

A RANDOM MATRIX THEORY PERSPECTIVE ON THE CONSISTENCY OF DIFFUSION MODELS

000
001
002
003
004
005 **Anonymous authors**
006 Paper under double-blind review
007
008
009
010

ABSTRACT

011 Diffusion models trained on different, non-overlapping subsets of a dataset often
012 produce strikingly similar outputs when given the same noise seed. We trace this
013 consistency to a simple linear effect: the shared Gaussian statistics across splits
014 already predict much of the generated images. To formalize this, we develop a
015 random matrix theory (RMT) framework that quantifies how finite datasets shape
016 the expectation and variance of the learned denoiser and sampling map in the linear
017 setting. For expectations, sampling variability acts as a renormalization of the
018 noise level through a self-consistent relation $\sigma^2 \mapsto \kappa(\sigma^2)$, explaining why limited
019 data overshrink low-variance directions and pull samples toward the dataset mean.
020 For fluctuations, our variance formulas reveal three key factors behind cross-split
021 disagreement: *anisotropy* across eigenmodes, *inhomogeneity* across inputs, and
022 overall scaling with dataset size. Extending deterministic-equivalence tools to
023 fractional matrix powers further allows us to analyze entire sampling trajectories.
024 The theory sharply predicts the behavior of linear diffusion models, and we validate
025 its predictions on UNet and DiT architectures in their non-memorization regime,
026 identifying where and how samples deviates across training data split. This provides
027 a principled baseline for reproducibility in diffusion training, linking spectral
028 properties of data to the stability of generative outputs.
029

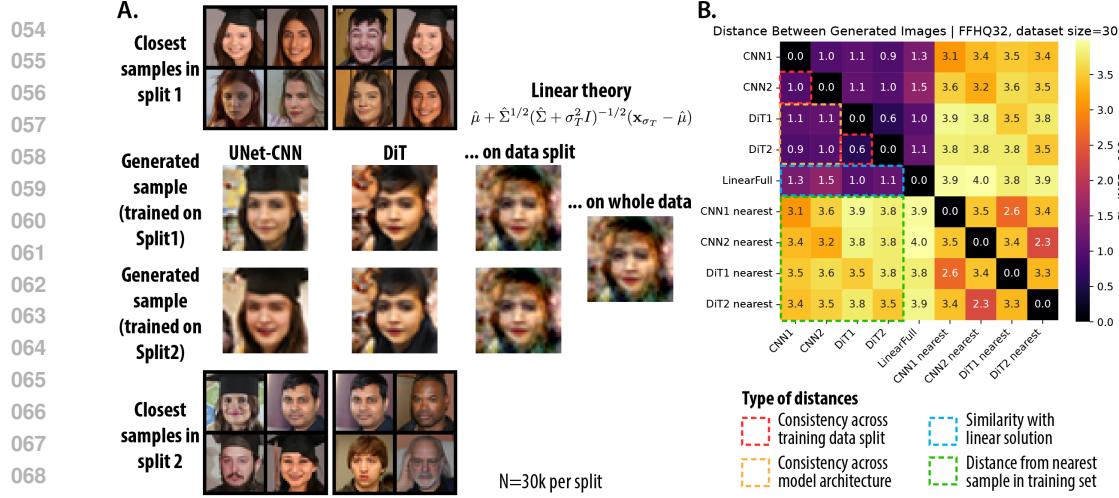
1 INTRODUCTION

030 Diffusion models and their relatives such as flow matching have become the dominant generative
031 modeling paradigm across diverse domains, including images, video, and proteins. By learning a
032 time-dependent vector field, these models transform Gaussian noise into structured samples through
033 an ordinary differential equation (ODE) or its stochastic variants (Song et al., 2021; Albergo et al.,
034 2023).
035

036 A distinctive feature of diffusion models is their striking *consistency across training runs* (Figure 1).
037 When trained on the same distribution, even with disjoint datasets, different architectures, or repeated
038 initializations, diffusion models often map the same noise seed to highly similar outputs under the
039 deterministic probability flow (Kadkhodaie et al., 2024; Zhang et al., 2024). This phenomenon
040 contrasts with other generative modeling frameworks including GANs and VAEs, where the isotropic
041 Gaussian latent space admits arbitrary rotations, leading to run-to-run variability in the mapping from
042 latent codes to data (Martinez & Pearson, 2022).
043

044 **Why does consistency matter?** Consistency across non-overlapping data splits suggests that
045 diffusion models recover aspects of the underlying *data manifold* that are insensitive to the specific
046 training set. This raises fundamental questions about how such models generalize beyond their
047 training samples, to what extent they memorize idiosyncratic data, and whether their outputs reflect
048 universal statistical regularities of the distribution. These issues connect to emerging theoretical
049 and empirical debates on generalization, memorization, and creativity in diffusion models (Kamb &
050 Ganguli, 2024; Niedoba et al., 2024; Kadkhodaie et al., 2024; Chen, 2025; Vastola, 2025; Bonnaire
051 et al., 2025); see also further discussion in App. A.

052 **Our approach.** We analyze this phenomenon through the lens of random matrix theory (RMT),
053 beginning with the observation that the consistency effect can already be predicted by a linear



070
071
072
073
074
075
076
077
078
079
080

Figure 1: Motivating observation and the linear theory. **A.** Diffusion models trained on non-overlapping data splits generate similar images from the same initial noise, even with different neural network architectures, consistent with results in Kadkhodaie et al. (2024); Zhang et al. (2024). Notably, generated samples from both splits are visually similar to the prediction from the Gaussian linear theory (Wang & Vastola, 2024b). **B.** Quantification of **A** by paired image distances (MSE) averaging from 512 initial noises. The low-MSE block structure of the four DNNs and linear solution emphasize that this consistency effect is related to the linear structure. CNN1 denotes the CNN trained on split1, similar for CNN2, DiT1, DiT2; CNN1 nearest denotes the set of closest training set sample for the 512 generated image. We hide results for linear predictor of two splits since their samples are nearly identical with the linear predictor for the full dataset. Similar analysis for FFHQ64 is showed in Fig. 6.

081 Gaussian model (Fig. 1). Building on the linear denoiser framework, we develop a precise RMT 082 analysis of how finite-sample variability in the empirical covariance affects both the expectation 083 and fluctuation of denoisers and sampling maps. We then validate these theoretical predictions 084 against deep diffusion models (CNNs and DiTs), showing that the same RMT principles still govern 085 their inhomogeneity of consistency across data splits. Our **main contributions** are as follows:

086

- **Linear origin of consistency:** show that shared Gaussian statistics i.e. linear denoiser already 087 predict cross-split agreement.
- **Finite-sample RMT:** prove that randomness enters through a renormalized noise scale $\sigma^2 \mapsto$ 088 $\kappa(\sigma^2)$, explaining overshrinkage of low-variance modes.
- **Variance law:** derive a factorized form for cross-split fluctuations—anisotropy across eigenmodes, 089 inhomogeneity across inputs, and global scaling with n .
- **Fractional-power DE:** extend deterministic equivalence to fractional matrix powers, enabling 090 analysis of full sampling trajectories.
- **Deep-net validation:** qualitatively confirm overshrinkage, anisotropy, and inhomogeneity phe- 091 nomenon in UNet and DiT models beyond the linear regime.

2 NOTATION AND SET UP

100 **Score-based Diffusion Models** Let $p_0(\mathbf{x})$ be the target data distribution. For each noise scale 101 $\sigma > 0$, define the noised distribution as $p(\mathbf{x}; \sigma) = (p_0 * \mathcal{N}(0, \sigma^2 \mathbf{I}))(\mathbf{x}) = \int p_0(\mathbf{y}) \mathcal{N}(\mathbf{x} | \mathbf{y}, \sigma^2 \mathbf{I}) d\mathbf{y}$. 102 The corresponding *score function* is $\nabla_{\mathbf{x}} \log p(\mathbf{x}; \sigma)$, i.e. the gradient of the log-density. In the EDM 103 formulation (Karras et al., 2022), the probability flow ODE (PF-ODE) reads,

$$\frac{d\mathbf{x}}{d\sigma} = -\sigma \nabla_{\mathbf{x}} \log p(\mathbf{x}; \sigma) \quad (\text{PF})$$

104
105
106
107 This ODE transports samples from $p(\cdot; \sigma_2)$ to $p(\cdot; \sigma_1)$ when integrating σ from σ_2 to σ_1 . In particular, by starting from Gaussian noise $\mathcal{N}(0, \sigma_T^2 I)$ and integrating the PF-ODE from a sufficiently

large σ_T down to $\sigma = 0$, one recovers clean samples from p_0 . We adopt the EDM parametrization for its notational simplicity; other common diffusion formalisms are equivalent up to simple rescalings of time and space (Karras et al., 2022).

To estimate the score function of distribution $p_0(\mathbf{x})$, we minimize the denoising score matching (DSM) objective (Vincent, 2011) with a function approximator. We reparametrize the score function via a ‘denoiser’ $\mathbf{s}_\theta(\mathbf{x}, \sigma) = (\mathbf{D}_\theta(\mathbf{x}, \sigma) - \mathbf{x})/\sigma^2$, then at noise level σ the DSM objective becomes

$$\mathcal{L}_\sigma = \mathbb{E}_{\mathbf{x}_0 \sim p_0, \mathbf{z} \sim \mathcal{N}(0, \mathbf{I})} \left\| \mathbf{D}_\theta(\mathbf{x}_0 + \sigma \mathbf{z}; \sigma) - \mathbf{x}_0 \right\|_2^2. \quad (\text{DSM})$$

In practice (Karras et al., 2022), diffusion models balance these scale-specific objectives with a weighting function $w(\sigma)$, yielding the overall training loss $\mathcal{L} = \int_\sigma d\sigma w(\sigma) \mathcal{L}_\sigma$.

Data distribution. Consider a ground truth data distribution $p_0(\mathbf{x})$, $\mathbf{x} \in \mathbb{R}^d$, with population mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$. From this ground truth distribution, we construct an empirical distribution $\{\mathbf{x}_i\}$ with n samples, stacked as $\mathbf{X} \in \mathbb{R}^{n \times d}$, then we denote the empirical mean $\hat{\boldsymbol{\mu}}$ and covariance $\hat{\boldsymbol{\Sigma}}$.

Here we are interested in the effect of the number of samples n , and different realizations of \mathbf{X} on the expectation (mean) and fluctuation (variance) of learned diffusion model. More specifically, we will study how randomness in the empirical covariance $\hat{\boldsymbol{\Sigma}}$ drives variability in the denoiser, relative to the population covariance $\boldsymbol{\Sigma}$.

Linear Denoiser A tractable setting for analytical study is the linear denoiser

$$\mathbf{D}(\mathbf{x}; \sigma) = \mathbf{W}_\sigma \mathbf{x} + \mathbf{b}_\sigma, \quad (1)$$

which is an affine function of the noised state, independent across noise scales. As in linear regression, the training data enters the learned denoiser only through their first two moments (Wang & Pehlevan, 2025; Hastie et al., 2019). More explicitly, minimizing DSM \mathcal{L}_σ for the empirical dataset $p_0 = \{\mathbf{x}_i\}$ ¹ yields the optimal empirical linear denoiser, depending on $\hat{\boldsymbol{\mu}}$, $\hat{\boldsymbol{\Sigma}}$.

$$\mathbf{D}_{\hat{\boldsymbol{\Sigma}}}^*(\mathbf{x}; \sigma) = \hat{\boldsymbol{\mu}} + (\hat{\boldsymbol{\Sigma}} + \sigma^2 \mathbf{I})^{-1} \hat{\boldsymbol{\Sigma}}(\mathbf{x} - \hat{\boldsymbol{\mu}}) \quad (2)$$

For simplicity, we will later set $\hat{\boldsymbol{\mu}} = \boldsymbol{\mu}$ to isolate the effect of the empirical covariance $\hat{\boldsymbol{\Sigma}}$.

Sampling trajectory and sampling map. Given an initial noise pattern $\mathbf{x}_{\sigma_T} \sim \mathcal{N}(0, \sigma_T^2 \mathbf{I})$, the PF-ODE evolves it to a final sample \mathbf{x}_0 . We refer to this mapping from \mathbf{x}_{σ_T} to \mathbf{x}_0 as the *sampling map*; the phenomenon of consistency is precisely about the stability of this mapping across different realizations of training data. When the denoiser is linear and optimal at each noise scale, the PF-ODE can be solved in closed-form by projecting onto the eigenbasis of the data, yielding the analytic sampling trajectory (Wang & Vastola, 2024b; Pierret & Galerne, 2024).

$$\mathbf{x}_{\hat{\boldsymbol{\Sigma}}}(\mathbf{x}_{\sigma_T}, \sigma) = \hat{\boldsymbol{\mu}} + (\hat{\boldsymbol{\Sigma}} + \sigma^2 \mathbf{I})^{1/2} (\hat{\boldsymbol{\Sigma}} + \sigma_T^2 \mathbf{I})^{-1/2} (\mathbf{x}_{\sigma_T} - \hat{\boldsymbol{\mu}}) \quad (3)$$

Taking $\sigma \rightarrow 0$ recovers the Wiener filter with Gaussian prior (Wiener, 1964), which has been shown to be a strong predictor of the sampling map of the learned diffusion networks (Wang & Vastola, 2024b; Lukoianov et al., 2025). In the linear case, the mapping remains affine in the initial state, with the matrix $\hat{\boldsymbol{\Sigma}}^{1/2} (\hat{\boldsymbol{\Sigma}} + \sigma_T^2 \mathbf{I})^{-1/2}$ emerging as the central object of analysis.

3 MOTIVATING EMPIRICAL OBSERVATION

We begin with a simple experiment illustrating the consistency phenomenon. We train UNet-CNN (Song & Ermon, 2019) and DiT (Peebles & Xie, 2023) diffusion models under the EDM framework (Karras et al., 2022), each on two non-overlapping splits of FFHQ32 (30k images each; details in App. D.3). When sampling from the same noise seed with a deterministic solver, the outputs are visually similar across both splits and architectures (Fig. 1A). Quantification via pixel MSE confirms this effect: generated images are more similar across splits than to their nearest neighbors in the training set (Fig. 1B), ruling out memorization (Kadkhodaie et al., 2024; Zhang et al., 2024).

¹With n samples, we average over infinite noise draws, so each sample is reused infinitely.

Strikingly, the linear Gaussian predictor (Wiener filter) (Wang & Vastola, 2024b) already accounts for much of this behavior. Using the empirical mean and covariance $(\hat{\mu}, \hat{\Sigma})$ of each split in Eq. 3, the linear predictor yields nearly identical outputs across splits, also sharing visual similarities with CNN and DiT results (Fig. 1A,B). This suggests that consistency arises because different data splits share nearly identical Gaussian statistics, the only feature the linear diffusion can absorb (Wang & Pehlevan, 2025). Pointwise, samples nearer to the Gaussian solution are also more consistent across splits (Pearson $r = 0.244, p = 5 \times 10^{-15}$), suggesting convergence toward the Gaussian predictor underlies consistency. More visual examples and quantitative comparisons for other datasets (CIFAR10, CIFAR100, FFHQ at 32 and 64 pixels, LSUN church and bedroom dataset at 32 and 64 pixels) are provided in Appendix B.1.

In summary, (i) diffusion models trained on independent splits converge to nearly identical sampling maps, (ii) this property holds across architectures, and (iii) a simple Gaussian predictor already captures much of the effect. While linear diffusion is more consistent than deep networks—which can exploit higher-order statistics—it provides a necessary baseline: if Gaussian statistics differ, deep models may not yield consistent samples. These observations motivate our random matrix theory analysis of finite-sample effects.

4 THEORY OF DIFFUSION CONSISTENCY ACROSS INDEPENDENT DATA

The goal of the study is to calculate the expectation and covariance of various quantities in diffusion model under independent instantiation of dataset.

4.1 SELF CONSISTENCY EQUATION AND RENORMALIZED NOISE SCALE

Deterministic equivalence of sample covariance Our central technical tool is deterministic equivalence (Potters & Bouchaud, 2020; Bun et al., 2015), which allows random matrices to be replaced by deterministic surrogates—an approximation that becomes exact in the large-dimensional limit. In particular, we rely on the deterministic equivalence relation for the empirical covariance matrix $\hat{\Sigma}$ (Atanasov et al., 2024b; Bach, 2024),

$$\hat{\Sigma}(\hat{\Sigma} + \lambda I)^{-1} \asymp \Sigma(\Sigma + \kappa(\lambda)I)^{-1} \quad (4)$$

where κ is the unique positive solution to the self-consistent equation (Silverstein, 1995; Marchenko & Pastur, 1967).

$$\kappa(\lambda) - \lambda = \gamma \kappa(\lambda) \int_0^\infty \frac{sd\mu(s)}{\kappa(\lambda) + s} = \gamma \kappa(\lambda) \text{tr}[\Sigma(\Sigma + \kappa(\lambda)I)^{-1}] \quad (5)$$

where $\gamma = d/n$ is the aspect ratio, and μ is the (limiting) spectral measure of Σ .² Note we use tr to denote the *normalized trace*, such that $\text{tr}[I] = 1$, and Tr the unnormalized one. More elaborate two-point deterministic equivalences (Bach, 2024; Atanasov et al., 2024a; 2025) are required to derive the variance results in the paper, which can be found in Appendix C.1.

Property of renormalized noise $\kappa(\sigma^2)$ As Eq. 4 suggests, with trace-like measurement, the stochastic effects of sample covariance $\hat{\Sigma}$ can be absorbed into the scalar $\kappa(\lambda)$ leaving the population covariance Σ otherwise unchanged, similar to the renormalization of self-energy in field theory (Atanasov et al., 2024b; Hastie et al., 2019; Bach, 2024). In our context, λ usually corresponds to noise variance σ^2 , so we could understand κ as the renormalized noise variance. To build intuition, we numerically evaluate this nonlinear mapping using the spectrum of natural images (FFHQ) (Fig. 2 A, Method in D.1). The renormalization effect $\kappa(\sigma^2)$ is most pronounced at low noise scales, and when the sample number is much fewer than the data dimension ($\gamma = d/n \gg 1$).

Notation Per conventions, we define $\text{df}_1(\lambda) := \text{Tr}[\Sigma(\Sigma + \lambda I)^{-1}]$, $\text{df}_2(\lambda) := \text{Tr}[\Sigma^2(\Sigma + \lambda I)^{-2}]$, $\text{df}_2(\lambda, \lambda') := \text{Tr}[\Sigma^2(\Sigma + \lambda I)^{-1}(\Sigma + \lambda' I)^{-1}]$. We have $\min(n, d) > \text{df}_2(\lambda) > \text{df}_1(\lambda) \geq 0$.

²We write $A_n \asymp B_n$ for *deterministic equivalence*: for any sequence of deterministic matrices C_n with uniformly bounded spectral norm, $\text{tr}[C_n(A_n - B_n)] \rightarrow 0$ as $d, n \rightarrow \infty$, $d/n \rightarrow \gamma$. Equivalences of scalar trace expressions are denoted similarly with \asymp .

