Manifolds, Random Matrices and Spectral Gaps: The geometric phases of Generative Diffusion

Enrico Ventura*[1], Beatrice Achilli*[1], Gianluigi Silvestri*[2,3], Carlo Lucibello[1] and Luca Ambrogioni[2]

[1]Bocconi University. [2]Donders Centre for Cognition, Radboud University. [3]OnePlanet Research Center. *Equal contribution

Diffusion Models

Generative Diffusion is the frontier of image and video generation [1].

Forward process: $\mathbf{x}_0 \sim p_0(\mathbf{x}), \quad \mathrm{d}\mathbf{x}_t = \mathrm{d}\mathbf{Z}_t$

Backward process:
$$\begin{cases} \mathbf{x}_{t_f} \sim \mathcal{N}(0, t_f I_d) \\ \mathrm{d}\mathbf{x}_t = -s(\mathbf{x}_t, t)\mathrm{d}t + d\mathbf{Z}_t \end{cases}$$

Score: $s(\mathbf{x}_t, t) = \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t)$

Jacobian of the Score function: $J(\mathbf{x}_t, t) = \nabla_{\mathbf{x}} s(\mathbf{x}_t, t)$

Manifolds and Spectral Gaps

We assume that the target distribution is supported on a **manifold** $\mathcal{M}$ with **latent dimension** $m < d$.
For a small perturbation **p** around a manifold point the score can be approximated as

$$s(\mathbf{x}^* + \mathbf{p}, t) \approx J(\mathbf{x}^*, t)\ \mathbf{p}$$

Perturbations aligned with the tangent space of $\mathcal{M}$ correspond to small eigenvalues, orthogonal perturbations to large **eigenvalues**.

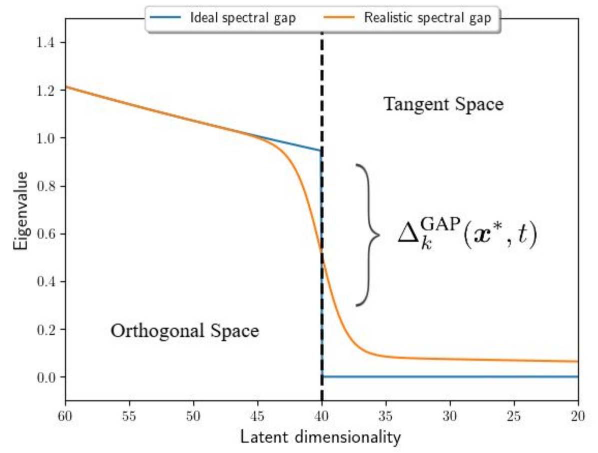

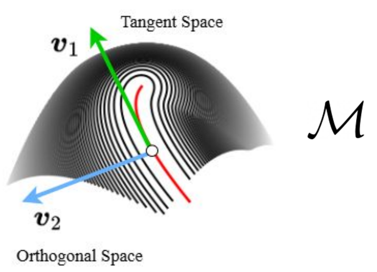

The **hidden dimensionality** of the manifold can be estimated from the **location of a drop** in the sorted spectrum of J [2].

Subspaces and Intermediate Gaps

The spectrum of the eigenvalues of J can show **sub-gaps**, i.e. subsets of null eigenvalues, indicating the presence of separate **subspaces** associated to certain **local variances**.

Phenomenology

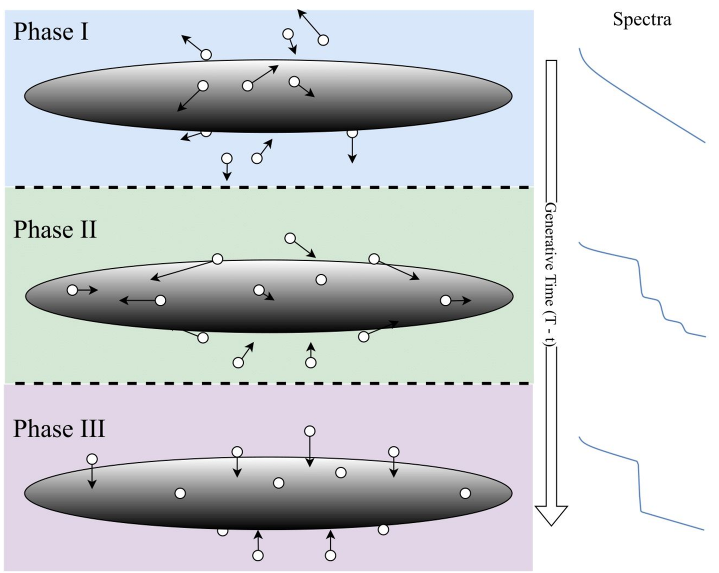

Trivial Phase (I)

The **diffusing trajectory is Brownian** regardless the manifold directions.

