product of the state σ .

Here, we have what one may call a **No Approximate Cloning theorem**.

No Approximate Cloning Theorem

No Approximate Cloning Theorem (Informal)

The optimal achievable fidelity for $t \rightarrow t + m$ approximate cloning of d -dimensional pure states is

$$\mathcal{F}^*(t \rightarrow t + m) = \sup_{\Lambda} \mathcal{F}(\Lambda) = \frac{d[t]}{d[t + m]}, \quad (3)$$

where $d[t] = \binom{t+d-1}{t}$ is the dimension of symmetric subspace of $\mathcal{H}^{\otimes t}$.

Cloning is as Hard as Learning for Pure States

- In sample complexity sense, to get an extra clone with fidelity $1 - \epsilon$ requires $\Theta\left(\frac{d}{\epsilon}\right)$ copies.
- This matches the sample complexity of learning general pure states in fidelity.

Dagmar Bruß and Chiara Macchiavello (Mar. 1999). "Optimal state estimation for d-dimensional quantum systems". In: *Physics Letters A* 253.5–6, pp. 249–251. ISSN: 0375-9601. DOI: 10.1016/S0375-9601(99)00099-7. URL: [http://dx.doi.org/10.1016/S0375-9601\(99\)00099-7](http://dx.doi.org/10.1016/S0375-9601(99)00099-7).

Henry Yuen (Jan. 2023). "An Improved Sample Complexity Lower Bound for (Fidelity) Quantum State Tomography". In: *Quantum* 7, p. 890. ISSN: 2521-327X. DOI: 10.22331/q-2023-01-03-890. URL: <http://dx.doi.org/10.22331/q-2023-01-03-890>.

Is Cloning as Hard as Learning?




Figure: Cloning structured states

Q. Can approximate cloning be easier than learning for structured classes of states?

Main Result

Cloning is as Hard as Learning for Stabilizer States

Note: For a set of states $\mathcal{S}$, we care about the following figure of merit for cloning:

$$\epsilon_{\text{Cl}}^*(\mathcal{S}, t, t+1) = \inf_{\Lambda} \sup_{\rho \in \mathcal{S}} d_{\text{TD}}(\Lambda(\rho^{\otimes t}), \rho^{\otimes t+1}) . \quad (4)$$

No Approximate Cloning for Stabilizer States (Informal)

Approximate cloning for the class of pure n -qubit stabilizer states from $< \lfloor n/4 \rfloor$ copies incurs at least a constant error of 0.14.

- Cloning stabilizer states **with small constant error requires $\Omega(n)$ copies.**
- This matches sample complexity of learning n -qubit stabilizer states: **$\Theta(n)$.**

Proof Overview

Bell Sample Amplification

Task: Given $2t$ copies of a stabilizer state, generate $t + 1$ Bell samples.

- The $t + 1$ Bell samples should look like as if drawn i.i.d. from Bell measurement distribution of the state.
- Given an n -qubit stabilizer state, the Bell measurement outcome distribution is a uniform distribution over an n -dimensional affine subspace of $\mathbb{Z}_2^{2n}$.
- One approach is to measure the copies of the state using Bell measurements and then amplify the t Bell samples:

Bell Sample Amplification!

Sample Amplification

Suppose we have a distribution family $\mathcal{D}$.

Task: Given t i.i.d. samples drawn from some unknown $D \in \mathcal{D}$, amplify it to $t + m$ samples.

- Figure of Merit:

$$\epsilon_{SA}^*(\mathcal{D}, t, t + m) = \min_{\Lambda} \max_{D \in \mathcal{D}} d_{TV} (D^{\otimes t} \circ \Lambda^{-1}, D^{\otimes t+m}) \quad (5)$$

$$\epsilon_{CI}^*(\mathcal{S}, t, t + m) = \inf_{\Lambda} \sup_{\rho \in \mathcal{S}} d_{TD} (\Lambda(\rho^{\otimes t}), \rho^{\otimes t+m}) . \quad (6)$$

- **Example:** For the family of **unstructured discrete distributions with bounded support k** , $t = \Theta(\sqrt{k})$ is the sample complexity for amplifying by one extra sample with small constant error.
- Learning with small constant error in TV distance requires $\Theta(k)$ samples.

Brian Axelrod, Shivam Garg, Vatsal Sharan, and Gregory Valiant (2020). "Sample amplification: increasing dataset size even when learning is impossible". In: *Proceedings of the 37th International Conference on Machine Learning*. ICML'20. JMLR.org. URL: <https://proceedings.mlr.press/v119/axelrod20a.html>.