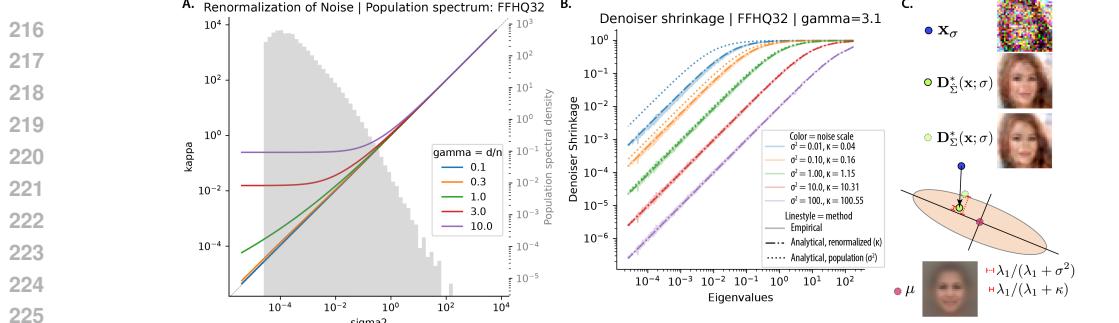


Figure 2: **Renormalization of noise and its effect on expectation of linear denoiser.** **A.** The relationship between the renormalized and raw noise variance $\kappa(\sigma^2)$ as a function of $\gamma = d/n$, using the empirical spectrum of FFHQ32 as the limiting spectrum (plot underneath). See D.1 for numerical methods. **B.** Shrinkage factor of linear denoiser along population eigenvectors at different noise scales. Empirical shows $\mathbf{v}^\top \hat{\Sigma}(\hat{\Sigma} + \sigma^2 I)^{-1} \mathbf{v}$, when $\mathbf{v} = \mathbf{u}_k$ population PCs, at dataset size $n = 1000$, $\gamma \approx 3.1$. **C.** Schematics showing the overshrinking effect at lower eigenspaces, using linear denoiser outcome of faces as example.

4.2 EXPECTATION: FINITE DATA RENORMALIZES NOISE SCALES

Next we apply these tools to compute the expectation and fluctuation of the denoiser under dataset realizations. The form of Eq. 2 naturally suggests the deterministic equivalence in Eq. 4, leading to the following result.

Proposition 1 (Deterministic equivalence of the denoiser expectation). *Assuming $\hat{\mu} = \mu$, and given a fixed probe vector $\mathbf{v} \in \mathbb{R}^d$, then the optimal empirical linear denoiser has the following deterministic equivalence. (Proof in App. C.2).*

$$\mathbb{E}_{\hat{\Sigma}} \left[\mathbf{v}^\top \mathbf{D}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma) \right] \asymp \mathbf{v}^\top \mathbf{D}_{\Sigma}^*(\mathbf{x}; \kappa^{1/2}(\sigma^2)) = \mathbf{v}^\top \left[\mu + \Sigma(\Sigma + \kappa(\sigma^2)I)^{-1}(\mathbf{x} - \mu) \right] \quad (6)$$

Interpretation In expectation, finite data act by renormalizing the noise scale, $\sigma^2 \rightarrow \kappa(\sigma^2)$, in the population denoiser. This is equivalent to adding an adaptive Ridge penalty to the DSM objective (Eq. DSM). Compared to the population solution \mathbf{D}_{Σ}^* , the finite-sample denoiser shrinks low-variance directions more aggressively, treating them as noise and pulling outputs toward the dataset mean (Fig. 2C). Numerically, deviations are indeed most pronounced in the lower spectrum and at lower noise levels, where the renormalization effect is the strongest (Fig. 2B). Since smaller noise scale is associated with generation of high frequency details in image, this result suggests these detail eigenmodes take more samples to be learned correctly, which we'll confirm in next section.

4.3 FLUCTUATION: ANISOTROPIC AND INHOMOGENEITY OF DENOISER CONSISTENCY

Next, we tackle the fluctuation due to dataset realizations, which addresses the consistency of diffusion models trained on independent data splits. We prove the following equivalence using two-point and one-point deterministic equivalence identities (Eq. 19, 17, Bach (2024)).

Proposition 2 (Deterministic equivalence of the denoiser variance). *Assuming $\hat{\mu} = \mu$, across dataset realizations of size n , the variance of the optimal empirical linear denoiser at point \mathbf{x} in direction \mathbf{v} , given by $\mathbf{v}^\top \mathcal{S}_D(\mathbf{x}) \mathbf{v}$, admits the following deterministic equivalence. Proof in App. C.3.*

$$\begin{aligned} \mathbf{v}^\top \mathcal{S}_D(\mathbf{x}) \mathbf{v} &= \text{Var}_{\hat{\Sigma}} [\mathbf{v}^\top \mathbf{D}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma)] \quad (7) \\ &\asymp \frac{\kappa(\sigma^2)^2}{n - \text{df}_2(\kappa(\sigma^2))} \underbrace{\left(\mathbf{v}^\top (\Sigma + \kappa(\sigma^2)I)^{-2} \Sigma \mathbf{v} \right)}_{\text{anisotropy: } \square(\mathbf{v}, \kappa, \Sigma)} \underbrace{\left((\mathbf{x} - \mu)^\top (\Sigma + \kappa(\sigma^2)I)^{-2} \Sigma (\mathbf{x} - \mu) \right)}_{\text{inhomogeneity: } \square(\mathbf{x} - \mu, \kappa, \Sigma)} \end{aligned}$$

Interpretation The variance of denoiser across dataset realizations factorizes into three interpretable components: a dependence on probe direction (*anisotropy*), a dependence on noised sample location (*inhomogeneity*), and an overall scale with n and σ (*global scaling*). Note, given the relation between score and denoiser, the score variance is $\sigma^{-4} \mathbf{v}^\top \mathcal{S}_D(\mathbf{x}) \mathbf{v}$, i.e. all results translate by scaling.

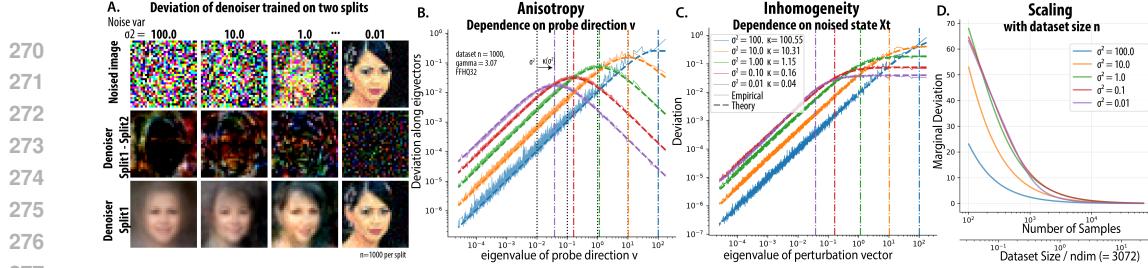


Figure 3: **Structure of denoiser deviation across dataset splits.** **A.** Visual examples of linear denoisers trained on two disjoint splits of FFHQ32 as noise variance σ^2 decreases, $n = 1000$. Top, \mathbf{x}_t noised sample; Bottom, output of linear denoiser (trained on split 1) $\mathbf{D}_{\hat{\Sigma}_1}(\mathbf{x}_t, \sigma)$; Middle, deviation between two denoisers (normalized) $\mathbf{D}_{\hat{\Sigma}_1}(\mathbf{x}_t, \sigma) - \mathbf{D}_{\hat{\Sigma}_2}(\mathbf{x}_t, \sigma)$. At high noise, denoisers diverge on global, low-frequency content; at low noise, they deviate at specular details. **B. Anisotropy:** variance depends on probe direction \mathbf{v} ; deviation is maximized when the eigenvalue λ_k of \mathbf{v} matches the renormalized noise $\kappa(\sigma^2)$, in agreement with theory. **C. Inhomogeneity:** variance depends on probe location \mathbf{x}_t ; samples displaced along high-variance eigenmodes induce larger deviations. **D. Global scaling:** marginal deviation decays with dataset size n , vanishing in the infinite-sample limit.

Anisotropy in probe direction. The anisotropy of consistency is governed by $\square(\mathbf{v}, \kappa, \Sigma)$. When the probe \mathbf{v} aligns with a principal component (PC) \mathbf{u}_k of Σ with eigenvalue λ_k , this reduces to $\chi(\lambda_k, \kappa) := \lambda_k / (\lambda_k + \kappa)^2$. The function $\chi(\lambda, \kappa)$ is bell-shaped in λ , uniquely maximized at $\lambda = \kappa$ with value $1/(4\kappa)$. Thus, for each noise scale, the directions of greatest uncertainty are precisely those whose variances match the renormalized noise $\kappa(\sigma^2)$ (Fig. 3B). This effect is evident visually. For linear denoisers trained on non-overlapping splits of human face dataset (FFHQ), their differences follow the spectral structure of natural images (Ruderman, 1994): at high noise the deviations appear as low-frequency facial envelopes, while at low noise they shift to high-frequency specular patterns (Fig. 3B). Quantitatively, the MSE between two denoisers along each PC matches the variance prediction of Eq. 7, with the expected factor of two from independent sampling (Lemma 1).

Inhomogeneity in input location. The inhomogeneity of denoiser variance across input space is governed by $\square(\mathbf{x} - \mu, \kappa, \Sigma)$. While structurally similar to the anisotropy factor, here $\mathbf{x} - \mu$ is drawn from the noised data distribution rather than a unit probe. Approximating $\mathbf{x} - \mu$ as lying on the ellipsoidal shell of $\mathcal{N}(0, \Sigma + \sigma^2 I)$, its displacement along eigenvector \mathbf{u}_k has typical radius $\sqrt{\sigma^2 + \lambda_k}$. Substituting gives $\square(\sqrt{\sigma^2 + \lambda_k} \mathbf{u}_k, \kappa, \Sigma) = (\sigma^2 + \lambda_k) \chi(\lambda_k, \kappa)$. Unlike the pure anisotropy factor, this expression grows monotonically with λ_k . Thus, denoiser variability is amplified for inputs displaced along high-variance modes, yielding larger uncertainty for such locations (Fig. 3C), which agree quantitatively with numerical results. Based on this factor, denoiser consistency can be predicted for noisy input point by point (e.g. Pearson $r = 0.94$ across noised images, at $\sigma^2 = 1$, $n = 1000$, Fig. 21).

Global scaling with sample size. Finally, marginalizing over all directions and noised samples yields a closed-form expression for the overall denoiser variance (Eq. 22, Fig. 3D). At large n limit, denoiser variance scale inversely with sample number n^{-1} , reminiscent of classic statistical laws; while at smaller n , the renormalization effects modify the scaling.

Summary. In sum, the variance structure reveals three key effects. *Anisotropy*: uncertainty is maximized along eigenmodes whose variance λ_k is comparable to the renormalized noise $\kappa(\sigma^2)$. *Inhomogeneity*: noised points displaced along high-variance directions experience larger uncertainty. *Scaling*: the overall variance shrinks with dataset size n , recovering the population model in the large-sample limit. Together, these predictions yield a detailed spatial and spectral map of where denoisers trained on different data splits are most likely to disagree.

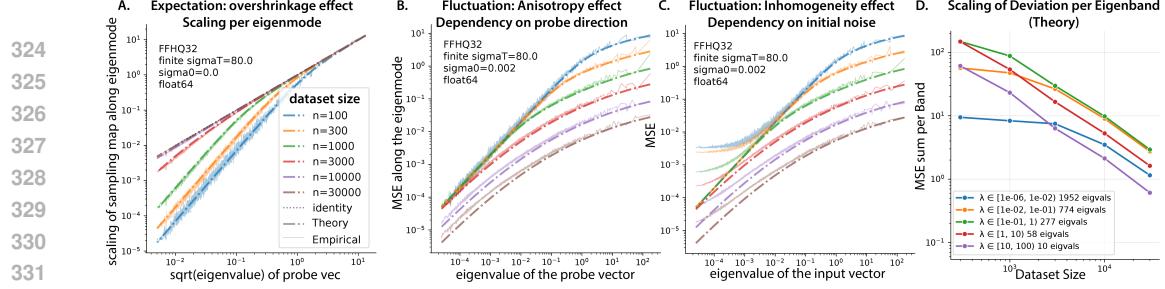


Figure 4: Finite sample effect on diffusion sampling map. A. Overshrinkage of expectation. Expected scaling along eigenmode of the empirical sampling map $\mathbf{u}_k^\top \hat{\Sigma}^{1/2} \mathbf{u}_k$ compared to the ideal $\sqrt{\lambda_k}$, showing overshrinking along lower eigenmodes. **B. Anisotropy of consistency.** Cross-split MSE depends on probe direction \mathbf{v} , with larger deviation on top eigenspaces. **C. Inhomogeneity of consistency.** Cross-split MSE depends on input location $\bar{\mathbf{x}}$; samples displaced along high-variance modes exhibit larger disagreement. Colors denote dataset size, shared across A, B, C. **D. Scaling of consistency by eigenband.** Decomposition of MSE across eigenbands shows that lower-variance modes require substantially more samples before cross-split MSE decays. See also Fig. 22.

5 CONSISTENCY OF DIFFUSION SAMPLES FOR LINEAR DENOISERS

Beyond the consistency of single-step denoiser output or score, we are interested in the final diffusion sample from the same initial noise seed \mathbf{x}_{σ_T} . For linear denoisers, sampling map from initial noise to generated sample is captured by Wiener filter (Eq. 3, $\sigma = 0$). However, unlike one-step denoiser, this mapping involves fractional power of covariances $\Sigma^{1/2}(\Sigma + \sigma^2 I)^{-1/2}$, for which the deterministic equivalence is not readily available. Here, we leveraged the integral representation of fractional power (Balakrishnan (1960)’s formula) and deterministic equivalence, and arrived at a few novel equivalence of these matrices (Prop. 6, 8, Proof in App. C.4). Using these developments, we can calculate the expectation and fluctuation of sampling map.

5.1 EXPECTATION OF DIFFUSION SAMPLE: OVER-SHRINKAGE TO THE MEAN

We note that when the initial noise scale σ_T is large, the sampling map admits the approximation

$$\mathbf{x}_{\hat{\Sigma}}(\mathbf{x}_{\sigma_T}, 0) = \boldsymbol{\mu} + \hat{\Sigma}^{1/2}(\hat{\Sigma} + \sigma_T^2 I)^{-1/2}(\mathbf{x}_{\sigma_T} - \boldsymbol{\mu}) \approx \boldsymbol{\mu} + \hat{\Sigma}^{1/2}\bar{\mathbf{x}} \quad (8)$$

where we define the shift and normalized noise $\bar{\mathbf{x}} := \frac{\mathbf{x}_{\sigma_T} - \boldsymbol{\mu}}{\sigma_T}$. At the $\sigma_T \rightarrow \infty$ limit, this approximation becomes exact, and $\bar{\mathbf{x}} \sim \mathcal{N}(0, I)$. For clarity, we present results under this infinite- σ_T approximation; the expressions accounting for finite σ_T effects are provided in App. C.6.

Proposition 3 (Deterministic equivalence for expectation of diffusion sampling map). *The sample generated from initial state \mathbf{x}_{σ_T} has the following deterministic equivalence. Proof in App. C.5.*

$$\mathbf{x}_{\hat{\Sigma}}[\mathbf{x}_{\hat{\Sigma}}(\mathbf{x}_{\sigma_T}, 0)] \approx \boldsymbol{\mu} + \mathbb{E}_{\hat{\Sigma}}[\hat{\Sigma}^{1/2}] \frac{\mathbf{x}_{\sigma_T} - \boldsymbol{\mu}}{\sigma_T} \asymp \boldsymbol{\mu} + \frac{2}{\pi} \int_0^\infty \Sigma \left(\Sigma + \kappa(u^2) I \right)^{-1} \bar{\mathbf{x}} du \quad (9)$$

Interpretation This expression mirrors the deterministic equivalence of denoisers (Eq. 6), but with an integration over effective noise scales. Comparing to the population sampling map, where $\kappa(u^2)$ reduces to u^2 , the finite data case integrates over a stronger shrink factor $\Sigma(\Sigma + \kappa I)^{-1}$ (since $\kappa(u^2) > u^2$), especially on the lower eigenmodes. This effect is confirmed with numerics of empirical covariance (Fig. 4A). This leads to a systematic overshrinkage toward the dataset mean along these modes, reducing the generated variance along lower-variance directions.³

5.2 VARIANCE OF DIFFUSION SAMPLE: ANISOTROPY AND INHOMOGENEITY

Proposition 4 (Deterministic equivalence for variance of diffusion sampling map). *Due to dataset realization, the variance of generated sample starting from initial state \mathbf{x}_{σ_T} , along vector \mathbf{v} admits*

³Note that, though the sample covariance $\hat{\Sigma}$ is an unbiased estimator of the population covariance Σ , taking the square root introduces this finite sample bias, i.e., $\Sigma = \mathbb{E}[\hat{\Sigma}] = \mathbb{E}[\hat{\Sigma}^{1/2} \hat{\Sigma}^{1/2}] \neq (\mathbb{E}[\hat{\Sigma}^{1/2}])^2$.

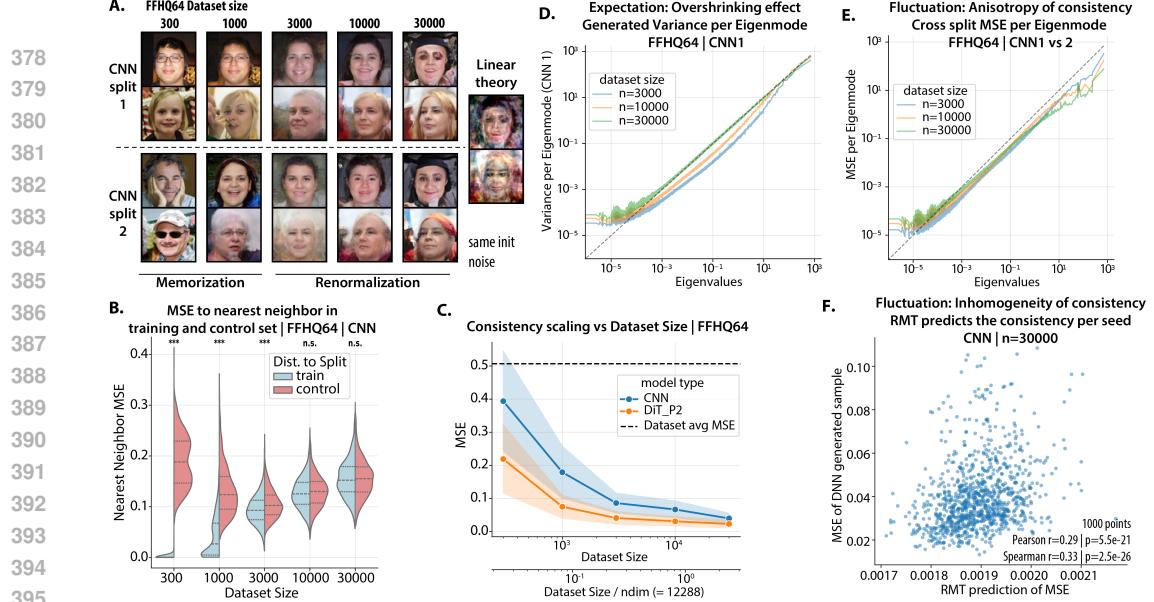


Figure 5: **DNN validation of theory.** **A.** Samples generated by UNet (same two seeds) across training set sizes and splits (FFHQ64); similarity increases with n , and increasingly matches the population linear predictor (right). **B.** Nearest-neighbor MSE in training vs. control sets reveals memorization at small n , $n > 3000$ shows no statistical difference between the splits. **C.** Overall consistency improve as a function of dataset size, with DiT more consistent than UNet at each n (cross split MSE, mean \pm std). **D.** Variance of generated samples per eigenmode highlight insufficient variance (*overshrinkage*) in mid-to-low eigenmodes with limited dataset size. **E.** Cross-split MSE per eigenmode shows *anisotropy* of consistency (Fig. 4B). Further, per dataset size, deviation in top eigenmodes decrease the most. **F.** In the renormalization regime ($n = 30k$), RMT predictions of seed-wise consistency correlate with empirical deviations (Spearman $r = 0.33$).

the following deterministic equivalence,

$$\begin{aligned} \text{Var}_{\hat{\Sigma}}[\mathbf{v}^\top \mathbf{x}_{\hat{\Sigma}}(\mathbf{x}_{\sigma_T}, 0)] &= \text{Var}_{\hat{\Sigma}}[\mathbf{v}^\top \hat{\Sigma}^{1/2} \bar{\mathbf{x}}] \\ &\asymp \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \frac{\kappa \kappa'}{n - \text{df}_2(\kappa, \kappa')} \underbrace{\triangle(\mathbf{v}; \kappa, \kappa', \Sigma)}_{\text{anisotropy}} \underbrace{\triangle(\bar{\mathbf{x}}; \kappa, \kappa', \Sigma)}_{\text{inhomogeneity}} du dv, \end{aligned} \quad (10)$$

where $\triangle(\mathbf{a}; \kappa, \kappa', \Sigma) := \mathbf{a}^\top \Sigma (\Sigma + \kappa I)^{-1} (\Sigma + \kappa' I)^{-1} \mathbf{a}$, and $\kappa := \kappa(u^2), \kappa' := \kappa(v^2)$ are variables to be integrated over. Proof in App. C.7.