Manifold Coverage (II)

The process progressively fits the target distribution defined on the manifold: **different subspaces with different variances emerge as intermediate gaps** in the spectrum of J. The time scale at which the *k*-th gap is maximally visible is:

$$t_k = \mathcal{O}(\sigma_k \cdot \sigma_{k+1})$$

Manifold Consolidation (III)

Asymptotic closure of the intermediate gaps and the sharpening of the total manifold gap, indicating the full dimensionality of $\mathcal{M}$

$$s(\mathbf{x}_t, t) \simeq \frac{1}{t}\Big[\Pi - I_d\Big]\mathbf{x}, \quad \Pi = F(F^\top F)^{-1}F^\top$$

The geometric phases and manifold overfitting

Likelihood-based generative models are prone to **manifold overfitting**: the trained model fits the manifold while ignoring its internal density, resulting in poor generation [3]. Our analysis suggests that **generative diffusion models overcome this limitation** since the **score function is sensitive to the target distribution,** defined on the manifold, **even at intermediate times.**

Theoretical analysis

Linear manifold data [4]: $\mathbf{y}^\mu = F\mathbf{z}^\mu, \quad \mathbf{z}^\mu \sim \mathcal{N}(0, I_m), \quad F \in \mathbb{R}^{d\times m}$

Jacobian of the Score function:

$$J_t = \frac{1}{t}F\left[I_m + \frac{1}{t}F^\top F\right]^{-1} F^\top - I_d$$

Choice of the matrix F as

$$F_{ij} \sim \mathcal{N}(0, \sigma_1^2/m) \text{ for } j \leq fm$$

$$F_{ij} \sim \mathcal{N}(0, \sigma_2^2/m) \text{ for } j > fm$$

The **spectrum of** $F^\top F$ at any time t can be computed via **Replica Method** [5].