Brian Axelrod, Shivam Garg, Yanjun Han, Vatsal Sharan, and Gregory Valiant (2024). "On the statistical complexity of sample amplification". In: *The Annals of Statistics* 52.6, pp. 2767–2790. doi: 10.1214/24-AOS2444.

Sample Amplification for structured distributions

Consider two classes of distributions with linear structure:

1. $\mathcal{D}_1$: Uniform distribution over n -dimensional subspace of $\mathbb{Z}_2^{2n}$.
2. $\mathcal{D}_2$: $(\mathcal{U}, f_a)$ which have samples $(x, f_a(x))$ where x is drawn uniformly from $\mathbb{Z}_2^n$ and f_a is parity function:

$$f_a(x) = a \cdot x \pmod{2}. \quad (7)$$

Learning in each case is **collecting sufficient linearly independent samples and then solving a resulting linear system!**

Sample Amplification is as Hard as Learning for Linear Structure

- In both cases, n samples are linearly independent w.p. ≥ 0.28
- Roughly: n samples have much more information about the subspace than $n - 1$ samples.
- And, sample Amplification of $n - 1$ samples can't create this new information.

No Sample Amplification Theorem (Informal)

For the distributions $\mathcal{D}_1$ and $\mathcal{D}_2$, sample amplification from $t \rightarrow t + 1$ incurs an error of at least

$$\epsilon_{SA}^*(\mathcal{D}_i, t, t + 1) \geq 0.14, \quad \text{for } t \leq n - 1. \quad (8)$$

Proof Idea: n samples have more information about the subspace than $n - 1$ samples. This extra information allows one to distinguish n true samples from n amplified ones.

Bell Sample Amplification

- Bell Sample amplification is just the sample amplification of $\mathcal{D}_1$.
- Bell Sample amplification with small constant error is as hard as learning!

Cloning via Bell Sampling

1. Given copies of a stabilizer state, measure in Bell Basis.
 2. On the measured data, do Bell sample amplification.
 3. Based on the samples at hand, prepare a stabilizer state.
- **Cloning via Bell Sampling** is as hard as learning because Bell sample amplification is as hard as learning.

More General Cloners?

But there can be more general cloners, potentially using global measurements.

- To get around this, we consider the following mixed $(2n)$ -qubit states:

$$\sigma_L = \frac{1}{2^n} \sum_{y \in L^\perp} |\Phi_y\rangle\langle\Phi_y|, \quad (9)$$

where $|\Phi_y\rangle$ are $2n$ -qubit Bell states, L is an n -dimensional subspace of $\mathbb{Z}_2^{2n}$ and $L^\perp$ is its dual subspace defined as,

$$L^\perp = \{y \in \mathbb{Z}_2^{2n} \mid (-1)^{[x,y]} = 1, \forall x \in L\} \quad (10)$$

where $[x, y]$ is the symplectic inner-product.

Optimality of Bell Sampling

- σ_L have a stabilizer group given by Pauli matrices labelled by L as $V_x^{\otimes 2}$ with $V_x = \bigotimes_{i=1}^n Z^{a_i} X^{b_i}$ where $x = (a, b) \in L$.
- These states are diagonal in Bell basis. Thus, **Bell Sampling is optimal to extract information**.
- Thus **Cloning via Bell Sampling** is optimal here.

No Approximate Cloning for $\{\sigma_L\}_{L \leq \mathbb{Z}_2^{2n}, |L|=2^n}$

For the set of states $\mathcal{S}_n = \{\sigma_L\}_{L \leq \mathbb{Z}_2^{2n}, |L|=2^n}$, the $t \rightarrow t+1$ cloning incurs an error of at least

$$\epsilon_{\text{Cl}}^*(\mathcal{S}_n, t, t+1) \geq 0.14, \quad \text{for } t \leq n-1. \quad (11)$$

- Cloning these states is as hard as learning.

Lifting to Pure States: Random Purification




Figure: Random Purification

Ewin Tang, John Wright, and Mark Zhandry (2025). *Conjugate queries can help*. arXiv: 2510.07622 [quant-ph]. URL: <https://arxiv.org/abs/2510.07622>.

Michael Walter and Freek Witteveen (2025). *A random purification channel for arbitrary symmetries with applications to fermions and bosons*. arXiv: 2512.15690 [quant-ph]. URL: <https://arxiv.org/abs/2512.15690>.