Interpretation The variance of sampling map Eq. 10 simplifies to a double integral of the denoiser-variance (Eq. 7). The integrand factorizes into a direction-dependent term (*anisotropy*), a initial noise-dependent term (*inhomogeneity*), and a scaling term. Note the anisotropy and inhomogeneity factors rely on the same $\triangle(\cdot; \kappa, \kappa', \Sigma)$ function, showing that dependency on \mathbf{v} and $\bar{\mathbf{x}}$ has the same spectral structure.

We resort to numerical simulation to provide more intuition. We note that integrals in Eqs. 9, 10 are nontrivial to evaluate; we describe our numerical scheme in App. D.1. Using this procedure, the theoretical predictions align closely with direct computations of linear diffusion (Fig. 4). **Inhomogeneity** Spatially, when initial noise $\bar{\mathbf{x}}$ deviates more along the top eigenspace of Σ , there will be larger uncertainty (Fig. 4C), this enables us to predict the sample difference point by point. **Anisotropy** Directionally, the dependency on \mathbf{v} has the same structure, in absolute term, the deviation is larger at higher eigenspace (Fig. 4B). Note that when comparing across the dataset size, the variance in the top eigenspace decay immediately from small sample size; while the deviation in mid to lower eigenspace will stay put and start decaying only later at larger dataset size (Fig. 4D). This shows that the fine detail of the samples needs a larger dataset size to be consistency across training.

432 6 VALIDATING PREDICTIONS ON DEEP NETWORKS

434 Finally, given that linear diffusion behavior is well captured by our random matrix theory (RMT), we
 435 test the applicability of its prediction to practical deep diffusion networks.

437 **Setup.** We trained UNet- and DiT-based denoisers under the EDM framework on FFHQ64, FFHQ32,
 438 AFHQ32 (Choi et al., 2020), and CIFAR10. For each dataset we trained on two non-overlapping
 439 splits at sizes $n = \{300, 1000, 3000, 10^5, 3 \times 10^5\}$ (10 runs total per architecture). Sampling was
 440 performed with the same random seed using the Heun solver (Karras et al., 2022). We train for
 441 50,000 steps with Adam optimizer, further details are provided in App. D.3.

442 **Expectation: from memorization to renormalization.** We observe a clear two-phase behavior as
 443 dataset size increases. *Memorization phase* ($n \leq 1000$): models largely reproduce training samples
 444 (Fig. 5A,B), and samples are much closer to the nearest neighbor in their training split than the control
 445 split, consistent with prior observations. This regime is outside the scope of linear theory, since
 446 linear score models cannot memorize individual points (Wang & Pehlevan, 2025). *Renormalization*
 447 *phase* ($n \geq 3000$): the samples have comparable distance to the neighbor in the training split and
 448 control split, showing generalization. Further, samples begin to resemble the linear predictors (Li
 449 et al., 2024b). In this regime, the overshrinkage predicted by Prop. 3 becomes visible: generated face
 450 samples resemble the average face (Langlois et al., 1994), with smoother textures and background
 451 (Fig. 5A, $n = 3000$). Quantitatively, we observe reduced variance along low- and mid-spectrum
 452 eigenmodes of the generated samples (Fig. 5 D). This bias decreases as dataset size increases, and
 453 vanishes when learned and population spectra coincide at $n \sim 30000$. The same transition occurs
 454 across architectures, though the dataset size at which it occurs depends on model capacity and image
 455 resolution (Fig. 23,25,24).

456 **Fluctuations: inhomogeneity of consistency.** Within the renormalization phase, RMT further
 457 predicts which initial noise and direction exhibit the largest discrepancies across data splits, due to
 458 their alignment with data covariance (Eq. 4). Spectrally, measuring the cross-split deviation along
 459 population eigenbases, we can see the characteristic anisotropy profile. Further the decrease of
 460 MSE majorly occurs in top eigenspace, while the middle or lower eigenspace remains unchanged or
 461 becomes less consistent when sample size increases (Fig. 5 E). This is consistent with the prediction of
 462 the theory that lower eigenmodes need more training samples to be consistent (Fig. 4B). Spatially, the
 463 inhomogeneity effect is borne out: RMT predictions correlate with observed cross-split deviations for
 464 each initial noise point by point; e.g., UNets trained on FFHQ64 with $n = 30000$ achieve a Spearman
 465 correlation of 0.33 ($p = 2.5 \times 10^{-26}$) over 1000 seeds (Fig. 5F). Remarkably, the prediction requires
 466 only the population covariance and dataset size, with no knowledge of split identities or network
 467 architecture. The absolute deviation magnitudes, however, are much larger in deep networks than
 468 predicted by linear theory, reflecting nonlinear source of variability. As controls, correlations collapse
 469 in the memorization regime (Fig. 42,43) and disappear when mismatched noise seeds are used.

470 **Summary.** Across architectures and datasets, the predictions of our linear RMT framework extend
 471 to deep diffusion models: limited data induce overshrinkage toward the mean, and the variance
 472 structure across splits exhibits inhomogeneity and anisotropy predicted by theory.

475 7 DISCUSSION

477 Our analysis shows that much of the consistency in diffusion models across training data is already
 478 captured by Gaussian statistics: if two data splits share their first two moments, the corresponding
 479 sampling maps nearly coincide. Random matrix theory sharpens this picture by showing that finite
 480 data act through a renormalized noise scale $\sigma^2 \mapsto \kappa(\sigma^2)$, and that fluctuations across splits factor into
 481 anisotropy over eigenmodes, inhomogeneity across inputs, and a global scaling with n . These results
 482 extend deterministic-equivalence tools to fractional matrix powers, allowing closed-form predictions
 483 for both denoisers and sampling trajectories, and align well with deep networks in terms of where
 484 deviation accentuates, even if nonlinear effects amplify the magnitudes.

485 At the same time, our framework has limitations. Linear surrogates underestimate variability in
 expressive models and do not capture architecture-specific inductive biases. Extending the theory

486 to random-feature models or mild non-Gaussian structure would better explain the transition from
487 memorization to renormalization (Bonnaire et al., 2025; George et al., 2025), and help quantify how
488 model capacity shifts the required dataset size. Another promising direction is to study the *anisotropy*
489 of the initial noise space and its alignment with the data manifold. The seemingly unstructured
490 noise space is already anchored by the data covariance before generation. Such anchoring might
491 explain why certain “magic” random seeds may consistently yield better generations, e.g. they
492 avoid directions where cross-split disagreement is largest (Xu et al., 2025). This echoes anisotropic
493 effects observed in GANs’ latent space, where noise vectors aligned strongly with top eigenspaces of
494 Jacobian can lead to degraded generations (Wang & Ponce, 2021). Such connections suggest that
495 spectral geometry of the input space deserves closer attention as a unifying factor across generative
496 models.

497 **ACKNOWLEDGMENTS**
498

499 B.W. was supported by the Kempner Fellowship from the Kempner Institute at Harvard. J.Z.-V. was
500 supported by a Junior Fellowship from the Harvard Society of Fellows. C.P. is supported by an
501 NSF CAREER Award (IIS-2239780), DARPA grants DIAL-FP-038 and AIQ-HR00112520041, the
502 Simons Collaboration on the Physics of Learning and Neural Computation, and the William F. Milton
503 Fund from Harvard University. This work has been made possible in part by a gift from the Chan
504 Zuckerberg Initiative Foundation to establish the Kempner Institute for the Study of Natural and
505 Artificial Intelligence, and by the generous computing resources provided by the Kempner Institute
506 and Harvard Research Computing.

507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

540 REFERENCES
541

542 Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying
543 framework for flows and diffusions. *arXiv*, 2023. URL <https://arxiv.org/abs/2303.08797>.

544 Luca Ambrogioni. In search of dispersed memories: Generative diffusion models are associative
545 memory networks, 2023. URL <https://arxiv.org/abs/2309.17290>.

546 Alexander Atanasov, Jacob A. Zavatone-Veth, and Cengiz Pehlevan. Risk and cross validation in
547 ridge regression with correlated samples, August 2024a.

548 Alexander Atanasov, Jacob A Zavatone-Veth, and Cengiz Pehlevan. Scaling and renormalization in
549 high-dimensional regression. *arXiv preprint arXiv:2405.00592*, 2024b.

550 Alexander Atanasov, Blake Bordelon, Jacob A. Zavatone-Veth, Courtney Paquette, and Cengiz
551 Pehlevan. Two-Point Deterministic Equivalence for Stochastic Gradient Dynamics in Linear
552 Models, February 2025.

553 Francis Bach. High-dimensional analysis of double descent for linear regression with random
554 projections. *SIAM Journal on Mathematics of Data Science*, 6(1):26–50, 2024.

555 Zhidong Bai, Jack William Silverstein, et al. *Spectral analysis of large dimensional random matrices*.
556 Springer, 2010.

557 A. V. Balakrishnan. Fractional powers of closed operators and the semigroups generated by them.
558 *Pacific Journal of Mathematics*, 10(2):419–437, January 1960. ISSN 0030-8730.

559 Tony Bonnaire, Raphaël Urfin, Giulio Biroli, and Marc Mézard. Why diffusion models don’t memo-
560 rize: The role of implicit dynamical regularization in training. *arXiv preprint arXiv:2505.17638*,
561 2025.

562 Joël Bun, Romain Allez, Jean-Philippe Bouchaud, and Marc Potters. Rotational invariant estimator
563 for general noisy matrices, February 2015.

564 Zhengdao Chen. On the interpolation effect of score smoothing. *arXiv preprint arXiv:2502.19499*,
565 2025.

566 Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for
567 multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern
568 recognition*, pp. 8188–8197, 2020.

569 Emma Finn, T. Anderson Keller, Manos Theodosis, and Demba E. Ba. Origins of creativity in
570 attention-based diffusion models, 2025. URL <https://arxiv.org/abs/2506.17324>.

571 Anand Jerry George, Rodrigo Veiga, and Nicolas Macris. Denoising score matching with random fea-
572 tures: Insights on diffusion models from precise learning curves. *arXiv preprint arXiv:2502.00336*,
573 2025.

574 Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J. Tibshirani. Surprises in High-
575 Dimensional Ridgeless Least Squares Interpolation, March 2019.

576 Zahra Kadkhodaie, Florentin Guth, Eero P Simoncelli, and Stéphane Mallat. Generalization in diffu-
577 sion models arises from geometry-adaptive harmonic representations. In *The Twelfth International
578 Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=ANvmVS2Yr0>.

579 Mason Kamb and Surya Ganguli. An analytic theory of creativity in convolutional diffusion models.
580 *arXiv e-prints*, art. arXiv:2412.20292, December 2024. doi: 10.48550/arXiv.2412.20292.

581 Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-
582 based generative models. *arXiv preprint arXiv:2206.00364*, 2022.

583 Judith H. Langlois, Lori A. Roggman, and Lisa Musselman. What is average and what is not average
584 about attractive faces? *Psychological Science*, 5(4):214–220, 1994. doi: 10.1111/j.1467-9280.
585 1994.tb00503.x. URL <https://doi.org/10.1111/j.1467-9280.1994.tb00503.x>.

594 Olivier Ledoit and Sandrine Péché. Eigenvectors of some large sample covariance matrix ensembles.
 595 *Probability Theory and Related Fields*, 151(1):233–264, 2011.
 596

597 Sixu Li, Shi Chen, and Qin Li. A good score does not lead to a good generative model, 2024a. URL
 598 <https://arxiv.org/abs/2401.04856>.

599 Xiang Li, Yixiang Dai, and Qing Qu. Understanding generalizability of diffusion models requires
 600 rethinking the hidden gaussian structure. *Advances in neural information processing systems*, 37:
 601 57499–57538, 2024b.

602 Xiang Li, Yixiang Dai, and Qing Qu. Understanding generalizability of diffusion models requires
 603 rethinking the hidden gaussian structure. *arXiv preprint arXiv:2410.24060*, 2024c.

604 Artem Lukoianov, Chenyang Yuan, Justin Solomon, and Vincent Sitzmann. Locality in Image
 605 Diffusion Models Emerges from Data Statistics, September 2025.

606 V. A. Marchenko and L. A. Pastur. Distribution of eigenvalues for some sets of random matrices. *Math-
 607 ematics of the USSR-Sbornik*, 1(4):457–483, 1967. doi: 10.1070/SM1967v001n04ABEH001994.
 608 URL <https://www.mathnet.ru/eng/sm4101>.

609 Miles Martinez and John Pearson. Reproducible, incremental representation learning with Rosetta
 610 VAE, January 2022.

611 Matthew Niedoba, Berend Zwartenberg, Kevin Murphy, and Frank Wood. Towards a Mechanistic
 612 Explanation of Diffusion Model Generalization. *arXiv e-prints*, art. arXiv:2411.19339, November
 613 2024. doi: 10.48550/arXiv.2411.19339.

614 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the
 615 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4195–4205, October 2023.

616 Emile Pierret and Bruno Galerne. Diffusion models for Gaussian distributions: Exact solutions and
 617 Wasserstein errors. *arXiv preprint arXiv:2405.14250*, 2024.

618 Marc Potters and Jean-Philippe Bouchaud. *A First Course in Random Matrix Theory: For Physicists,
 619 Engineers and Data Scientists*. Cambridge University Press, Cambridge, 2020. ISBN 978-1-108-
 620 48808-2. doi: 10.1017/9781108768900.

621 Daniel L Ruderman. The statistics of natural images. *Network: computation in neural systems*, 5(4):
 622 517, 1994.

623 Jack W Silverstein. Strong convergence of the empirical distribution of eigenvalues of large dimen-
 624 sional random matrices. *Journal of Multivariate Analysis*, 55(2):331–339, 1995.

625 Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution.
 626 In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

627 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 628 Poole. Score-based generative modeling through stochastic differential equations. In *International
 629 Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=PxTIG12RRHS>.

630 John Vastola. Generalization through variance: how noise shapes inductive biases in diffusion
 631 models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL
 632 <https://openreview.net/forum?id=71Udo8Vuqa>.

633 Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computa-
 634 tion*, 23(7):1661–1674, 2011.

635 Binxu Wang. An analytical theory of power law spectral bias in the learning dynamics of diffusion
 636 models. *arXiv preprint arXiv:2503.03206*, 2025.

637 Binxu Wang and Cengiz Pehlevan. An Analytical Theory of Spectral Bias in the Learning Dynamics
 638 of Diffusion Models, March 2025. URL <http://arxiv.org/abs/2503.03206v2>.

648 Binxu Wang and Carlos R Ponce. A geometric analysis of deep generative image models and
649 its applications. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=GH7QRzUDdXG>.
650

651 Binxu Wang and John Vastola. The unreasonable effectiveness of gaussian score approximation for
652 diffusion models and its applications. *Transactions on Machine Learning Research*, December
653 2024a. arXiv preprint arXiv:2412.09726.
654

655 Binxu Wang and John Vastola. The unreasonable effectiveness of gaussian score approximation for
656 diffusion models and its applications. *Transactions on Machine Learning Research*, 2024b. ISSN
657 2835-8856. URL <https://openreview.net/forum?id=I0uKnSHM2j>.
658

659 Binxu Wang and John J. Vastola. The Hidden Linear Structure in Score-Based Models and its
660 Application. *arXiv e-prints*, art. arXiv:2311.10892, November 2023. doi: 10.48550/arXiv.2311.
661 10892.

662 Norbert Wiener. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*. The MIT
663 press, 1964.

664 Katherine Xu, Lingzhi Zhang, and Jianbo Shi. Good seed makes a good crop: Discovering secret
665 seeds in text-to-image diffusion models. In *2025 IEEE/CVF Winter Conference on Applications of*
666 *Computer Vision (WACV)*, pp. 3024–3034. IEEE, 2025.
667

668 Huijie Zhang, Jinfan Zhou, Yifu Lu, Minzhe Guo, Peng Wang, Liyue Shen, and Qing Qu. The
669 emergence of reproducibility and consistency in diffusion models. In *Forty-first International Con-*
670 *ference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=Hsli0qZkc0>.
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

702	CONTENTS	
703		
704		
705	1 Introduction	1
706		
707	2 Notation and Set up	2
708		
709	3 Motivating Empirical Observation	3
710		
711	4 Theory of Diffusion Consistency Across Independent Data	4
712		
713	4.1 Self consistency equation and renormalized noise scale	4
714	4.2 Expectation: Finite Data Renormalizes Noise Scales	5
715	4.3 Fluctuation: Anisotropic and Inhomogeneity of Denoiser Consistency	5
716		
717		
718	5 Consistency of Diffusion Samples for Linear Denoisers	7
719		
720	5.1 Expectation of diffusion sample: over-shrinkage to the mean	7
721	5.2 Variance of diffusion sample: Anisotropy and inhomogeneity	7
722		
723	6 Validating Predictions on Deep Networks	9
724		
725		
726	7 Discussion	9
727		
728	A Extended Related Works	15
729		
730		
731	B Extended Results and Figures	16
732	B.1 Extended visual examples for motivating observation	16
733	B.2 Extended validation of the theory	25
734	B.3 Extended evidence from the DNN validation experiments	27
735		
736		
737	C Proof and Derivations	41
738		
739	C.1 Deterministic equivalence relations	41
740	C.2 Proof for Deterministic equivalence of denoiser expectation (proposition 1)	43
741	C.3 Proof for Deterministic equivalence of denoiser fluctuation (proposition 2)	44
742	C.3.1 Interpretation and derivations	45
743	C.4 Integral representation of matrix fractional power (Balakrishnan formula)	49
744	C.5 Proof for expectation of the sampling mapping (approximate version, infinite σ_T , Proposition 3)	51
745	C.6 Proof for expectation of the sampling mapping (full version, finite σ_T)	53
746	C.7 Proof for fluctuation of the sampling mapping (approximate version, infinite σ_T , Proposition 4)	55
747	C.7.1 Interpretation	56
748		
749		
750		
751		
752		
753		
754	D Experimental Details	58
755		
	D.1 Numerical methods	58

756	D.2 Linear denoiser experiments	58
757	D.3 Deep neural network experiments	59
758		
759		
760	E Usage of LLMs	61
761		

A EXTENDED RELATED WORKS

Consistency and Reproducibility in Diffusion As a motivating observation, Kadkhodaie et al. (2024) found that diffusion models trained on non overlapping splits of training data could produce visually highly similar images. The seminal paper studying this effect is Zhang et al. (2024), there, the authors found that different models trained on the same dataset across architecture (transformer vs UNet), across objectives, across training runs, and across sampler and noising kernel, have consistent mapping from noise to sample as long as an ODE deterministic sampler is used. In their appendix B, they also made detailed discussion about lack of reproducibility in VAE and GANs. The consistency studied in our paper is more related to the reproducibility in the generalization regime.

