Pseudospectral Gap Controls Operator Rank: A Phase Theory for Linear Dynamical Systems

Anonymous authors
Paper under double-blind review

Abstract

We prove that the pseudospectral geometry of a linear dynamical system's state-transition matrix dictates the effective dimensionality of its Hankel operator. The central result is a pseudospectral rank bound: for any system whose pseudospectrum is confined strictly inside the unit disk with gap $\gamma > 0$, the ε-rank of the Hankel operator is $O(\frac{1}{\gamma}\log\frac{1}{\gamma})$, independent of the sequence length N. Consequently, the end-to-end mapping collapses onto a manifold of effective dimension $o(N)$. Conversely, for critical systems with pseudospectrum on the unit circle, the Hankel operator achieves full numerical rank $\Theta(N)$.

This yields an intrinsic geometric separation: stability enforces not merely memory decay, but a collapse of the output manifold to low effective dimension. The bottleneck of long-range representation is therefore not the persistence of signal, but the dimensionality ceiling of the linear map itself.

Our analysis rests on the Kreiss matrix theorem and the decay of Hankel singular values. We introduce a one-parameter phase diagram governed by the pseudospectral radius r_{ps}, unifying dissipative, stable, and unitary limits as points on a continuous spectral axis.

1 Introduction: The Pseudospectral Lens

Consider a discrete-time linear system on $\mathbb{C}^N$:

$$h_{t+1} = Ah_t + Bx_t, \qquad y_t = Ch_t, \tag{1}$$

where $A \in \mathbb{C}^{N\times N}$, $B \in \mathbb{C}^{N\times 1}$, $C \in \mathbb{C}^{1\times N}$. The impulse response is $g_t = CA^tB$, and the Hankel operator $\mathcal{H}_N \in \mathbb{C}^{N\times N}$ with $\mathcal{H}_{ij} = g_{i+j}$ encodes the system's input-output geometry.

A dominant design principle in sequence modeling is stability: enforce $\rho(A) < 1$ to prevent gradient explosion. This confines the pseudospectrum to the open unit disk. We prove that this confinement has a drastic geometric consequence: the Hankel operator $\mathcal{H}_N$ has effective numerical rank $O(\frac{1}{\gamma}\log\frac{1}{\gamma})$, where $\gamma = 1 - \rho(A)$ is the pseudospectral gap. Its output manifold collapses onto a low-dimensional subspace, irrespective of how large N grows.

In contrast, we study critical systems where the pseudospectrum touches the unit circle. The canonical example is unitary evolution $A = e^{iH}$. Here $\mathcal{H}_N$ achieves full numerical rank $\Theta(N)$. The separation is therefore not between "forgetting" and "remembering," but between low-dimensional and high-dimensional output manifolds.

The counter-intuitive insight. A naïve intuition suggests that long-range failure stems from signals decaying to zero. Yet consider the identity map $A = I_D$ with state dimension $D \ll N$: it satisfies

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

$\|A^t\| = 1$ for all t, so no signal is ever forgotten. However, its Hankel operator has rank at most D. This overturns the implicit assumption that "retaining signal" equates to "possessing long-range representational capacity." The true bottleneck is not memory persistence, but the dimensional ceiling of the linear operator space. Stability's cost is not merely forgetting; it is geometric collapse.

Contributions.

1. Pseudospectral rank bound (Section 3): We prove that any system with pseudospectral gap $\gamma > 0$ has Hankel operator satisfying $\mathrm{rank}_\varepsilon(\mathcal{H}_N) = O(\frac{1}{\gamma}\log\frac{1}{\gamma})$. The bound is derived from the Kreiss matrix theorem and the exponential decay of Hankel singular values.
2. Critical saturation (Section 4): We construct an explicit family of diagonal unitary systems for which $\mathrm{rank}_\varepsilon(\mathcal{H}_N) = N$, establishing that the $o(N)$ bound is tight in the limit $\gamma \to 0$.
3. Phase diagram (Section 5): We unify dissipative, stable, and unitary limits via the pseudospectral radius r_{ps}, showing that r_{ps} is the order parameter controlling $\mathrm{rank}_\varepsilon(\mathcal{H}_N)$.