Lifting to Pure States: Random Purification

Random Purification (Informal)

Let $t, r \in \mathbb{N}$. There is a quantum channel $C^{(t)}$ such that for all mixed states $\rho \in \mathbb{C}^{d \times d}$ of rank at most r ,

$$C^{(t)}(\rho^{\otimes t}) = \mathbf{E}_{|\rho_U\rangle} |\rho_U\rangle\langle\rho_U|^{\otimes t}, \quad (12)$$

where the expectation is over random purifications $|\rho_U\rangle \in \mathbb{C}^d \otimes \mathbb{C}^r$ of ρ .

- $C^{(t)}$ takes t i.i.d. copies of a mixed state to t i.i.d. copies of a random purification.

Ewin Tang, John Wright, and Mark Zhandry (2025). *Conjugate queries can help*. arXiv: 2510.07622 [quant-ph]. URL: <https://arxiv.org/abs/2510.07622>.

Michael Walter and Freek Witteveen (2025). *A random purification channel for arbitrary symmetries with applications to fermions and bosons*. arXiv: 2512.15690 [quant-ph]. URL: <https://arxiv.org/abs/2512.15690>.

Random purification of σ_L states

Structured Random Purification for Stabilizer States

Let $t \in \mathbb{N}$. There is a quantum channel $C^{(t)}$ such that for all states of the form $\sigma_L = \frac{1}{2^n} \sum_{y \in L^\perp} |\Phi_y\rangle\langle\Phi_y|$ with $L \leq \mathbb{Z}_2^{2n}$,

$$C^{(t)}(\sigma_L^{\otimes t}) = \mathbf{E}_{x \in \mathbb{Z}_2^{2n}} |\sigma_{L,x}\rangle\langle\sigma_{L,x}|^{\otimes t}, \quad (13)$$

where $|\sigma_{L,x}\rangle = \frac{1}{\sqrt{2^n}} \sum_{y \in L^\perp} |\Phi_y\rangle \otimes V_x^{\otimes 2} |\Phi_y\rangle$.

- Moreover, we show that these random purifications $|\sigma_{L,x}\rangle$ are all **pure $(4n)$ -qubit stabilizer states**.

Lifting the Cloning Lower Bound to pure states

Cloning σ_L via Random Purification

1. Given t copies of σ_L , obtain t copies of random purification using the structured random purification channel.
2. Use the optimal cloning map for pure stabilizer states.
3. Trace out purifying register from all copies.

- This means that

$$\epsilon_{\text{Cl}}^*(\text{Stab}_{4n}, t, t+1) \geq \epsilon_{\text{Cl}}^*(\mathcal{S}_n, t, t+1) \geq 0.14, \quad \text{for } t \leq n-1. \quad (14)$$

Summary

To summarize, we require three ingredients for our lower bound:




Figure: Proof Ingredients

Summary

Three ingredients:




Figure: Proof Ingredients

Beyond Stabilizer States: Abelian StateHSPs

- Abelian StateHSP: (G, μ, Ψ) where G is an Abelian group, μ is its unitary representation and Ψ is the set of states that hide some subgroup of G .
- Our methods provide cloning lower bound for Abelian StateHSPs with non-isomorphic irreps.

Adam Bouland, Tudor Giurgică-Tiron, and John Wright (2025). “The State Hidden Subgroup Problem and an Efficient Algorithm for Locating Unentanglement”. In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. STOC '25. Prague, Czechia: Association for Computing Machinery, pp. 463–470. ISBN: 9798400715105. DOI: 10.1145/3717823.3718118. URL: <https://doi.org/10.1145/3717823.3718118>.

Marcel Hinsche, Jens Eisert, and Jose Carrasco (2025). *The abelian state hidden subgroup problem: Learning stabilizer groups and beyond*. arXiv: 2505.15770 [quant-ph]. URL: <https://arxiv.org/abs/2505.15770>.

Beyond Stabilizer States: Abelian StateHSPs




Figure: Proof Ingredients for Cloning Lower Bound for Abelian StateHSP

Summary of Results

Problem Instance (Classical/Quantum)	Learning (with small constant error)	Sample Amplification/Cloning (with small constant error)
n -bit parities	$\Theta(n)$	$\Omega(n)$
Bell measurement outcome distribution	$\Theta(n)$	$\Omega(n)$
n -qubit stabilizer states	$\Theta(n)$	$\Omega(n)$

Table: **Overview of our main results:** A comparison between known sample complexity bounds for learning (with small constant error) and our sample complexity lower bounds for amplification/cloning.