Hidden Linear Score Structure in Diffusion Models Recent work has shown, for much of diffusion times (*i.e.*, signal to noise ratio), the learned neural score is closely approximated by the linear score of a Gaussian fit to the data, which is usually the best linear approximation (Wang & Vastola, 2023; Li et al., 2024c). Crucially, this Gaussian linear score admits a closed-form solution to the probability-flow ODE, which can be exploited to accelerate sampling and improve its quality (Wang & Vastola, 2024a). Moreover, this same linear structure has been linked to the generalization–memorization transition in diffusion models (Li et al., 2024c). In sum, across many noise levels, the Gaussian linear approximation captures many salient aspects of the learned score. Here, we leverage it to explain the observed consistency across splits and as a tractable set up for random matrix theory analysis.

Memorization, Generalization and Creativity in Diffusion The question of when diffusion models are able to generate genuinely novel samples matters both scientifically and for mitigating data leakage. From the score-matching perspective, if the learned score exactly matches that of the empirical data distribution, then the reverse process reproduces that empirical distribution, and thus does not create new samples beyond the training set (Kamb & Ganguli, 2024; Li et al., 2024a; Wang & Vastola, 2024b). Yet high-quality diffusion models routinely generate images that are not identical copies of images from the training set. Kamb & Ganguli (2024) take an important step toward reconciling this: when the score network is a simple CNN, its inductive biases (locality and translation equivariance) favor patch wise composition, enabling global samples that are novel while remaining locally consistent “mosaics.” Similarly, Wang (2025) noticed that score networks with different architectural constraints will learn various approximation of the dataset, and therefore generalize: e.g. linear networks learn the Gaussian approximation, and circular convolutional networks learn the stationary Gaussian process approximation. Finn et al. (2025) provided evidence that adding a final self-attention layer promotes global consistency across distant regions, organizing locally plausible features into coherent layouts that move beyond purely patch-level mosaics. This result is consistent with preliminary observations by Kamb & Ganguli (2024) regarding cases in which their purely convolutional models fail to generate coherent images, while models including attention succeed. Related theoretical work further probes why well-trained diffusion models can generalize despite apparent memorization pressures (Bonnaire et al., 2025; Vastola, 2025; Chen, 2025). These results suggest that departures from exact empirical-score fitting—mediated by inductive biases (both architectural and training dynamics) can explain how diffusion models avoid pure memorization while maintaining visual plausibility (Ambrogioni, 2023).

803
804
805
806
807
808
809

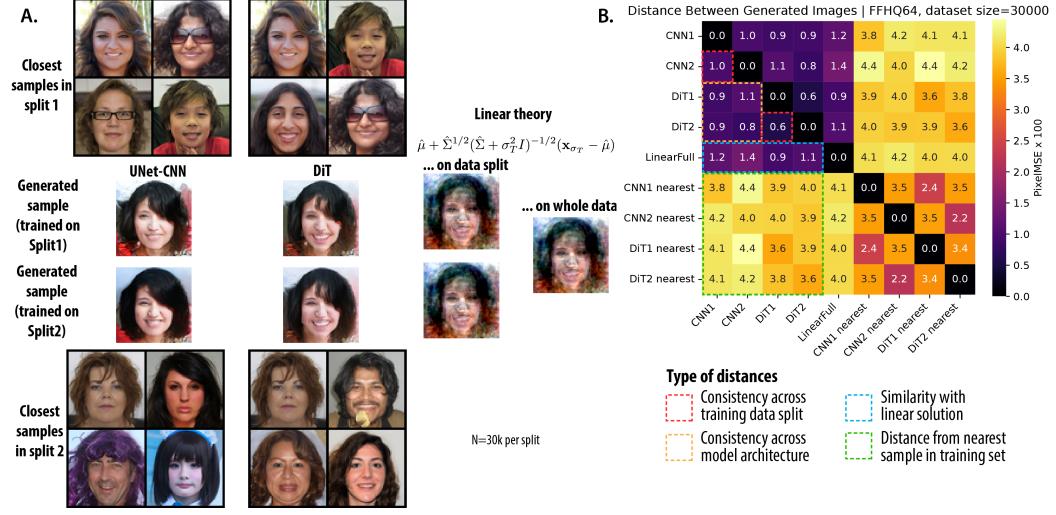
810 B EXTENDED RESULTS AND FIGURES
811812 B.1 EXTENDED VISUAL EXAMPLES FOR MOTIVATING OBSERVATION
813

Figure 6: **Motivating observation and the linear theory for FFHQ64 dataset.** Similar format to Fig. 1, but for FFHQ64 dataset. **A.** Examples of generated samples from the same noise seed, for UNet, DiT, and linear denoiser on split 1 and split 2 of data, each with 30k non overlapping samples. The closest 4 samples in its training set are shown above and below the generated sample. One can appreciate the visual similarity of samples generated from models trained on separate splits and even with different neural architectures, and also with the linear denoiser on each split. Admittedly, the generated outcomes of linear denoisers at 64 resolution look not as good, esp. for edges, showing signatures of non-Gaussian statistics, as Wang & Vastola (2024b) has pointed out. **B.** Quantification of **A**, paired image distances (MSE) averaging from 512 initial noises.

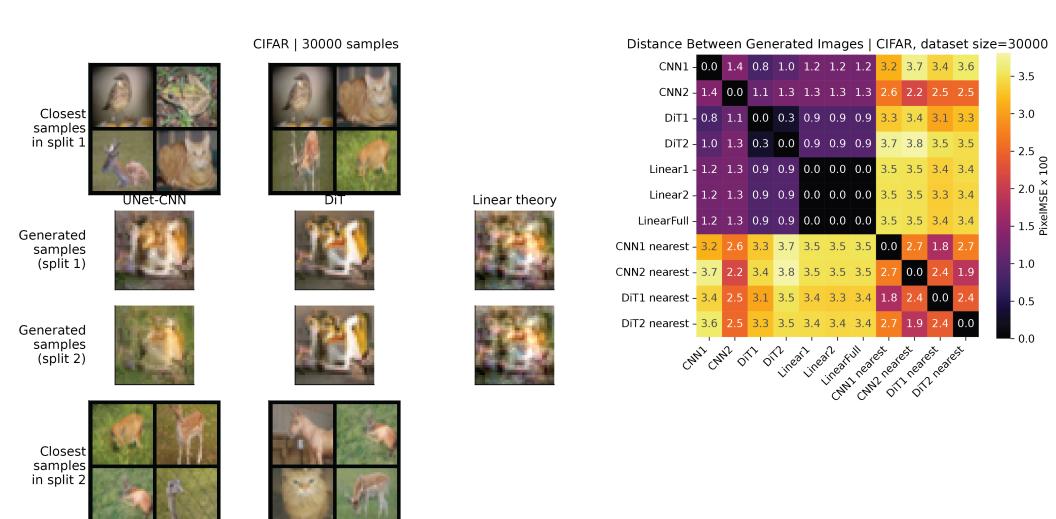
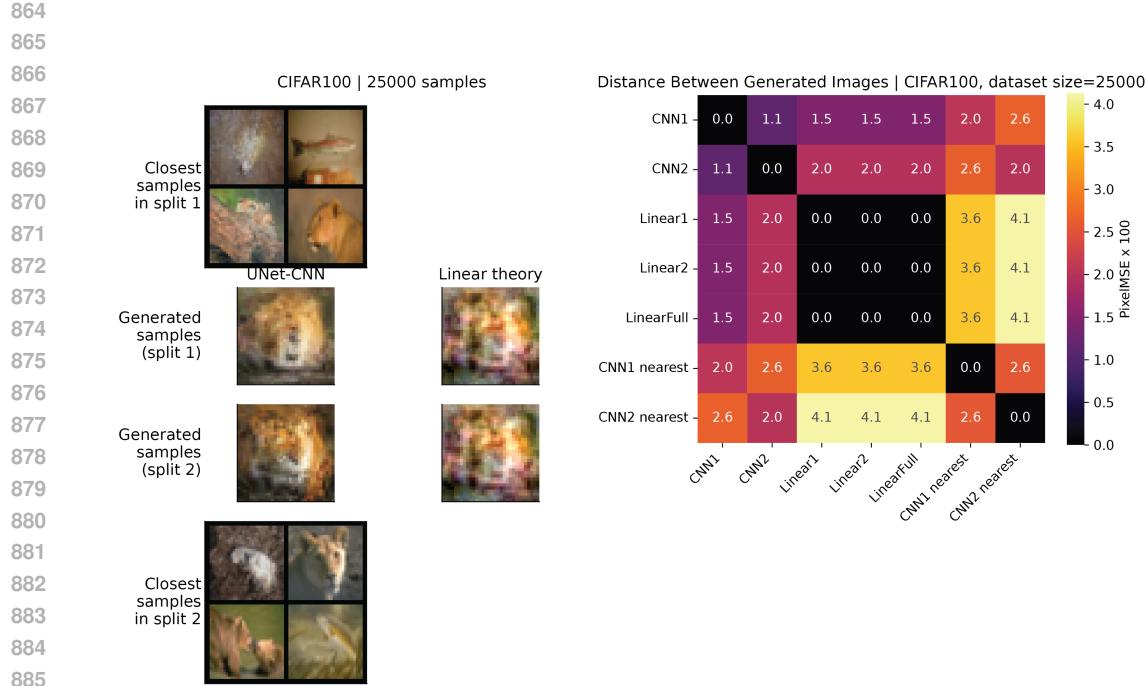
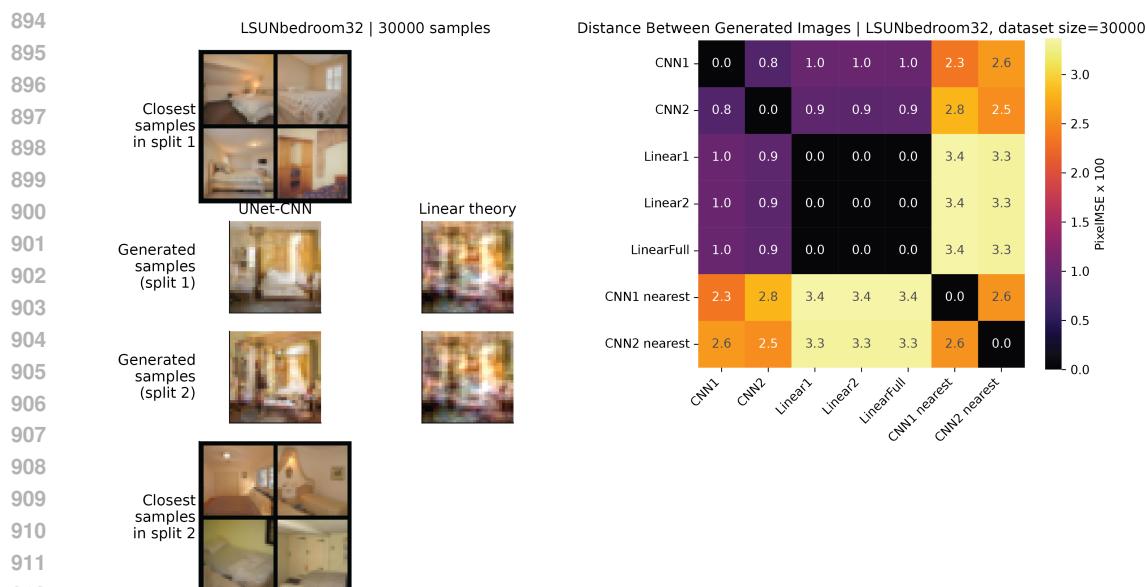


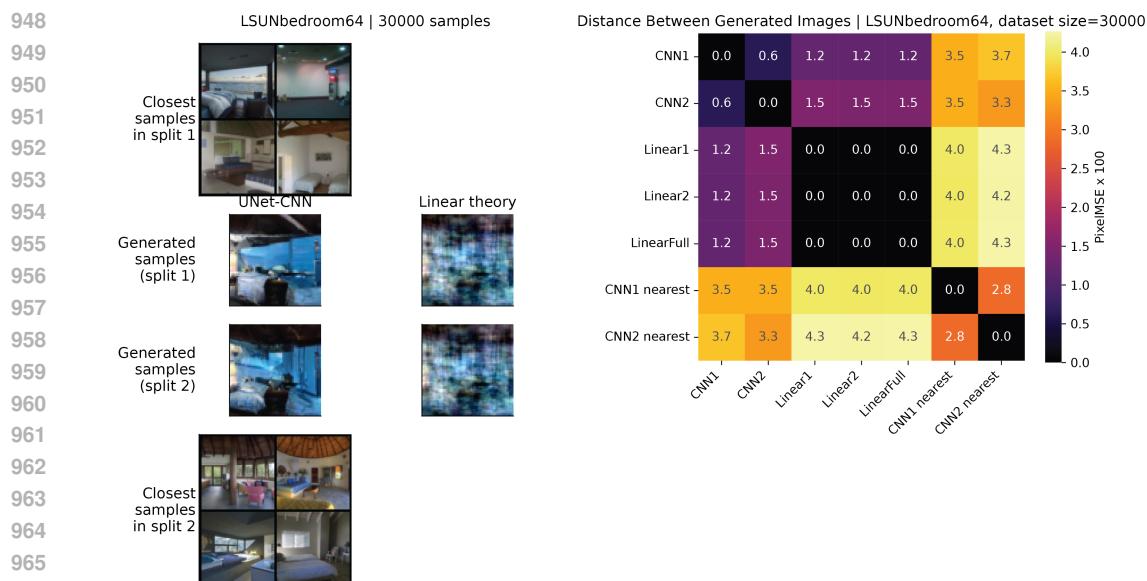
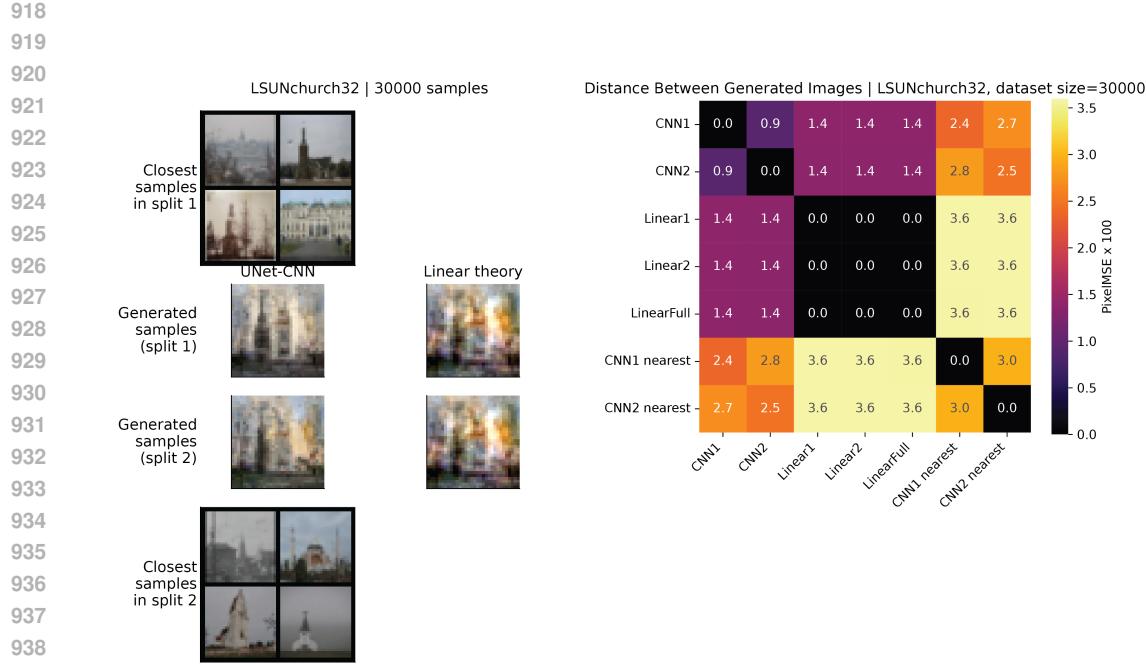
Figure 7: **Motivating observation and the linear theory for CIFAR10 dataset.** Similar format to Fig. 1. **Left.** Generated samples from DNN and linear theory from initial noise seed 2. **Right.** Paired image distance MSE averaging from 1000 initial noises.



887 **Figure 8: Motivating observation and the linear theory for CIFAR100 dataset.** Similar format to
888 Fig. 1. **Left.** Generated samples from DNN and linear theory from initial noise seed 2. **Right.** Paired
889 image distance MSE averaging from 1000 initial noises.



914 **Figure 9: Motivating observation and the linear theory for LSUN bedroom dataset (32 pixel).**
915 Similar format to Fig. 1. **Left.** Generated samples from DNN and linear theory from initial noise
916 seed 2. **Right.** Paired image distance MSE averaging from 1000 initial noises.



972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987

LSUNchurch64 | 30000 samples

Closest samples in split 1

UNet-CNN

Generated samples (split 1)

Generated samples (split 2)

Linear theory

Closest samples in split 2

Distance Between Generated Images | LSUNchurch64, dataset size=30000

	CNN1	CNN2	Linear1	Linear2	LinearFull	CNN1 nearest	CNN2 nearest
CNN1	0.0	1.9	2.3	2.3	2.3	5.2	5.9
CNN2	1.9	0.0	1.8	1.8	1.8	6.1	5.4
Linear1	2.3	1.8	0.0	0.2	0.1	5.5	5.2
Linear2	2.3	1.8	0.2	0.0	0.1	5.5	5.2
LinearFull	2.3	1.8	0.1	0.1	0.0	5.5	5.2
CNN1 nearest	5.2	6.1	5.5	5.5	5.5	0.0	4.3
CNN2 nearest	5.9	5.4	5.2	5.2	5.2	4.3	0.0

PixelMSE x 100

Figure 12: **Motivating observation and the linear theory for LSUN church dataset (64 pixel).** Similar format to Fig. 1. **Left.** Generated samples from DNN and linear theory from initial noise seed 2. **Right.** Paired image distance MSE averaging from 1000 initial noises.