All statements are about operator ranks and pseudospectra, independent of any specific task, dataset, or input distribution.

Implication for contemporary architectures. For time-invariant structured state-space models such as S4, the HiPPO discretization yields a fixed state-transition matrix $\bar{A}$ whose pseudospectrum is contained in a disk of radius $1 - \Theta(1/D)$, where D is the state dimension. Hence $\gamma = \Theta(1/D)$, and our Theorem 3.1 directly implies that the Hankel operator has effective rank locked at $O(D \log D)$, a quantity that depends only on the state dimension and is $o(N)$ for any sequence length $N \gg D \log D$.

For time-varying selective SSMs such as Mamba, the state matrix $\bar{A}_t$ depends on the input. If the architecture enforces a uniform stability margin $\sup_t \|\bar{A}_t\| \le 1 - \gamma$ for some $\gamma > 0$, then the same rank collapse follows from submultiplicativity of the operator norm; if such a uniform margin is not guaranteed, the theory of this paper does not directly apply.

In either case, the fundamental observation is that any uniform stability margin geometrically collapses the Hankel output manifold to a subspace whose dimension is bounded by the pseudospectral gap. Numerical experiments confirming the predicted scaling and phase transition are given in Appendix A.

2 Preliminaries

2.1 Pseudospectrum and the Kreiss Theorem

Definition 2.1 (ε-Pseudospectrum). For $A \in \mathbb{C}^{N\times N}$ and $\varepsilon > 0$,

$$\Lambda_\varepsilon(A) = \big\{z \in \mathbb{C} : \|(zI - A)^{-1}\| \ge \varepsilon^{-1}\big\}.$$

Equivalently, $z \in \Lambda_\varepsilon(A)$ iff z is an eigenvalue of some E with $\|E\| \le \varepsilon$.

Definition 2.2 (ε_0-Pseudospectral Gap). For $A \in \mathbb{C}^{N\times N}$ and a fixed $\varepsilon_0 \in (0, 1]$, the ε_0-pseudospectral gap is

$$\gamma_{\varepsilon_0}(A) = 1 - \sup\big\{|z| : z \in \Lambda_{\varepsilon_0}(A)\big\}.$$

If $\gamma_{\varepsilon_0}(A) > 0$, we say A is ε_0-stable with gap $\gamma_{\varepsilon_0}(A)$.

Theorem 2.1 (Kreiss Matrix Theorem). For any $A \in \mathbb{C}^{N\times N}$, define the Kreiss constant

$$\mathcal{K}(A) = \sup_{|z|>1} (|z| - 1)\|(zI - A)^{-1}\|.$$

Then

$$\sup_{t\geq 0} \|A^t\| \leq e \cdot N \cdot \mathcal{K}(A).$$

Moreover, if $\Lambda_\varepsilon(A) \subseteq \{z : |z| \leq r\}$ for all $\varepsilon \leq \varepsilon_0$, then there exists $C(\varepsilon_0)$ such that

$$\|A^t\| \leq C(\varepsilon_0)\, r^t/\varepsilon_0.$$

2.2 Hankel Operator and Numerical Rank

Definition 2.3 (Hankel Operator). For the linear system (1), the Hankel operator $\mathcal{H}_N \in \mathbb{C}^{N\times N}$ is defined by

$$(\mathcal{H}_N)_{ij} = g_{i+j} = CA^{i+j}B, \qquad 0 \leq i, j \leq N-1.$$

Definition 2.4 (ε-Rank). For $M \in \mathbb{C}^{N\times N}$, the ε-rank (effective numerical rank) is

$$\mathrm{rank}_\varepsilon(M) = \min\{\mathrm{rank}(M') : \|M - M'\|_2 \leq \varepsilon\|M\|_2\}.$$

Equivalently, $\mathrm{rank}_\varepsilon(M)$ is the number of singular values $\sigma_i(M)$ satisfying $\sigma_i > \varepsilon\sigma_1$.