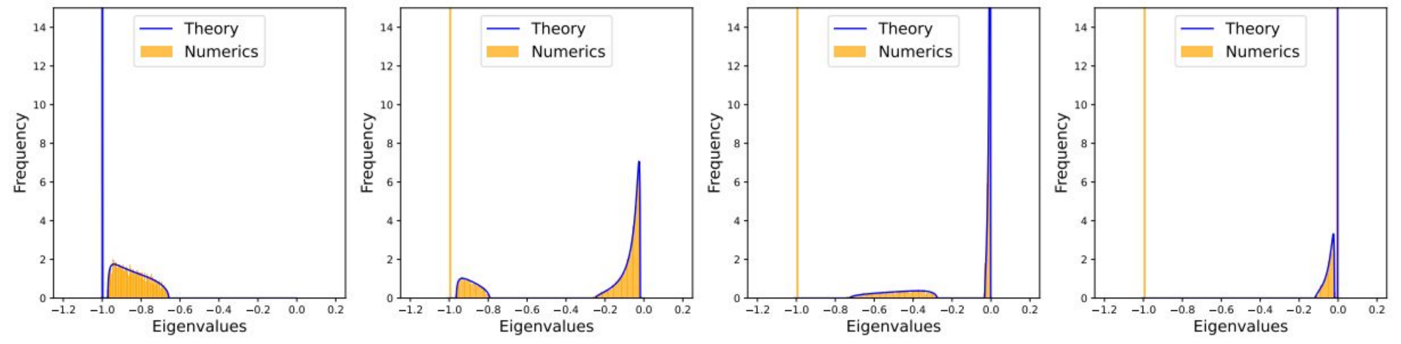

(a) t = 10 (b) t = 0.1 (c) t = 0.01 (d) t = 5e-4

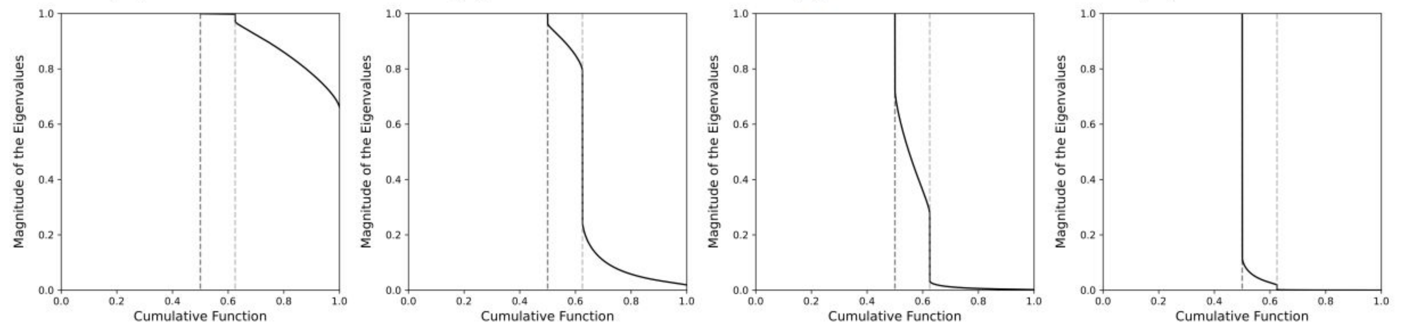

(e) t = 10 (f) t = 0.1 (g) t = 0.01 (h) t = 5e-4

This same analysis can be extended to the more **general case** where the spectral density is known to be formed by different detached bulks, associated with **hierarchically smaller variances** of the data.

Experiments

Synthetic Linear Datasets

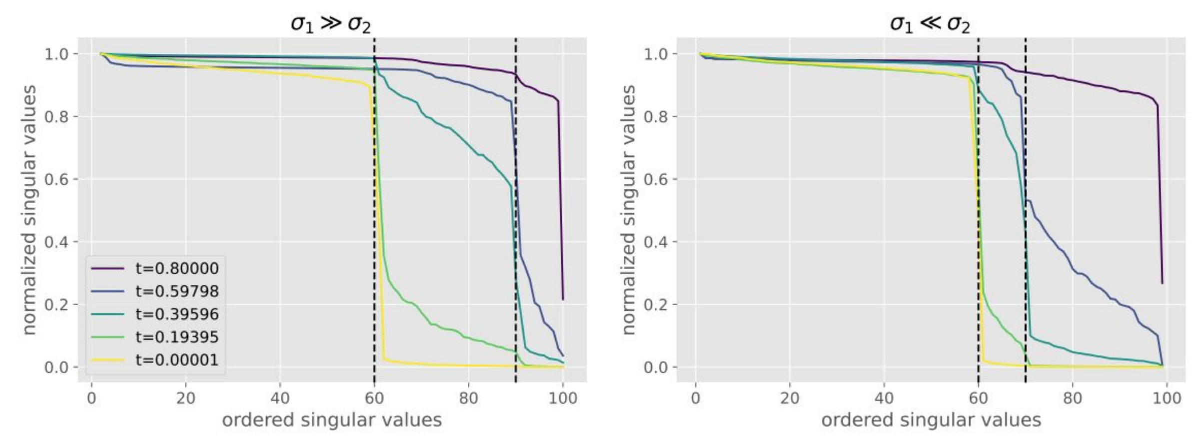

Theory Vs. Neural Networks:

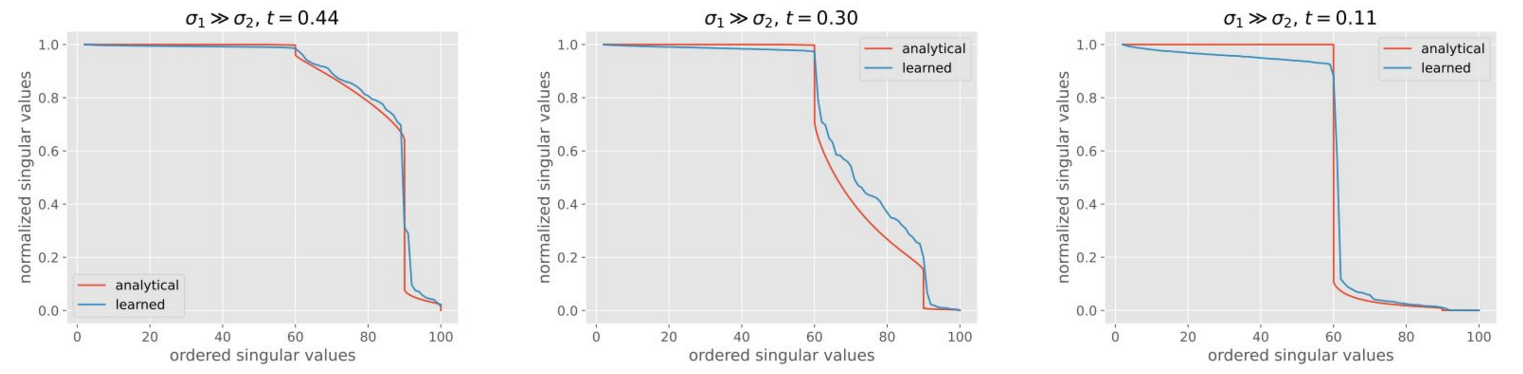

Natural Image Datasets

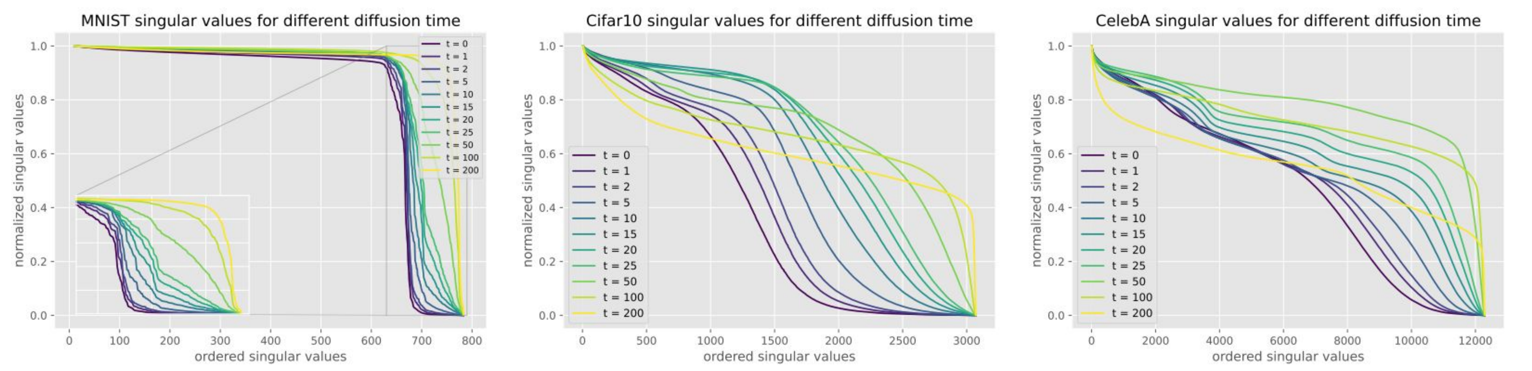

References & Acknowledgments

Enrico Ventura, Beatrice Achilli, Gianluigi Silvestri, Carlo Lucibello, Luca Ambrogioni (2025). "Manifolds, Random Matrices and Spectral Gaps: The geometric phases of Generative Diffusion".
[1] Ling Yang et al. (2024). "Diffusion Models: A Comprehensive Survey of Methods and Applications".
[2] Yan Stanczuk et al. (2023). "Your diffusion model secretly knows the dimension of the data manifold".
[3] Gabriel Loaiza-Ganem et al. (2022). "Diagnosing and fixing manifold overfitting in deep generative models".
[4] Sebastian Goldt et al. (2020). "Modeling the Influence of Data Structure on Learning in Neural Networks: The Hidden Manifold Model".
[5] Jean-Philippe Bouchaud and Marc Potters (2020). "A First Course in Random Matrix Theory: for Physicists, Engineers and Data Scientists".

This poster benefited from European Union - Next Generation EU funds, component M4.C2, investment 1.1. - CUP J53D23001330001.