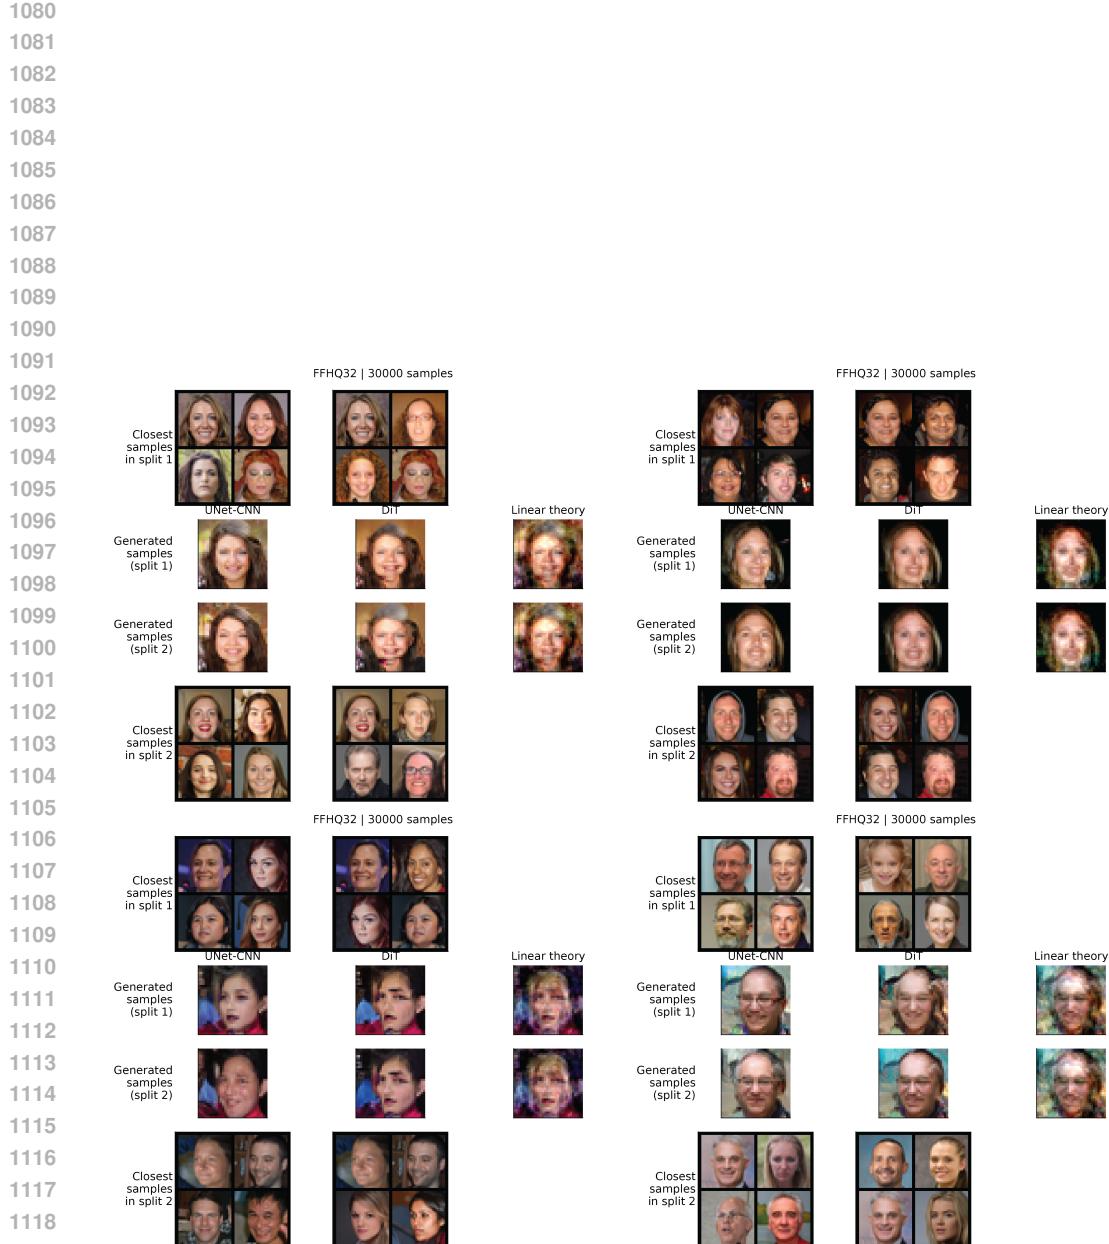
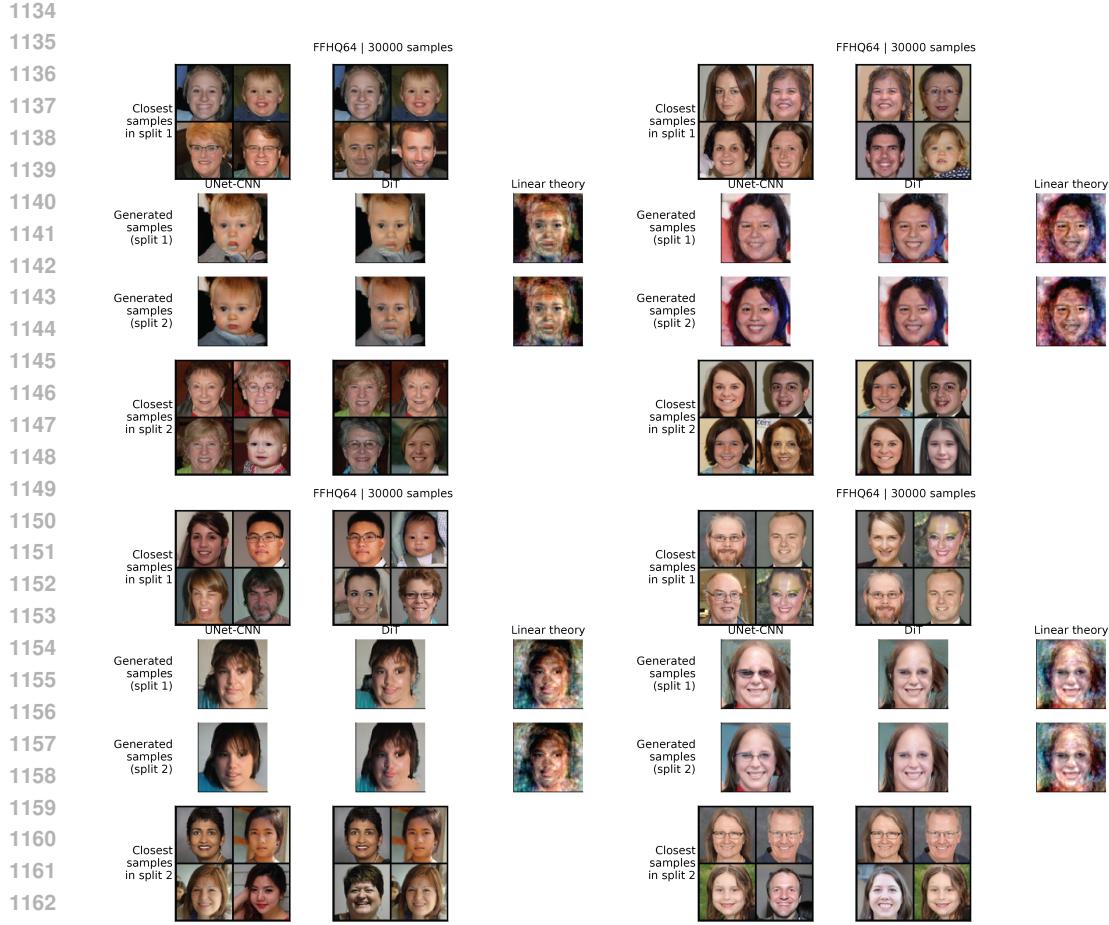


Figure 15: **Extended visual comparison of generation consistency and the linear theory for FFHQ32 dataset.** Similar format to Fig. 1A. Generated samples from DNN and linear theory from initial noise seed 0,1,3,4

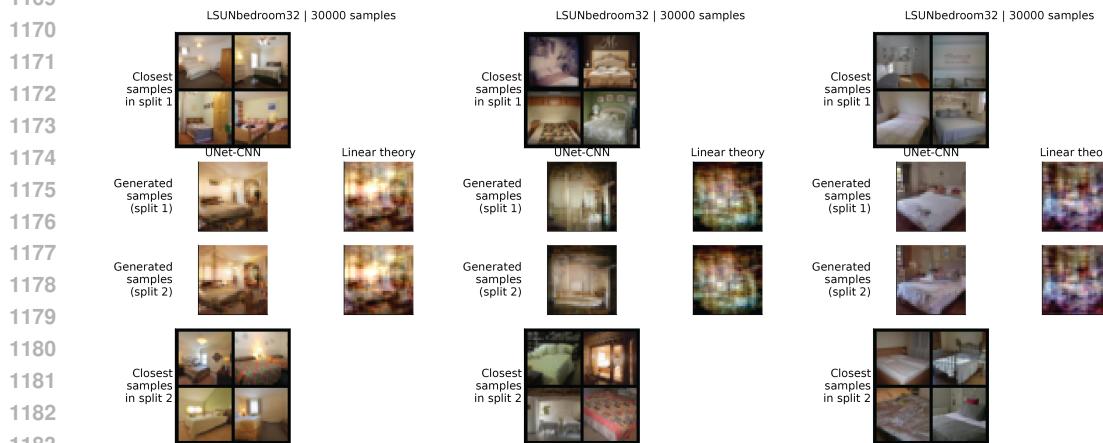


1164

1165

1166

Figure 16: **Extended visual comparison of generation consistency and the linear theory for**
FFHQ64 dataset. Similar format to Fig. 1A. Generated samples from DNN and linear theory from
initial noise seed 0,1,3,4



1185

1186

1187

Figure 17: **Extended visual comparison of generation consistency and the linear theory for**
LSUN bedroom dataset (32 pixel). Similar format to Fig. 1A. Generated samples from DNN and
linear theory from initial noise seed 0,1,3

1188

1189

1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

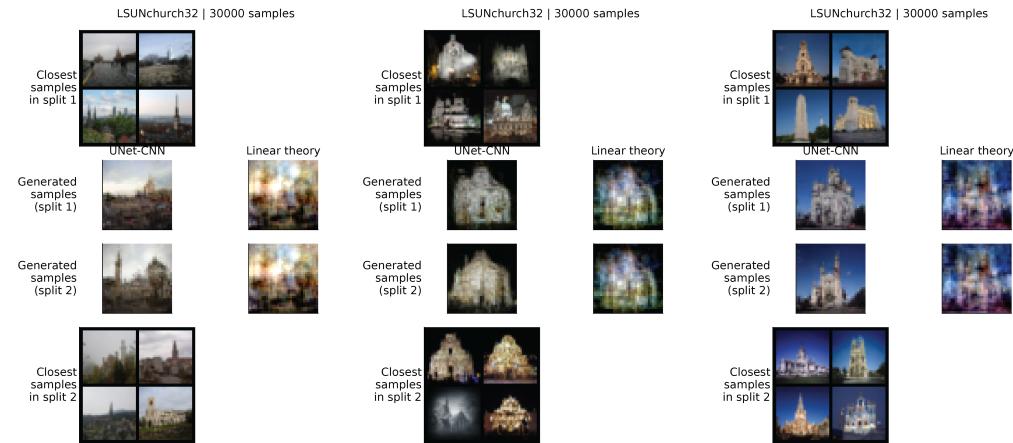


Figure 18: **Extended visual comparison of generation consistency and the linear theory for LSUN church dataset (32 pixel).** Similar format to Fig. 1A. Generated samples from DNN and linear theory from initial noise seed 0,1,3

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

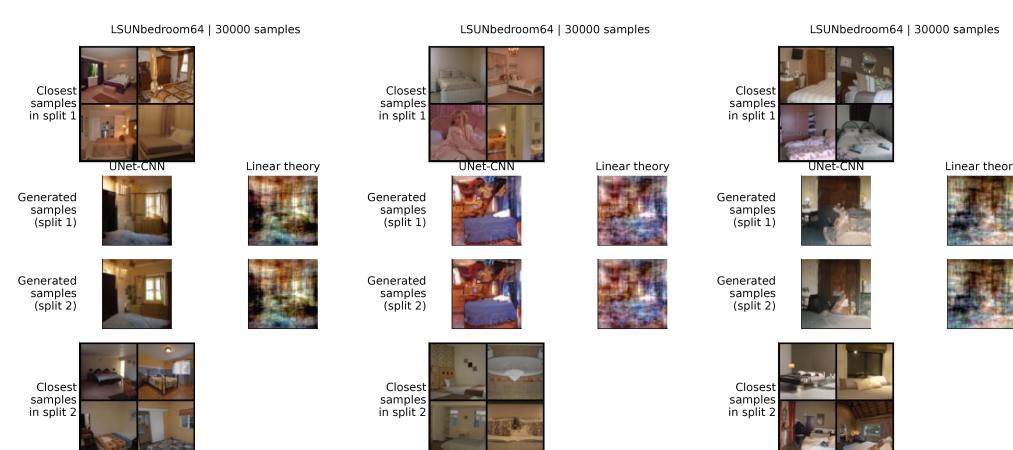


Figure 19: **Extended visual comparison of generation consistency and the linear theory for LSUN bedroom dataset (64 pixel).** Similar format to Fig. 1A. Generated samples from DNN and linear theory from initial noise seed 0,1,3

1238

1239

1240

1241

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261 LSUNchurch64 | 30000 samples



1262 LSUNchurch64 | 30000 samples



1263 LSUNchurch64 | 30000 samples



1264 Linear theory



Figure 20: **Extended visual comparison of generation consistency and the linear theory for LSUN church dataset (64 pixel).** Similar format to Fig. 1A. Generated samples from DNN and linear theory from initial noise seed 0,1,3

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295

1296
1297

B.2 EXTENDED VALIDATION OF THE THEORY

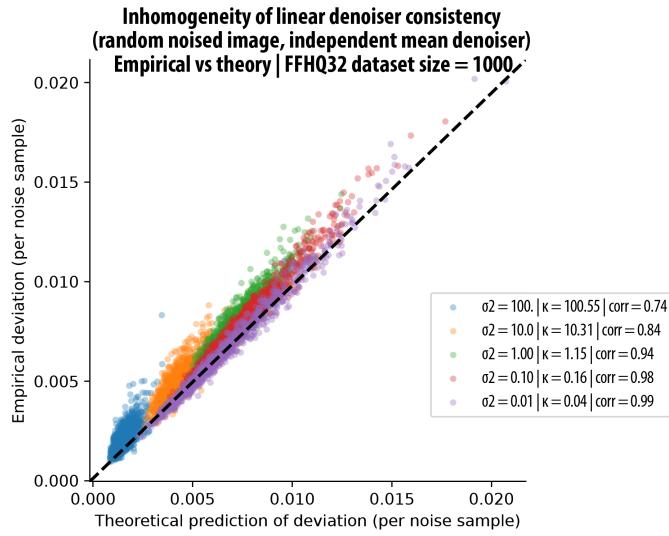
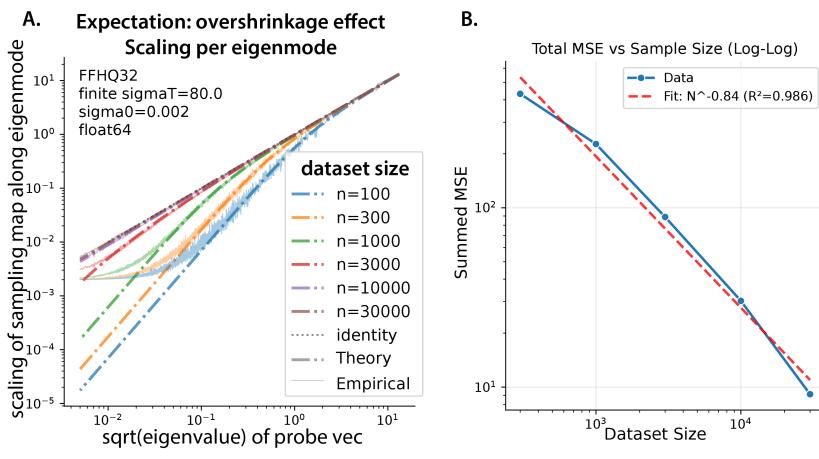
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
13151316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

Figure 21: **Point by point prediction of denoiser consistency. (FFHQ32 dataset, $n = 1000$)**
 Each dot denotes one noised image sample, x-axis shows the theoretical prediction from Eq. 7, after marginalizing over \mathbf{v} ; y-axis shows the empirical measurement of their MSE after training two linear denoiser on non-overlapping data splits. We note that, the RMT theory prediction is more precise for lower noise scales; at higher noise scales, we think the effect of different empirical means $\hat{\mu}$ kicks in, resulting in deviation from the theory that only considers $\hat{\Sigma}$.

1350
 1351
 1352
 1353
 1354
 1355
 1356
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380

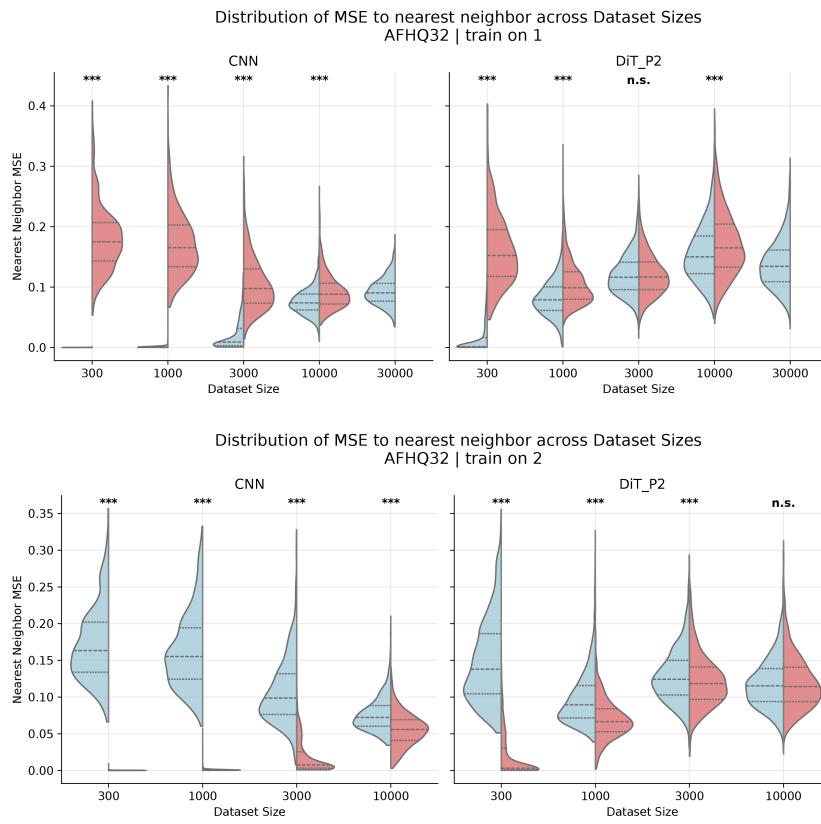


1381 **Figure 22: Finite sample effect on diffusion sampling map. (extended)** **A.** *Overshinkage of*
 1382 *expectation.* The expected scaling along PC $\mathbf{u}_k^\top \hat{\Sigma}^{1/2} \mathbf{u}_k$ of empirical sampling map compared to the
 1383 *ideal scaling* $\sqrt{\lambda_k}$, here we used $\sigma_0 = 0.002$ for empirical matrix computation. The σ_0 is smallest
 1384 noise scale that probability flow ODE integration stops, for numerical reasons. This floors the smallest
 1385 scaling factor it could generate, making the mismatch with theory at the low eigen space. **B.** Overall
 1386 MSE scaling with respect to dataset size, roughly scales at $1/n$ at large data, but the scaling is
 1387 shallower at smaller data scale.