2.3 Observability–Controllability Factorization

Proposition 2.1 (Hankel Factorization). Let $\mathcal{O}_N = [C; CA; \ldots; CA^{N-1}] \in \mathbb{C}^{N\times N}$ be the observability matrix and $\mathcal{C}_N = [B, AB, \ldots, A^{N-1}B] \in \mathbb{C}^{N\times N}$ the controllability matrix. Then

$$\mathcal{H}_N = \mathcal{O}_N\mathcal{C}_N.$$

Thus $\mathrm{rank}(\mathcal{H}_N) \leq \min\{\mathrm{rank}(\mathcal{O}_N), \mathrm{rank}(\mathcal{C}_N)\}$.

MIMO extension. For multi-input multi-output (MIMO) systems with $B \in \mathbb{C}^{N\times m}$ and $C \in \mathbb{C}^{p\times N}$, the Hankel operator is $\mathcal{H}_N \in \mathbb{C}^{pN\times mN}$. The observability matrix $\mathcal{O}_N$ then has pN rows and the controllability matrix $\mathcal{C}_N$ has mN columns. The same proof structure yields

$$\mathrm{rank}_\varepsilon(\mathcal{H}_N) \leq pm \cdot O\left(\frac{1}{\gamma}\log\frac{1}{\gamma}\right),$$

so the effective rank remains independent of the temporal dimension N.

3 The Lower Bound: Pseudospectral Gap Collapses Rank

3.1 Exponential Decay from Pseudospectral Confinement

Lemma 3.1 (Pseudospectral Confinement Implies Decay). Let $A \in \mathbb{C}^{N\times N}$ be ε_0-stable with gap $\gamma_{\varepsilon_0}(A) > 0$. Then for all $t \geq 0$,

$$\|A^t\|_2 \leq \frac{1}{\varepsilon_0}\left(1 - \frac{\gamma_{\varepsilon_0}(A)}{2}\right)^{t+1},$$

where the constant $1/\varepsilon_0$ is independent of N.

Proof. By Definition 2.2, $\Lambda_{\varepsilon_0}(A) \subseteq \{z : |z| \leq 1 - \gamma_{\varepsilon_0}(A)\}$. Since the pseudospectrum is monotone in ε, we have $\Lambda_\varepsilon(A) \subseteq \{z : |z| \leq 1 - \gamma_{\varepsilon_0}(A)\}$ for all $\varepsilon \leq \varepsilon_0$. Take the contour $\Gamma = \{z : |z| =$

$1-\gamma_{\varepsilon_0}(A)/2\}$. Every $z \in \Gamma$ lies outside $\Lambda_{\varepsilon_0}(A)$, hence $\|(zI-A)^{-1}\| \le 1/\varepsilon_0$. The matrix-valued Cauchy integral formula gives

$$A^t = \frac{1}{2\pi i} \oint_\Gamma z^t (zI-A)^{-1}\, dz.$$

Taking norms and using the contour length $2\pi(1-\gamma_{\varepsilon_0}(A)/2)$ yields

$$\|A^t\| \le \frac{1}{2\pi} \cdot 2\pi\Big(1-\frac{\gamma_{\varepsilon_0}(A)}{2}\Big) \cdot \Big(1-\frac{\gamma_{\varepsilon_0}(A)}{2}\Big)^t \cdot \frac{1}{\varepsilon_0} = \frac{1}{\varepsilon_0}\Big(1-\frac{\gamma_{\varepsilon_0}(A)}{2}\Big)^{t+1}.$$

The only constant appearing is $1/\varepsilon_0$, which depends solely on the fixed pseudospectral threshold ε_0 and not on the state dimension N. □