Open Questions

1. Cloning Lower Bounds for other relevant classes of states such as
 - Graph states/Hypergraph states
 - Output states of shallow depth circuits
 - States prepared by circuits with few non-Clifford gates etc...
2. Non-trivial classes of states where cloning is easier than learning?
We have a first (somewhat) non-trivial example based on mutually unbiased bases but are looking for more.
3. Average-Case Cloning?

Thank you!

Backup

Beyond Stabilizer States: Abelian StateHSPs

- Abelian StateHSP: (G, μ, Ψ) where G is an Abelian group, μ is its unitary representation and Ψ is the set of states that hide some subgroup of G .

Proof Ingredients (Disclaimer: The unitary representation should have non-isomorphic irreps!)

- The **common eigenbasis measurement** $\{\Pi_\lambda = |\lambda\rangle\langle\lambda|\}_\lambda$: Measurement outcome is over elements of $H^\perp$ which have structure similar to linear.
- Let H be some subgroup of G of order α , then define mixed states

$$\sigma_H = \frac{1}{|H^\perp|} \sum_{\lambda \in H^\perp} |\lambda\rangle\langle\lambda|. \quad (15)$$

For these states, **Cloning via $\{\Pi_\lambda\}_\lambda$ is optimal** and a lower bound can be obtained from classical sample amplification.

- A structured random purification channel which lifts t copies of mixed state σ_H to t copies of random purification $|\sigma_{H,g}\rangle = \frac{1}{\sqrt{2^n}} \sum_{\lambda \in H^\perp} |\lambda\rangle \otimes \mu(g)|\lambda\rangle$.

Adam Bouland, Tudor Giurgică-Tiron, and John Wright (2025). "The State Hidden Subgroup Problem and an Efficient Algorithm for Locating Unentanglement". In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. STOC '25. Prague, Czechia: Association for Computing Machinery, pp. 463–470. ISBN: 9798400715105. DOI: 10.1145/3717823.3718118. URL: <https://doi.org/10.1145/3717823.3718118>.

Marcel Hinsche, Jens Eisert, and Jose Carrasco (2025). *The abelian state hidden subgroup problem: Learning stabilizer groups and beyond*. arXiv: 2505.15770 [quant-ph]. URL: <https://arxiv.org/abs/2505.15770>.

Cloning Lower Bounds for Abelian StateHSP

- We also show that random purification states $|\sigma_{H,g}\rangle$ are instances of a new Abelian StateHSP $(G \times G, \mu', \Psi')$ for some new unitary representation μ' .
- Thus, by using the ingredients mentioned previously, we can get cloning lower bounds for instances of this new Abelian StateHSP.

Cloning Lower Bounds for Abelian StateHSP

The random purifications of $\mathcal{S}_\alpha = \{\sigma_H\}_{H \leq G, |H|=\alpha}$ are instances of an Abelian StateHSP $(G \times G, \mu', \Psi')$ and their optimal cloning error is lower bounded as

$$\epsilon_{\text{Cl}}^*(\mathcal{S}, t, t+1) \geq \min_{\sigma_H \in \mathcal{S}_\alpha} p_{\sigma_H}/2, \quad \text{for } t \leq n_{H^\perp}(\alpha) - 1, \quad (16)$$

where p_{σ_H} is the probability of exactly learning H from $n_{H^\perp}(\alpha)$ copies of the state σ_H , with $n_{H^\perp}(\alpha)$ representing the number of generators of $H^\perp$.

Proof: No Sample Amplification Theorem

- Sample amplification can alternatively be thought of as a game between amplifier and distinguisher who knows the true distribution.




Figure: Distinguisher picture

Proof: No Sample Amplification Theorem

Distinguisher

Input: n samples, description of f .

Output: Accept or Reject

1. Run a consistent learning algorithm that outputs a uniformly random consistent hypothesis $\hat{f}$ from $\mathcal{F}$ using the t samples.
2. Accept (aka, output 1) if $\hat{f} = f$, else reject (aka, output 0).

The the sample amplification error can be lower bounded as:

$$\epsilon_{SA}^*(\mathcal{D}, n-1, n) \geq \min_{\Lambda} \left| \Pr_{Z_n}[\mathcal{A}^f(Z_n) = 1] - \Pr_{Z_{n-1}}[\mathcal{A}^f(\Lambda(Z_{n-1})) = 1] \right|. \quad (17)$$

where Z_t represents t labelled samples $(x_i, f(x_i))$.