1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403

1404
1405

B.3 EXTENDED EVIDENCE FROM THE DNN VALIDATION EXPERIMENTS

1406
14071408
14091410
14111412
14131414
14151416
14171418
14191420
14211422
14231424
14251426
14271428
14291430
14311432
1433Figure 23: **DNN validation experiments (AFHQ32), nearest neighbor in training and control set**1434
14351436
14371438
14391440
14411442
14431444
14451446
14471448
14491450
14511452
14531454
14551456
1457

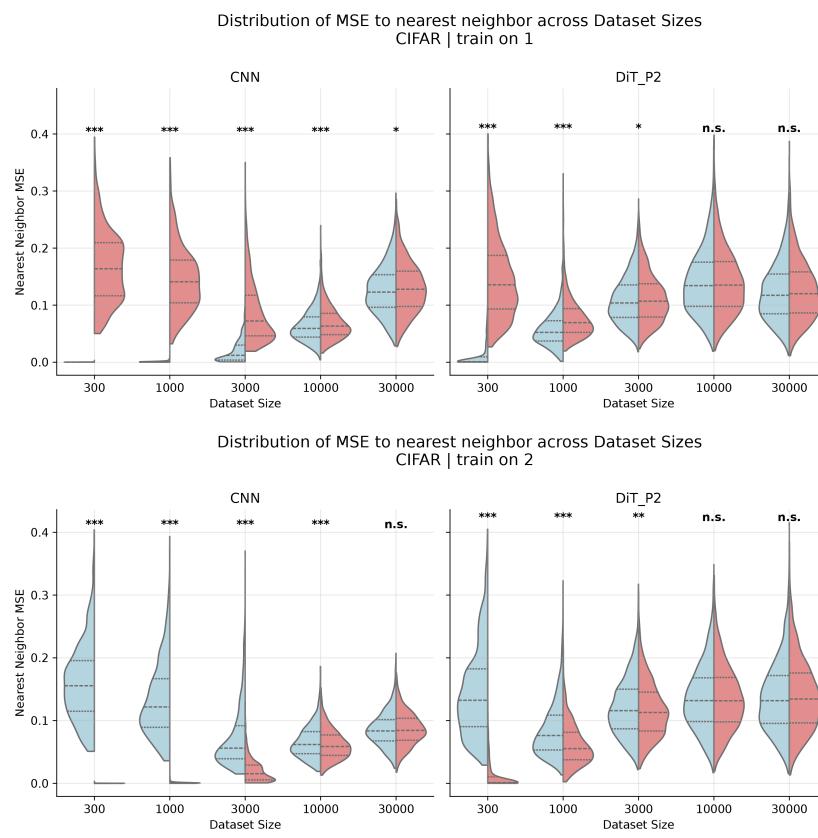


Figure 24: **DNN validation experiments (CIFAR), nearest neighbor in training and control set**

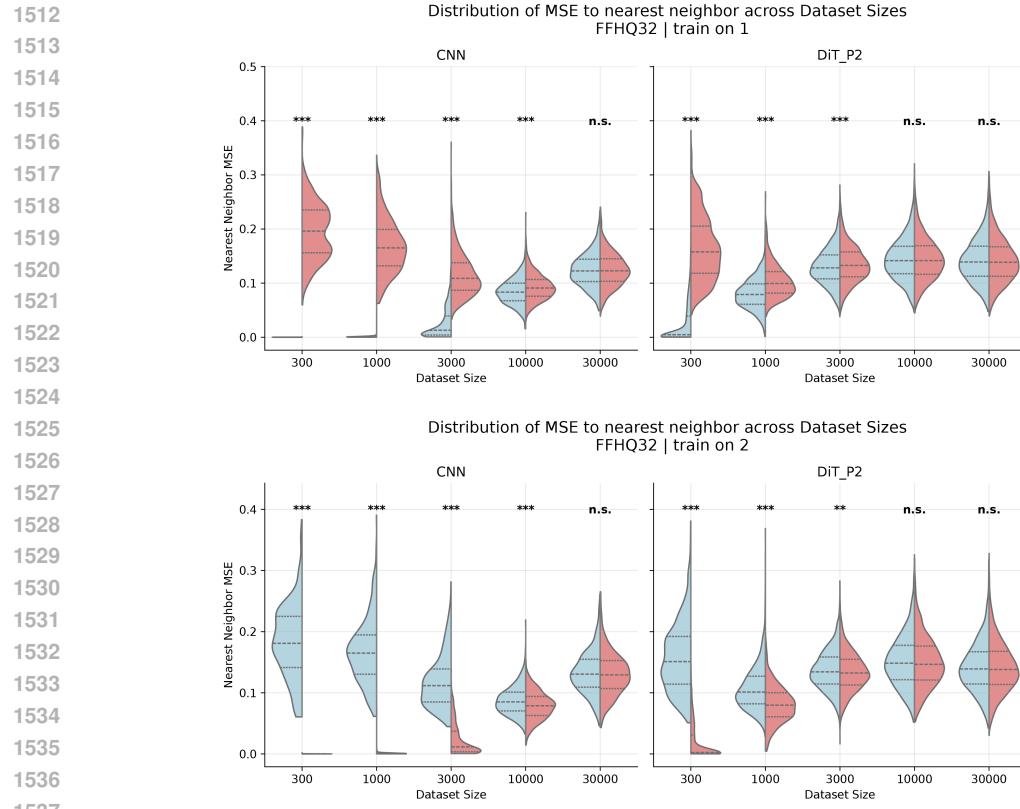


Figure 25: DNN validation experiments (FFHQ32), nearest neighbor in training and control set

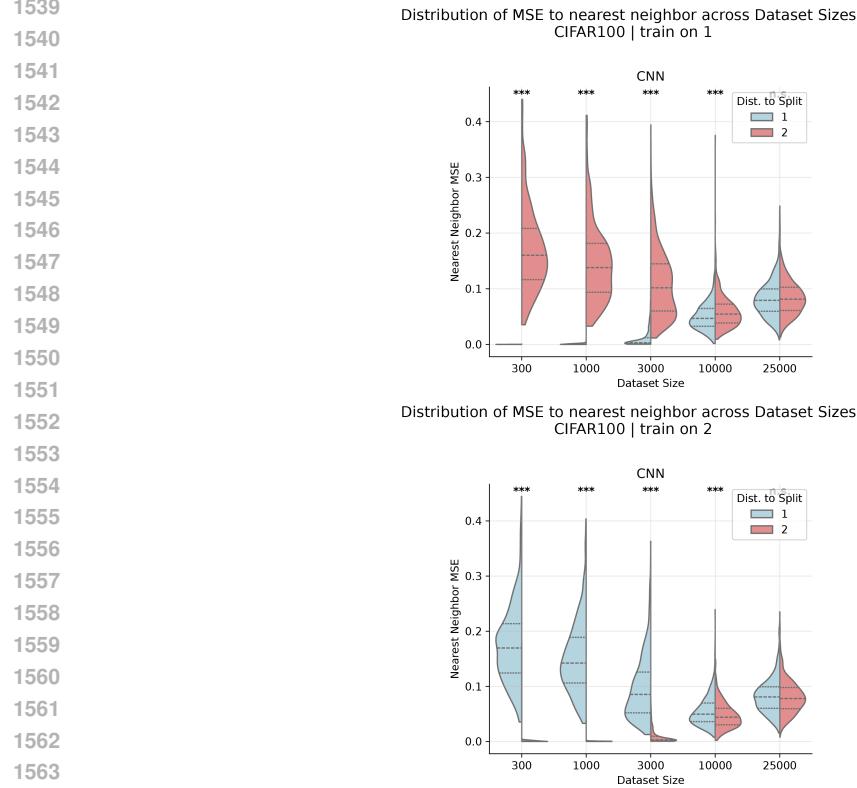
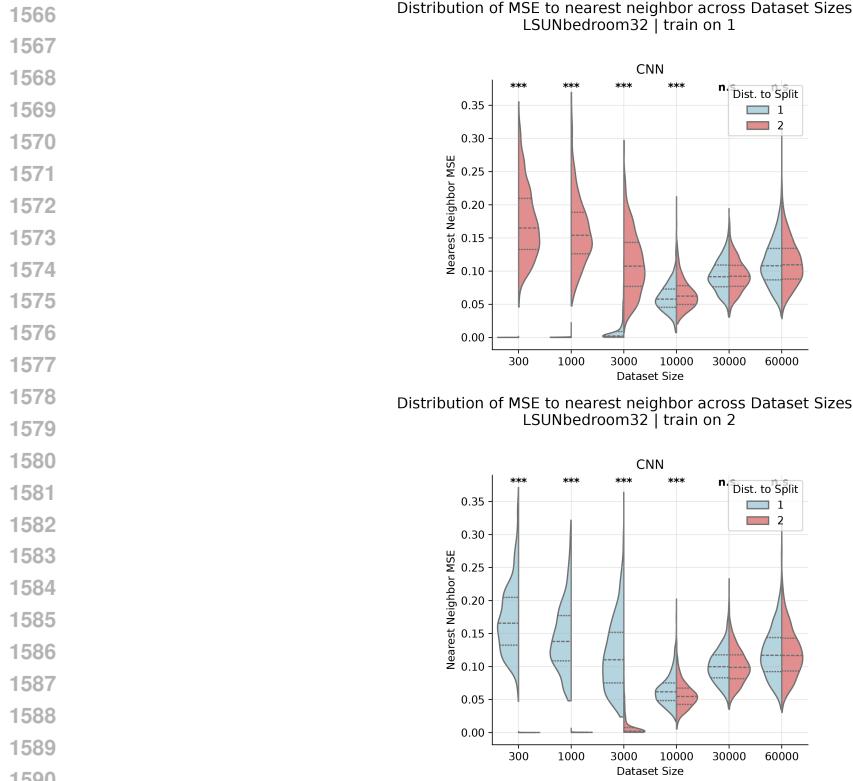
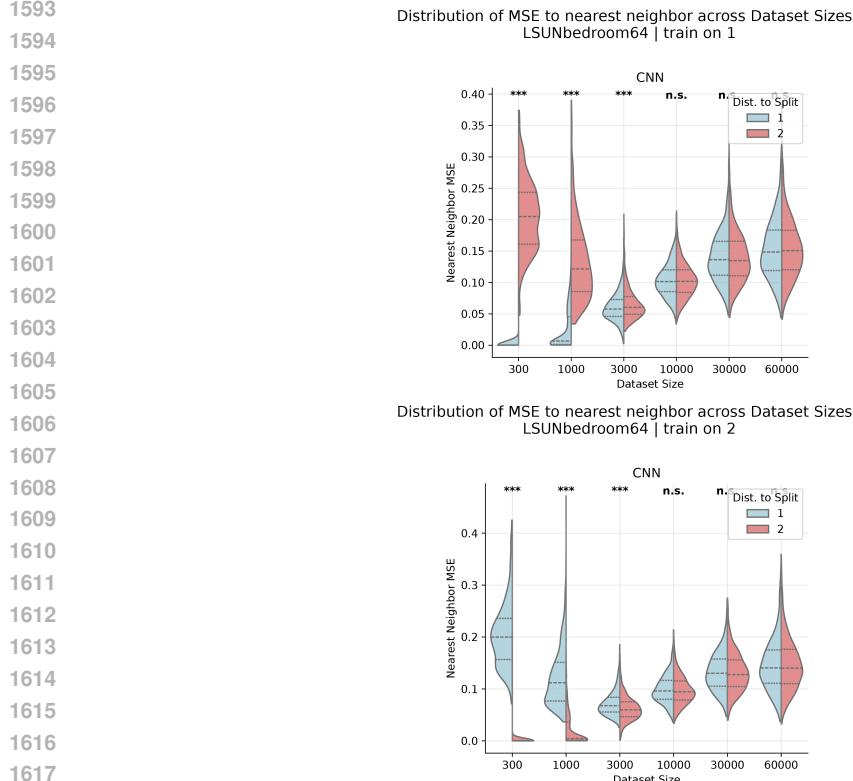


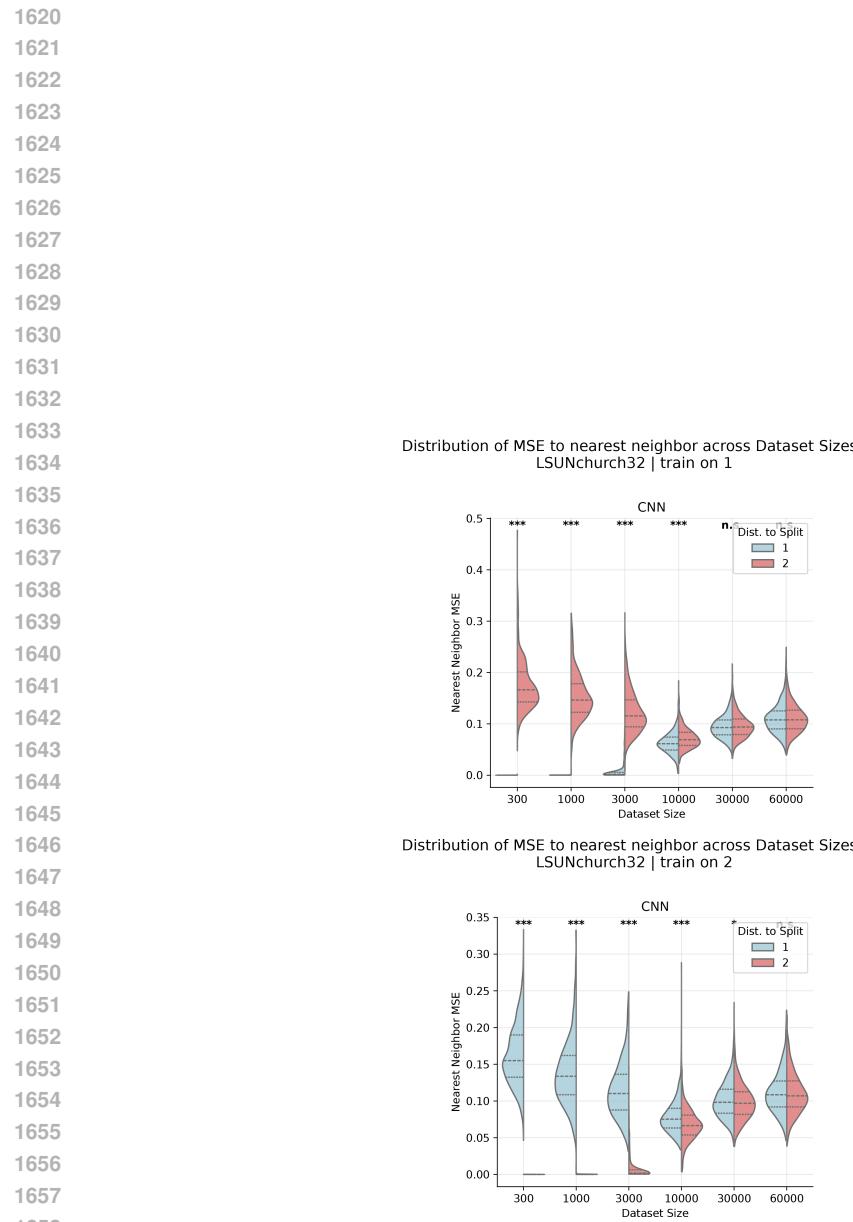
Figure 26: DNN validation experiments (CIFAR100), nearest neighbor in training and control set



1591 **Figure 27: DNN validation experiments (LSUN bedroom 32), nearest neighbor in training and
1592 control set**

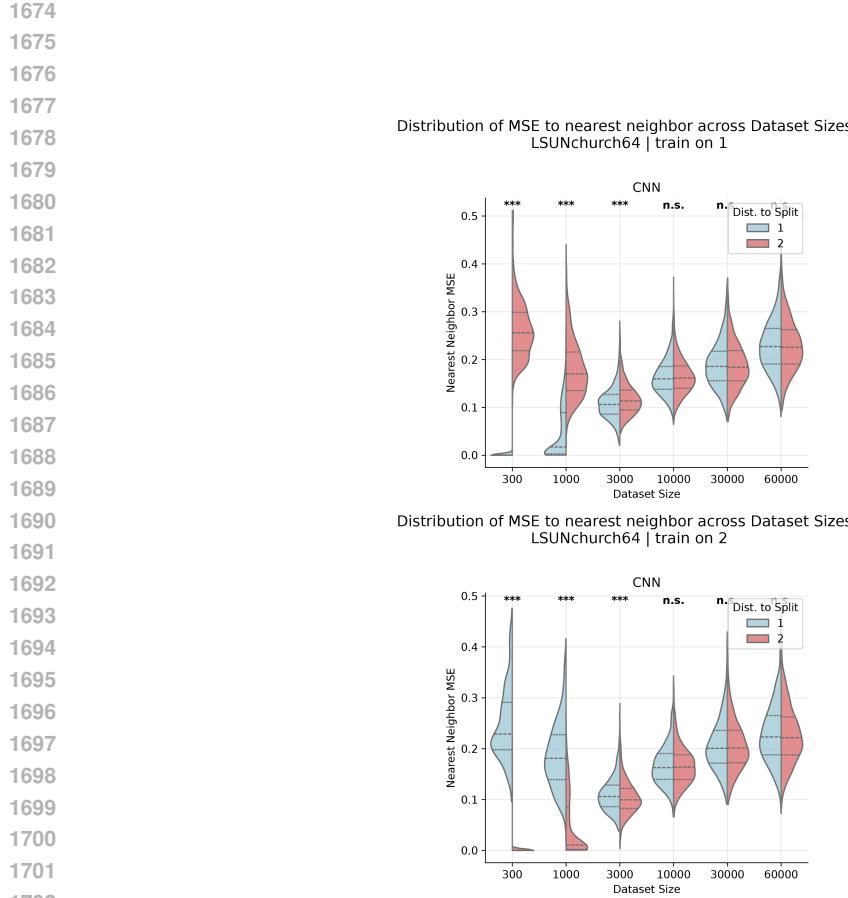


1618 **Figure 28: DNN validation experiments (LSUN bedroom 64), nearest neighbor in training and
1619 control set**



1659 **Figure 29: DNN validation experiments (LSUN church 32), nearest neighbor in training and
 1660 control set**