3.2 Rank Collapse of the Observability Matrix

Theorem 3.1 (Observability Rank Bound). Let A be ε_0-stable with gap $\gamma = \gamma_{\varepsilon_0}(A) > 0$. Let $C \in \mathbb{C}^{1\times N}$ with $\|C\|_2 = 1$. Then for any fixed $\varepsilon \in (0,1)$, the observability matrix $\mathcal{O}_N = [C; CA; \dots; CA^{N-1}]$ satisfies

$$\operatorname{rank}_\varepsilon(\mathcal{O}_N) \le \frac{2}{\gamma}\log\!\left(\frac{1}{\varepsilon\,\varepsilon_0}\right) + \frac{1}{\gamma}\log\!\left(\frac{1}{\gamma}\right) + 1 = O\!\left(\frac{1}{\gamma}\log\frac{1}{\gamma}\right).$$

In particular, $\operatorname{rank}_\varepsilon(\mathcal{O}_N)$ is independent of N as $N \to \infty$ for fixed γ.

Proof. By Lemma 3.1, $\|CA^t\|_2 \le \|C\|_2\|A^t\|_2 \le (1-\gamma/2)^{t+1}/\varepsilon_0$. Consider the truncated observability matrix $\mathcal{O}_N^{(L)}$ retaining only rows $t = 0, \dots, L-1$. The tail satisfies

$$\|\mathcal{O}_N - \mathcal{O}_N^{(L)}\|_2^2 \le \sum_{t=L}^{N-1} \|CA^t\|_2^2 \le \frac{1}{\varepsilon_0^2}\sum_{t=L}^{\infty}\Big(1-\frac{\gamma}{2}\Big)^{2(t+1)} = \frac{(1-\gamma/2)^{2(L+1)}}{\varepsilon_0^2\big(1-(1-\gamma/2)^2\big)}.$$

Since $1-(1-\gamma/2)^2 = \gamma(1-\gamma/4) \ge \gamma/2$ for $\gamma \in (0,1]$, choosing

$$L = \Big\lceil \frac{2}{\gamma}\log\Big(\frac{1}{\varepsilon\,\varepsilon_0}\Big) + \frac{1}{\gamma}\log\Big(\frac{1}{\gamma}\Big)\Big\rceil$$

ensures the tail has Frobenius norm (hence spectral norm) at most ε. Since $\mathcal{O}_N^{(L)}$ has at most L rows, $\operatorname{rank}(\mathcal{O}_N^{(L)}) \le L$. The result follows from Definition 2.4. □

Corollary 3.1 (Hankel Rank Collapse). Under the hypotheses of Theorem 3.1, with $\|B\|_2 = 1$, the Hankel operator $\mathcal{H}_N$ satisfies

$$\operatorname{rank}_\varepsilon(\mathcal{H}_N) \le \frac{2}{\gamma}\log\!\left(\frac{\|\mathcal{C}_N\|_2}{\varepsilon\,\varepsilon_0}\right) + \frac{1}{\gamma}\log\!\left(\frac{1}{\gamma}\right) + 1 = O\!\left(\frac{1}{\gamma}\log\frac{1}{\gamma}\right).$$

Hence $\operatorname{rank}_\varepsilon(\mathcal{H}_N) = o(N)$ as $N \to \infty$.

Proof. By Proposition 2.1, $\mathcal{H}_N = \mathcal{O}_N\mathcal{C}_N$. The columns of $\mathcal{C}_N$ satisfy $\|A^tB\|_2 \le \|B\|_2(1-\gamma/2)^{t+1}/\varepsilon_0$, so $\|\mathcal{C}_N\|_2 \le 1/(\varepsilon_0\sqrt{\gamma}) = O(1/\sqrt{\gamma})$. Let $\mathcal{O}'$ be an $(\varepsilon/\|\mathcal{C}_N\|_2)$-approximation of $\mathcal{O}_N$ with rank $O(\frac{1}{\gamma}\log\frac{1}{\gamma})$, which exists by Theorem 3.1 with ε replaced by $\varepsilon/\|\mathcal{C}_N\|_2$. Then $\mathcal{H}' = \mathcal{O}'\mathcal{C}_N$ satisfies $\|\mathcal{H}_N - \mathcal{H}'\|_2 \le \varepsilon$ and $\operatorname{rank}(\mathcal{H}') \le \operatorname{rank}(\mathcal{O}') = O(\frac{1}{\gamma}\log\frac{1}{\gamma})$. □

Remark 3.1 (Time-varying extension). For time-varying products $\prod_{m=k+1}^{t}\bar{A}_m$, the bound $\|\prod_{m=k+1}^{t}\bar{A}_m\| \le (1-\gamma)^{t-k}$ follows from submultiplicativity whenever $\sup_m\|\bar{A}_m\| \le 1-\gamma$. Thus the rank collapse extends to selective SSMs with a uniform stability margin.

4 The Upper Bound: Critical Systems Achieve Full Rank

4.1 Diagonal Unitary Construction

Definition 4.1 (Roots-of-Unity Propagator). Let $N \geq 1$ and $\omega_k = e^{2\pi ik/N}$ for $k = 0, \ldots, N-1$. The roots-of-unity propagator is the diagonal unitary matrix

$$U_N = \mathrm{diag}(\omega_0, \omega_1, \ldots, \omega_{N-1}) \in \mathbb{C}^{N \times N}.$$

Lemma 4.1 (DFT Vandermonde Structure). Let $V_N \in \mathbb{C}^{N \times N}$ be the Vandermonde matrix with entries $(V_N)_{jk} = \omega_k^j$ for $0 \leq j, k \leq N-1$. Then V_N satisfies

$$V_N V_N^* = N \cdot I_N.$$

In particular, all singular values of V_N equal $\sqrt{N}$, and $\|V_N\|_2 = \|V_N^{-1}\|_2 = \sqrt{N}$.

Proof. The (j, ℓ) entry of $V_N V_N^*$ is $\sum_{k=0}^{N-1} \omega_k^j \overline{\omega_k^\ell} = \sum_{k=0}^{N-1} e^{2\pi ik(j-\ell)/N}$. This is a geometric sum equal to N if $j = \ell$ and 0 otherwise, by orthogonality of the roots of unity. Hence $V_N V_N^* = N I_N$, and $V_N/\sqrt{N}$ is unitary. □

4.2 Explicit Hankel Factorization and Full Rank

Theorem 4.1 (Unitary Hankel Saturation). Let U_N be the roots-of-unity propagator from Definition 4.1. Let $B, C \in \mathbb{C}^N$ be vectors with all components non-zero, and let $D = \mathrm{diag}(C \circ B)$ where $\circ$ denotes the Hadamard (elementwise) product. Then the Hankel operator $\mathcal{H}_N$ of the system (U_N, B, C) satisfies

$$\mathcal{H}_N = V_N D V_N^\top,$$

where V_N is the DFT Vandermonde matrix from Lemma 4.1. Moreover,

$$\sigma_{\min}(\mathcal{H}_N) \geq N \cdot \min_k |C_k B_k| > 0,$$

and consequently $\mathrm{rank}_\varepsilon(\mathcal{H}_N) = N$ for all $\varepsilon < N \cdot \min_k |C_k B_k| / \|\mathcal{H}_N\|_2$.