1661
 1662
 1663
 1664
 1665
 1666
 1667
 1668
 1669
 1670
 1671
 1672
 1673



1703 Figure 30: DNN validation experiments (LSUN church 64), nearest neighbor in training and
1704 control set

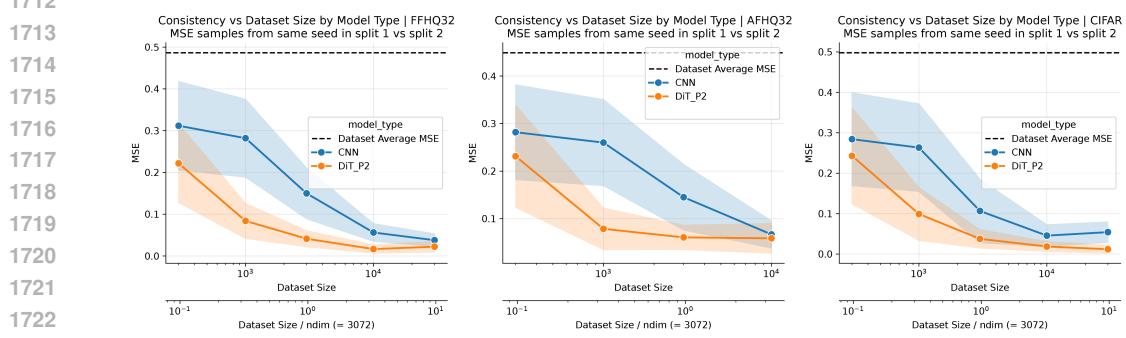


Figure 31: DNN validation experiments, scaling of consistency with dataset size

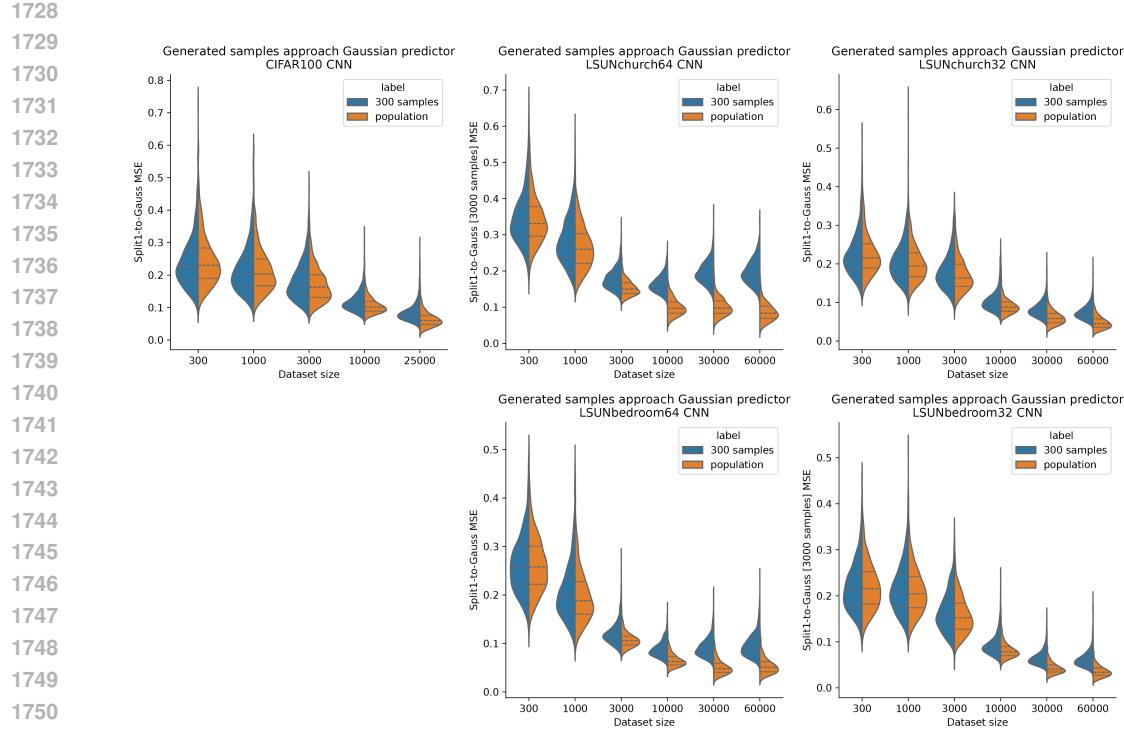


Figure 32: **DNN generated samples approach the linear theory predictor (with finite sample or population covariance).** With increasing dataset size n , the generated sample from DNN (trained on split 1) with a fixed noise seed gradually approach the linear theory predictor using the same initial seed. Consistent across datasets.

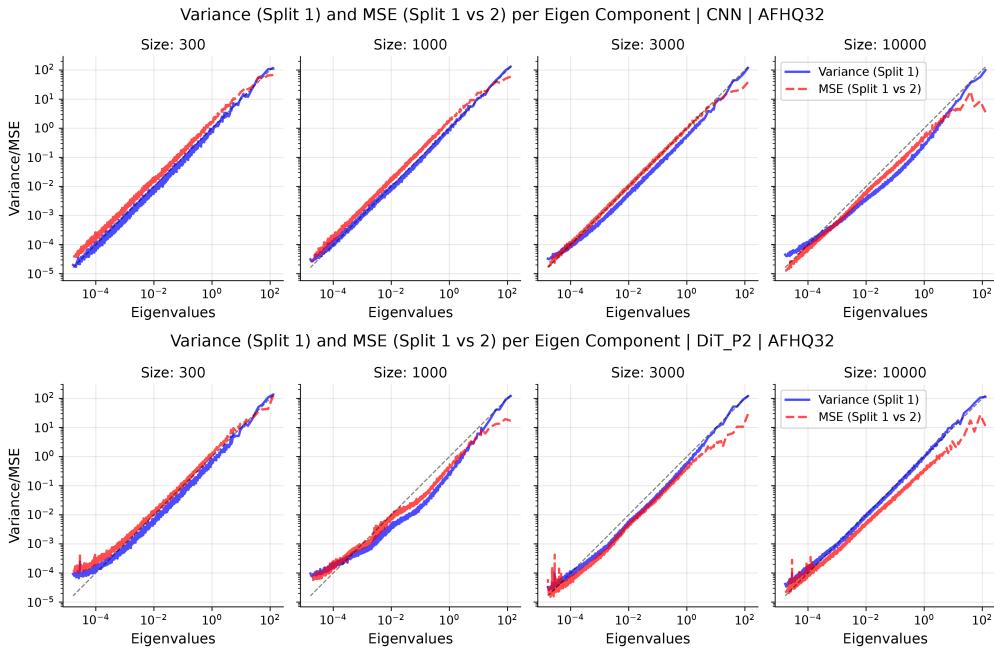


Figure 33: **DNN validation experiments, Anisotropy and overshrinking (AFHQ32)**

1782

1783

1784

1785

1786

1787

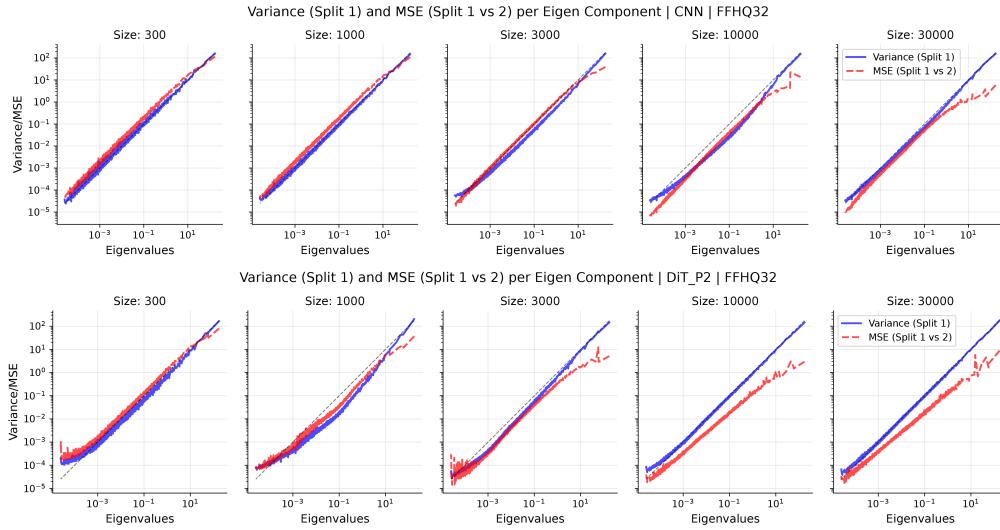


Figure 34: DNN validation experiments, Anisotropy and overshrinking (FFHQ32)

1803

1804

1805

1806

1807

1808

1809

1810

1811

1812

1813

1814

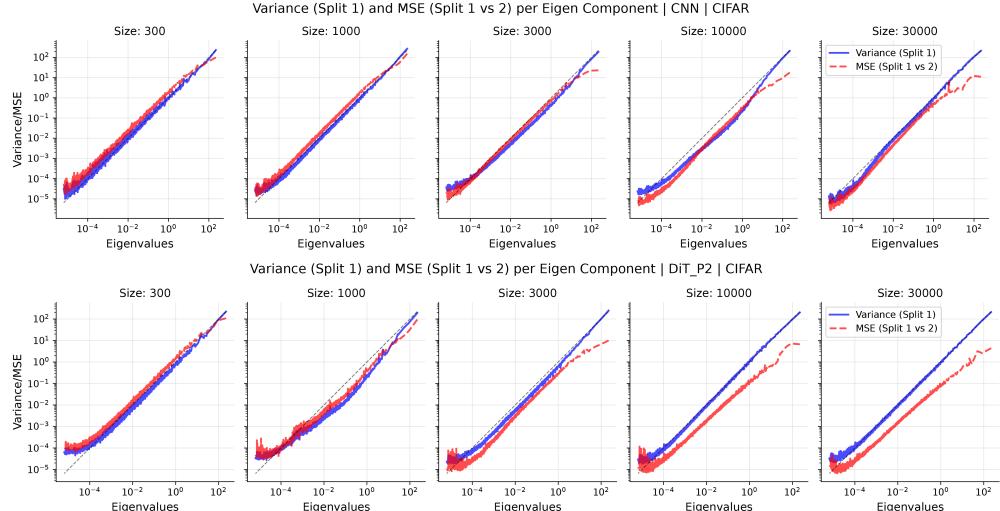


Figure 35: DNN validation experiments, Anisotropy and overshrinking (CIFAR10)

1831

1832

1833

1834

1835

1836

1837

1838

1839

1840

1841

1842

1843

1844

1845

1846

1847

1848

1849

1850

1851

1852

1853

1854

1855

1856

1857

1858

1859

1860

1861

1862

1863

1864

1865

1866

1867

1868

1869

1870

1871

1872

1873

1874

1875

1876

1877

1878

1879

1880

1881

1882

1883

1884

1885

1886

1887

1888

1889

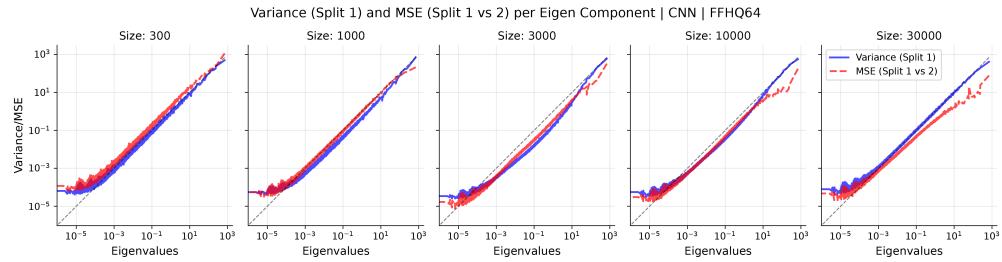


Figure 36: DNN validation experiments, Anisotropy and overshrinking (FFHQ64)

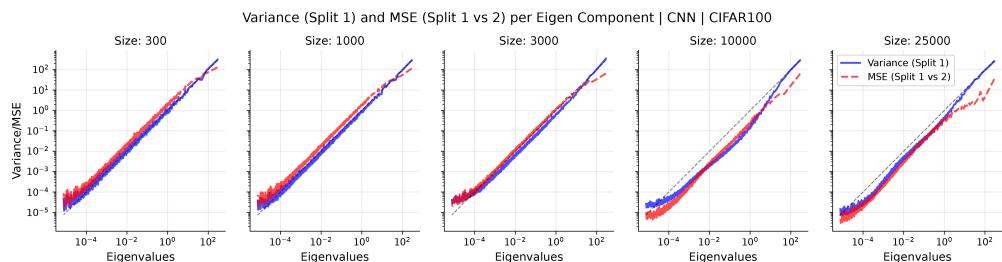


Figure 37: DNN validation experiments, Anisotropy and overshrinking (CIFAR100)

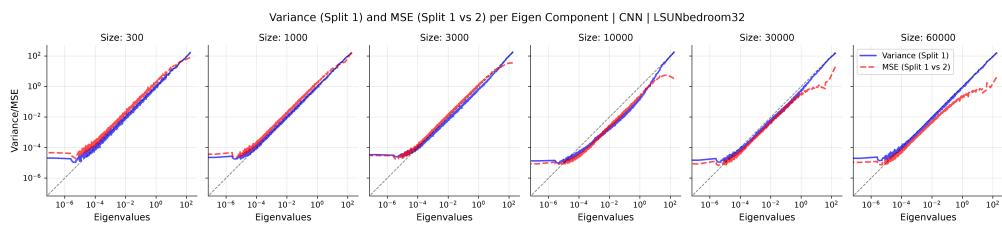


Figure 38: DNN validation experiments, Anisotropy and overshrinking (LSUN bedroom 32)

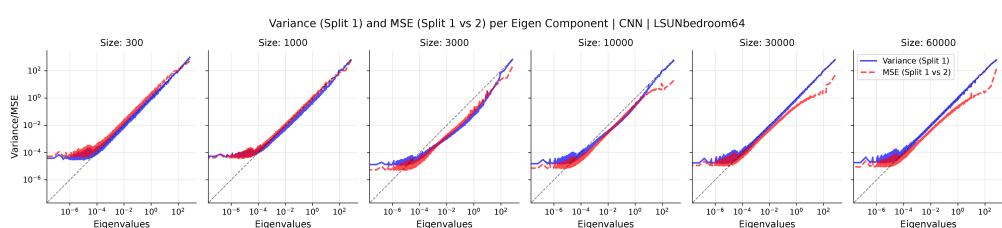


Figure 39: DNN validation experiments, Anisotropy and overshrinking (LSUN bedroom 64)

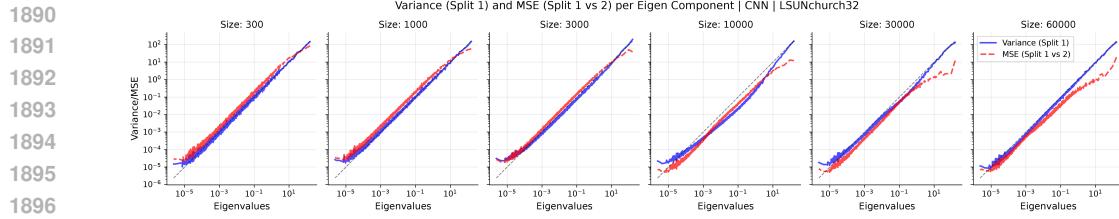


Figure 40: DNN validation experiments, Anisotropy and overshrinking (LSUN church 32)

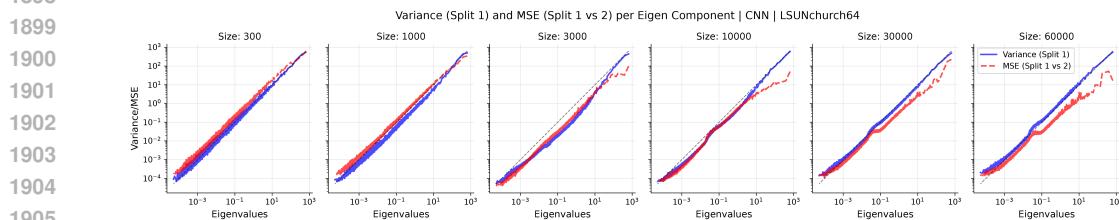


Figure 41: DNN validation experiments, Anisotropy and overshrinking (LSUN church 64)

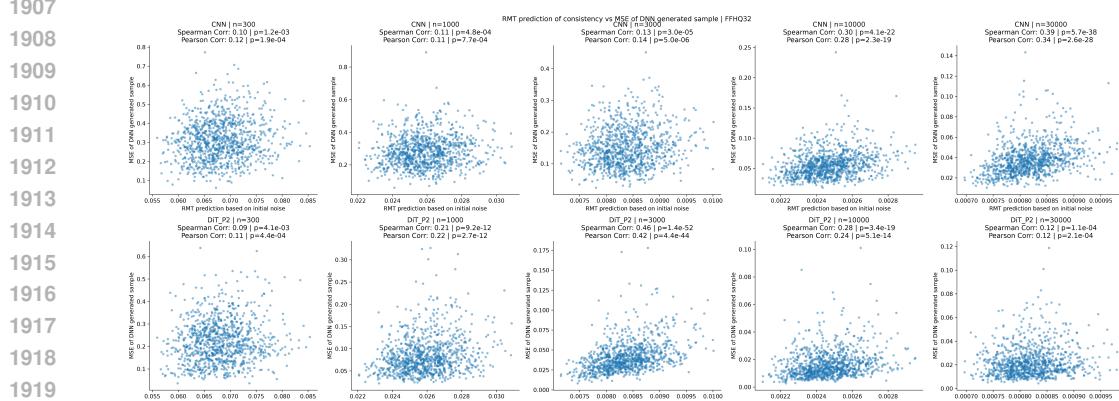


Figure 42: DNN validation experiments, RMT predicting inhomogeneity (FFHQ32)

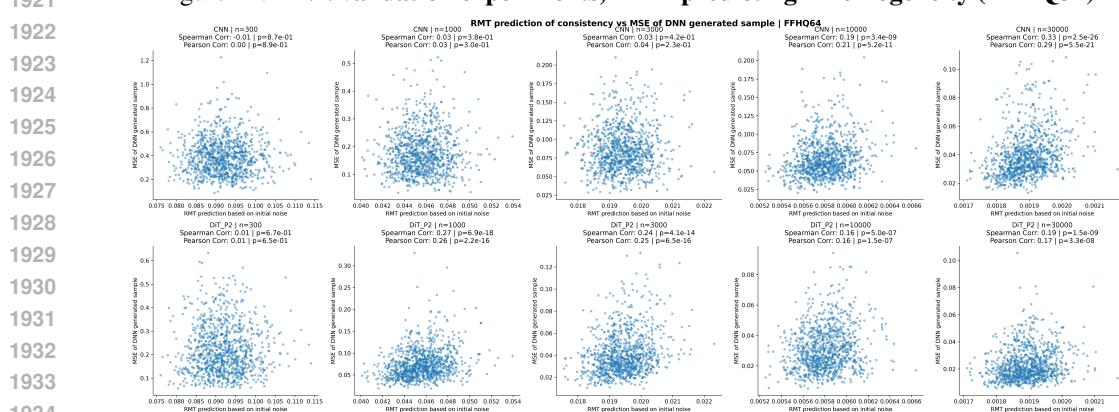


Figure 43: DNN validation experiments, RMT predicting inhomogeneity (FFHQ64)