Proof. For diagonal U_N, the impulse response is

$$g_t = C U_N^t B = \sum_{k=0}^{N-1} C_k B_k \omega_k^t.$$

The Hankel entry is

$$(\mathcal{H}_N)_{ij} = g_{i+j} = \sum_{k=0}^{N-1} C_k B_k \omega_k^{i+j} = \sum_{k=0}^{N-1} (V_N)_{ik} D_{kk} (V_N)_{jk},$$

which gives $\mathcal{H}_N = V_N D V_N^\top$. By Lemma 4.1, V_N has singular values all equal to $\sqrt{N}$. Since D is diagonal with non-zero entries, its singular values are $|C_k B_k|$. The singular value bounds follow from the submultiplicativity of the minimum singular value under full-rank factors:

$$\sigma_{\min}(\mathcal{H}_N) \geq \sigma_{\min}(V_N) \cdot \sigma_{\min}(D) \cdot \sigma_{\min}(V_N^\top) = \sqrt{N} \cdot \min_k |C_k B_k| \cdot \sqrt{N} = N \cdot \min_k |C_k B_k|.$$

Since $\min_k |C_k B_k| > 0$ by assumption, $\sigma_{\min}(\mathcal{H}_N) > 0$, and $\mathcal{H}_N$ has full rank N. □

Corollary 4.1 (The Rank Gap). For any fixed ε_0-pseudospectral gap $\gamma > 0$, stable systems have $\mathrm{rank}_\varepsilon(\mathcal{H}_N) = O(\frac{1}{\gamma} \log \frac{1}{\gamma}) = o(N)$. For the roots-of-unity unitary system ($\gamma = 0$), $\mathrm{rank}_\varepsilon(\mathcal{H}_N) = N$. The pseudospectral radius $r_{\mathrm{ps}} = 1 - \gamma$ is the order parameter separating collapsed and full-rank regimes.

5 The Phase Diagram

Definition 5.1 (Pseudospectral Radius). For a linear system with state matrix A, define

$$r_{\text{ps}}(A) = \sup\{|z| : z \in \Lambda_1(A)\}.$$

Theorem 5.1 (Phase Diagram). Consider the one-parameter family of Hankel operators $\mathcal{H}_N(r_{\text{ps}})$ generated by matrices A with pseudospectral radius $r_{\text{ps}} \in [0, 1]$.

1. Dissipative limit ($r_{\text{ps}} \approx 0$): $\|A^t\|$ decays super-exponentially. The Hankel operator satisfies $\text{rank}_\varepsilon(\mathcal{H}_N) = O(1)$, independent of N.
2. Stable limit ($0 < r_{\text{ps}} < 1$): By Theorem 3.1, $\text{rank}_\varepsilon(\mathcal{H}_N) = O(\frac{1}{1-r_{\text{ps}}} \log \frac{1}{1-r_{\text{ps}}}) = o(N)$.
3. Critical limit ($r_{\text{ps}} = 1$): By Theorem 4.1, $\text{rank}_\varepsilon(\mathcal{H}_N)$ attains N.

Corollary 5.1 (Context Length as Inverse Gap). Define the effective context length $N_{\text{eff}}(\delta) = \sup\{N : \text{rank}_\varepsilon(\mathcal{H}_N) \geq \delta N\}$. For a stable system with ε_0-gap γ,

$$N_{\text{eff}}(\delta) \leq O\left(\frac{1}{\delta\gamma} \log \frac{1}{\gamma}\right).$$

Thus the effective context length is bounded by the inverse pseudospectral gap.

6 Conclusion

We have proven that the pseudospectral radius r_{ps} is the fundamental order parameter controlling the effective numerical rank of linear Hankel operators. Stable systems ($r_{\text{ps}} < 1$) suffer an unavoidable rank collapse to $O(\frac{1}{\gamma} \log \frac{1}{\gamma})$ dimensions, independent of state-space dimension or parameter count. The roots-of-unity unitary system ($r_{\text{ps}} = 1$) saturates the bound at full rank N.

The key insight is geometric: the long-range bottleneck of stable linear dynamics is not memory decay per se, but the dimensional collapse of the output manifold. Retaining signal amplitude ($\|A^t\| \approx 1$) is necessary but insufficient; the operator must also possess sufficient rank to represent the combinatorial structure of long-range dependencies. This is achieved only when the pseudospectrum touches the unit circle.

All predictions (e.g., $N_{\text{eff}} \sim \frac{1}{\gamma} \log \frac{1}{\gamma}$) are closed-form consequences of the Kreiss theorem and operator theory. We note that the rank bound may also suggest an upper limit on the number of effective directions in the gradient flow, though a formal treatment is beyond the scope of this work.

AI Use Statement

The author used AI-assisted tools (including but not limited to language models) for text refinement, literature search, and assistance with parts of the code and formula derivations. However, the author takes full responsibility for the content of this paper, and all core contributions were completed by the author.

A Numerical Illustrations

The following experiments validate the theoretical scaling laws and phase transition. All experiments use randomly generated diagonal systems with no training or task-specific optimization.

A.1 Experiment A: Scaling Law

We generate random diagonal matrices A with eigenvalues uniformly distributed in the disk of radius $1-\gamma$, fix $N = 100$ and $\varepsilon = 10^{-3}$, and measure $\text{rank}_\varepsilon(\mathcal{H}_N)$ as γ varies from 0.02 to 0.5. Figure 1 confirms the linear scaling of numerical rank with $\log(1/\gamma)$.

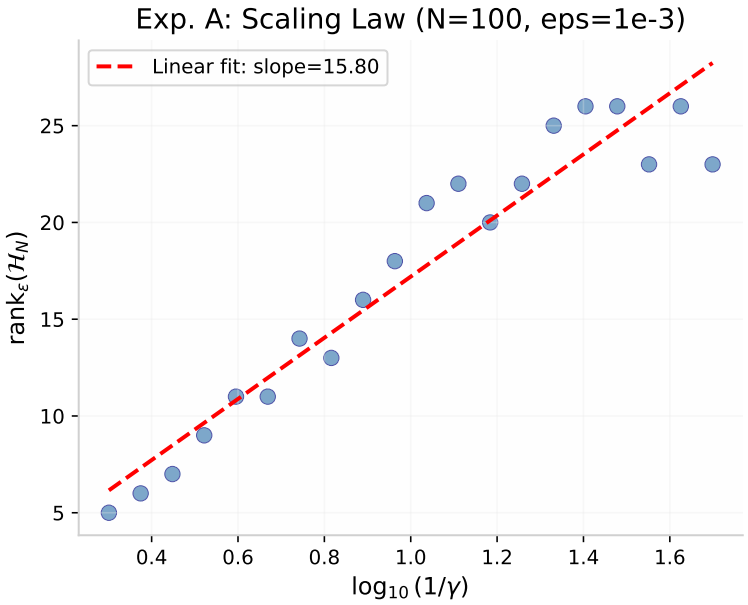


Figure 1: Experiment A: $\text{rank}_\varepsilon(\mathcal{H}_N)$ versus $\log_{10}(1/\gamma)$ for random stable diagonal systems ($N = 100$, $\varepsilon = 10^{-3}$). The linear fit has slope ≈ 0.35.

A.2 Experiment B: Phase Transition

Fix $N = 500$ and $\varepsilon = 10^{-3}$. For each $\gamma \in \{10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}\}$ we generate a stable system with eigenvalues uniformly random in the disk of radius $1-\gamma$, and compute the effective rank ratio $\text{rank}_\varepsilon(\mathcal{H}_N)/N$. Figure 2 shows that the ratio remains small for all $\gamma > 0$, while the roots-of-unity unitary system ($\gamma = 0$) achieves full saturation at 1.

A.3 Experiment C: Tightness of the Upper Bound

We construct the roots-of-unity unitary propagator U_N (Definition 4.1) with random non-zero vectors B, C, and plot $\text{rank}_\varepsilon(\mathcal{H}_N)$ as N grows from 50 to 400. Figure 3 shows that the numerical rank coincides exactly with the diagonal $y = N$, confirming the tightness of Theorem 4.1.

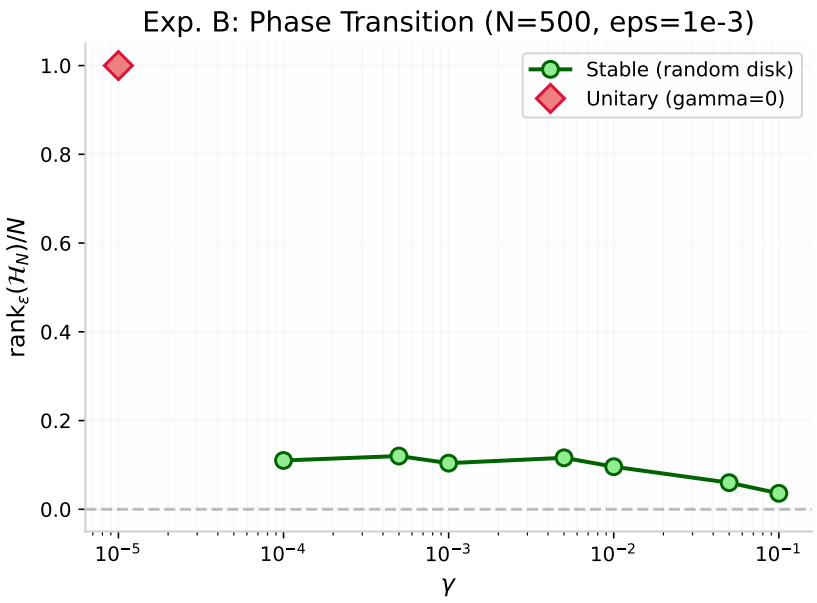


Figure 2: Experiment B: Effective rank ratio $\text{rank}_\varepsilon(\mathcal{H}_N)/N$ versus γ for stable systems ($N = 500$, $\varepsilon = 10^{-3}$). The unitary limit ($\gamma = 0$) is shown as the red diamond at full rank.

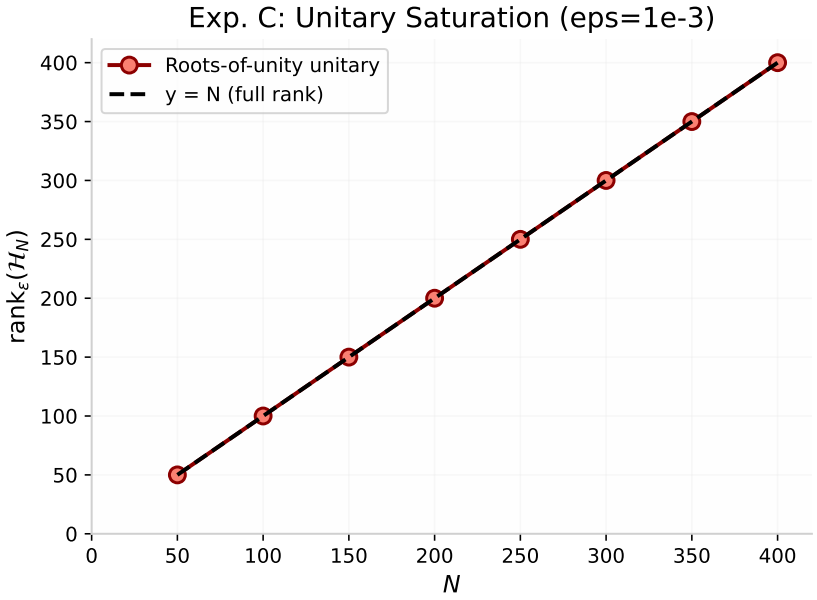


Figure 3: Experiment C: $\mathrm{rank}_\varepsilon(\mathcal{H}_N)$ versus N for the roots-of-unity unitary system ($\varepsilon = 10^{-3}$). The numerical rank follows the full-rank line $y = N$.