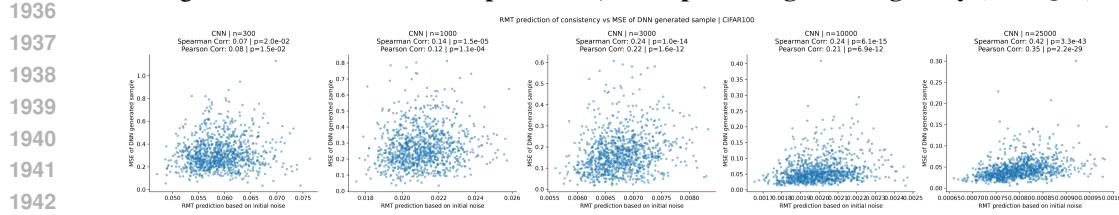


Figure 44: DNN validation experiments, RMT predicting inhomogeneity (CIFAR100)

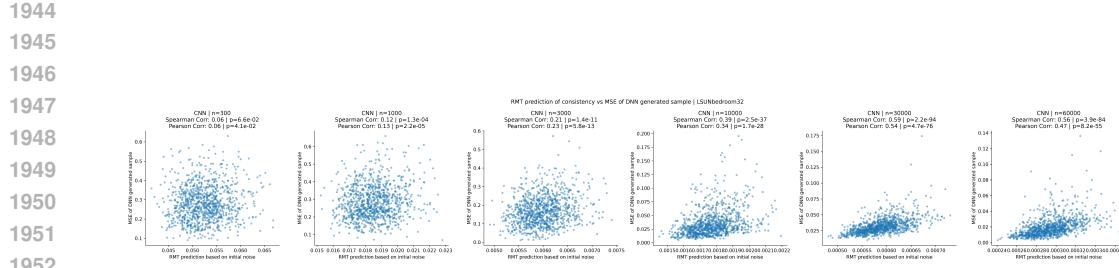


Figure 45: DNN validation experiments, RMT predicting inhomogeneity (LSUN bedroom 32)

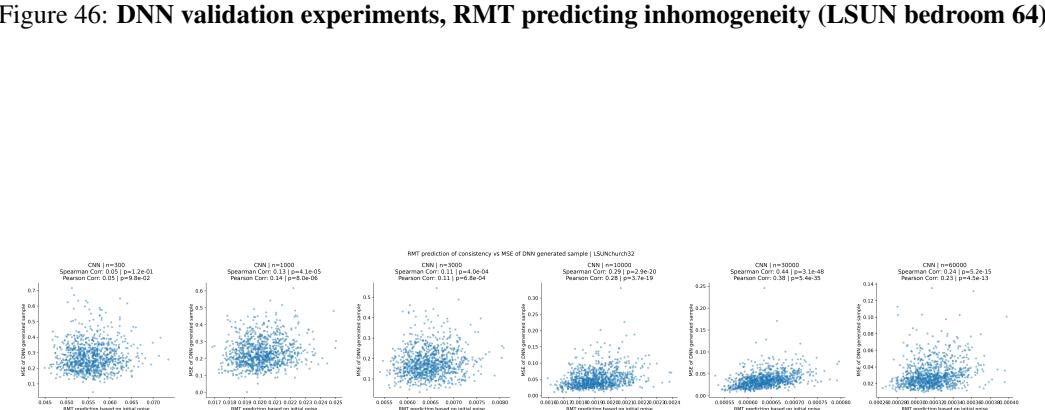


Figure 46: DNN validation experiments, RMT predicting inhomogeneity (LSUN bedroom 64)

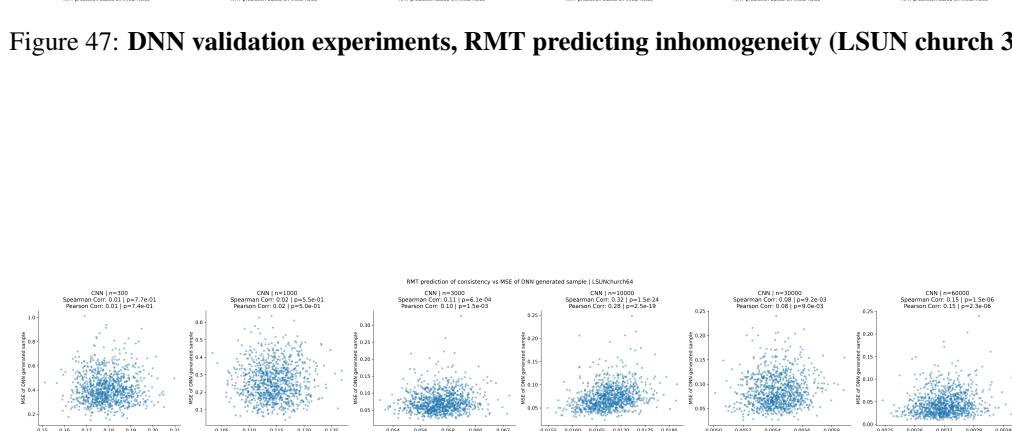


Figure 47: DNN validation experiments, RMT predicting inhomogeneity (LSUN church 32)



Figure 48: DNN validation experiments, RMT predicting inhomogeneity (LSUN church 64)

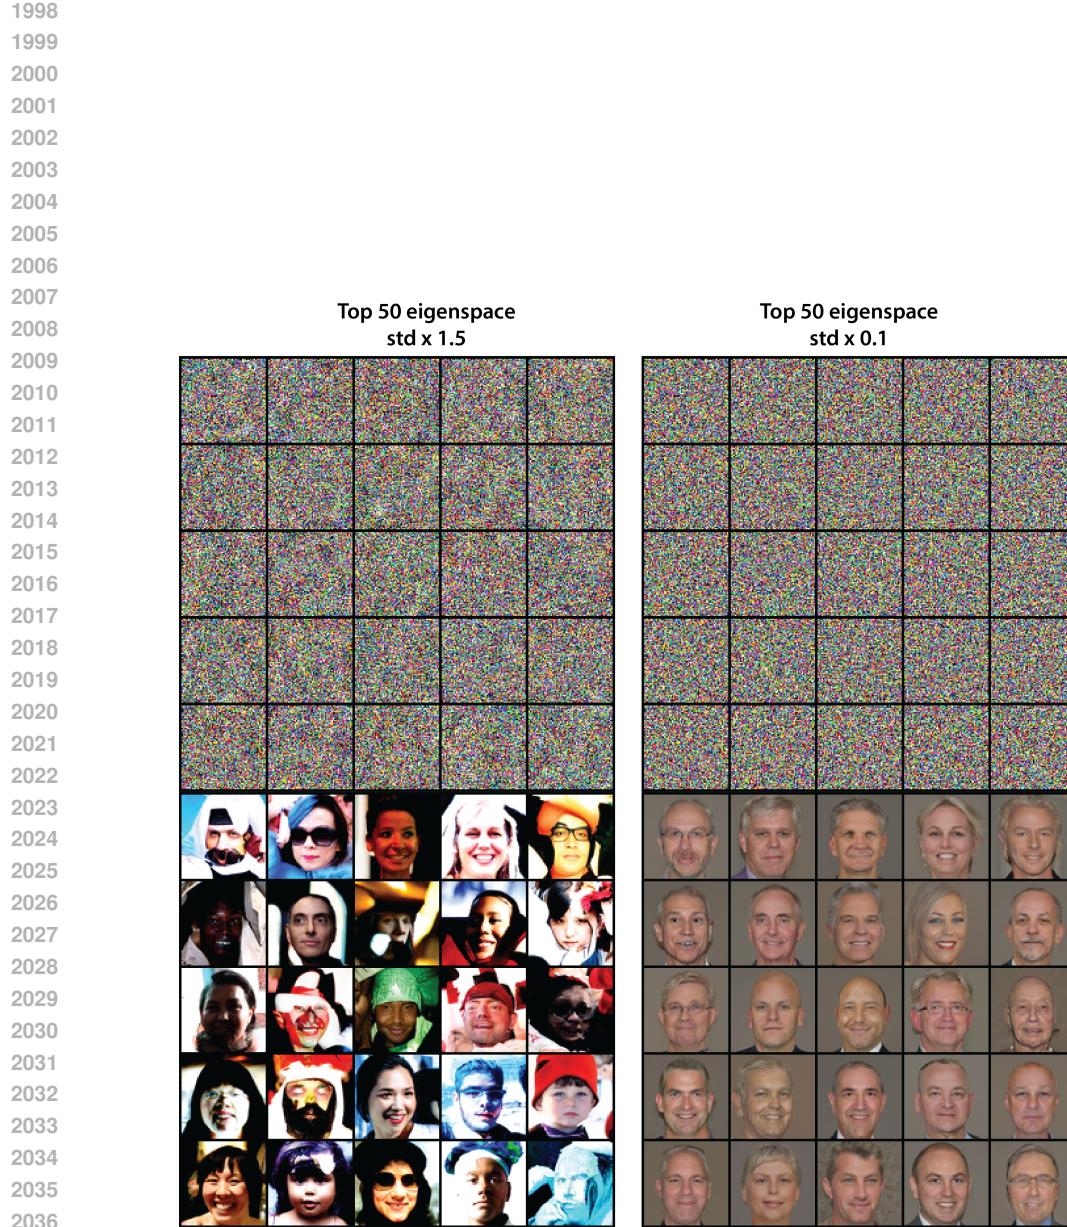


Figure 49: **Anisotropic structure of the initial noise space (CNN-UNet FFHQ64).** **Left:** Amplifying the initial-noise amplitude in the top-50 eigenspace. **Right:** Decreasing the initial-noise amplitude in the top-50 eigenspace. Top: initial noise; Bottom: generated samples (same noise seed). Increasing noise in the dominant eigendirections introduces more visual artifacts by amplifying the top eigen-structure of the generative map (Eq. 3). Conversely, reducing noise in these dimensions yields a cleanly segmented face against a gray background.

2043
2044
2045
2046
2047
2048
2049
2050
2051

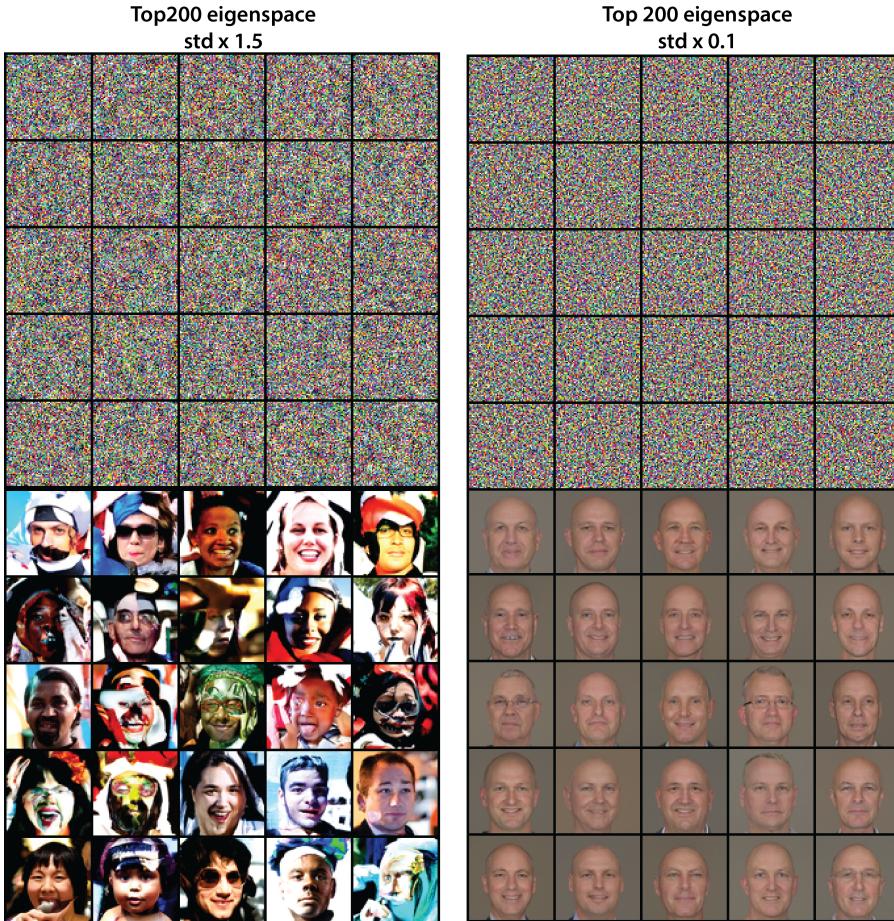


Figure 50: **Anisotropic structure of the initial noise space (CNN-UNet FFHQ64).** **Left:** Amplifying the initial-noise amplitude in the top-200 eigenspace. **Right:** Decreasing the initial-noise amplitude in the top-200 eigenspace. Top: initial noise; Bottom: generated samples (same noise seed). Similar to the top-50 case: stronger noise in leading eigendirections amplifies dominant structure, producing visible artifacts. Suppressing noise yields an even more homogeneous and simplified face, reflecting the removal of additional variation modes.

2160 **C PROOF AND DERIVATIONS**
 2161

2162 **C.1 DETERMINISTIC EQUIVALENCE RELATIONS**
 2163

2164 Here we collect the one-point and two point deterministic equivalence relationships adopted from
 2165 [Atanasov et al. \(2024b; 2025\); Bach \(2024\)](#), under the same notations.
 2166

2167 **Set up** Using similar notation as [Bach \(2024\)](#), we consider data matrix $X \in \mathbb{R}^{n \times d}$, where each row
 2168 is an i.i.d. sample \mathbf{x}_i . The population covariance of these samples is denoted as Σ . The key object of
 2169 analysis is their empirical covariance
 2170

$$\hat{\Sigma} = \frac{1}{n} X^\top X$$

2172
 2173 **Self-consistency equation for renormalized variable** The spectral properties of a matrix are
 2174 determined by the Stieltjes transform. We consider the Stieltjes transform of the kernel matrix
 2175 $\frac{1}{n} X X^\top$, defined as $\varphi(z) := \text{Tr}[(X X^\top - nzI)^{-1}]$. At the large matrix limit, the limiting variable
 2176 satisfy the following self consistent equation,
 2177

$$\frac{1}{\varphi(z)} + z = \gamma \int_0^\infty \frac{sd\mu(s)}{1 + s\varphi(z)} \quad (11)$$

2178 where $\mu(s)$ is the limiting spectral measure of the population covariance Σ . This follows from the
 2179 arguments in the Appendix of [Bach \(2024\)](#), as well as [Bai et al. \(2010\); Ledoit & Péché \(2011\)](#).
 2180

2181 This can be translated to the self-consistent equation of the renormalized ridge variable $\kappa(z) := \frac{1}{\varphi(-z)}$,
 2182 which is used throughout the paper,
 2183

$$\begin{aligned} \frac{1}{\varphi(-z)} - z &= \gamma \int_0^\infty \frac{sd\mu(s)}{1 + s\varphi(-z)} \\ \kappa(z) - z &= \gamma \int_0^\infty \frac{sd\mu(s)}{1 + s\frac{1}{\kappa(z)}} \\ \kappa(z) - z &= \gamma \kappa(z) \int_0^\infty \frac{sd\mu(s)}{\kappa(z) + s} \\ z &= \kappa(z) \left[1 - \gamma \int_0^\infty \frac{sd\mu(s)}{\kappa(z) + s} \right] \end{aligned}$$

2184 Practically, when solving such equations, given a finite size population covariance matrix, the integral
 2185 over the spectral measure can be represented as normalized trace, leading to the Silverstein equation
 2186 (Eq.5).
 2187

$$\kappa(\lambda) - \lambda = \gamma \kappa(\lambda) \text{tr}[\Sigma(\Sigma + \kappa(\lambda)I)^{-1}] \quad (12)$$

2188 **Degrees of Freedom** We define the degree of freedom functions with unnormalized trace, similar
 2189 to convention in [Bach \(2024\)](#), unlike [Atanasov et al. \(2025\)](#).
 2190

$$\text{df}_1(\lambda) := \text{Tr}[\Sigma(\Sigma + \lambda I)^{-1}] \quad (13)$$

$$\text{df}_2(\lambda) := \text{Tr}[\Sigma^2(\Sigma + \lambda I)^{-2}]. \quad (14)$$

2191 We see that
 2192

$$\begin{aligned} \text{df}_2(\kappa) - \text{df}_1(\kappa) &= \text{Tr}[\Sigma^2(\Sigma + \kappa I)^{-2}] - \text{Tr}[\Sigma(\Sigma + \kappa I)^{-1}] \\ &= \text{Tr}[(\Sigma(\Sigma + \kappa I)^{-1} - I)\Sigma(\Sigma + \kappa I)^{-1}] \\ &= \kappa \text{Tr}[\Sigma(\Sigma + \kappa I)^{-2}] \\ &> 0. \end{aligned}$$

2193 Note that both $\text{df}_2(\kappa), \text{df}_1(\kappa)$ are smaller than the number on non-zero eigenvalues of Σ , i.e. $\text{rank}(\Sigma)$.
 2194 Thus, we have the chain of inequalities
 2195

$$\min(n, p) \geq \text{rank}(\Sigma) > \text{df}_2(\kappa) > \text{df}_1(\kappa)$$

2214 **Basic equivalences** Following Proposition 1 of [Bach \(2024\)](#), we use the shorthand $\kappa(z) :=$
 2215 $1/\varphi(-z)$ to express the deterministic equivalences in the more convenient forms below. In what
 2216 follows, A and B are test matrices of bounded spectral norm. For the resolvent, we have:
 2217

$$2218 \quad \text{Tr}[A(\hat{\Sigma} + \lambda I)^{-1}] \asymp \frac{\kappa(\lambda)}{\lambda} \text{Tr}[A(\Sigma + \kappa(\lambda)I)^{-1}] \quad (15)$$

2220 and

$$\begin{aligned} 2221 \quad & \text{Tr}[A(\hat{\Sigma} + \lambda I)^{-1}B(\hat{\Sigma} + \lambda I)^{-1}] \\ 2222 \quad & \asymp \frac{\kappa(\lambda)^2}{\lambda^2} \text{Tr}[A(\Sigma + \kappa(\lambda)I)^{-1}B(\Sigma + \kappa(\lambda)I)^{-1}] \\ 2223 \quad & + \frac{\kappa(\lambda)^2}{\lambda^2} \frac{1}{n - \text{df}_2(\kappa(\lambda))} \text{Tr}[A(\Sigma + \kappa(\lambda)I)^{-2}\Sigma] \text{Tr}[B(\Sigma + \kappa(\lambda)I)^{-2}\Sigma]. \end{aligned} \quad (16)$$

2227 Equivalently,

$$2228 \quad \text{Tr}[A\hat{\Sigma}(\hat{\Sigma} + \lambda I)^{-1}] \asymp \text{Tr}[A\Sigma(\Sigma + \kappa(\lambda)I)^{-1}] \quad (17)$$

2230 and

$$\begin{aligned} 2231 \quad & \text{Tr}[A\hat{\Sigma}(\hat{\Sigma} + \lambda I)^{-1}B\hat{\Sigma}(\hat{\Sigma} + \lambda I)^{-1}] \\ 2232 \quad & \asymp \text{Tr}[A\Sigma(\Sigma + \kappa(\lambda)I)^{-1}B\Sigma(\Sigma + \kappa(\lambda)I)^{-1}] \\ 2233 \quad & + \frac{\kappa^2(\lambda)}{n - \text{df}_2(\kappa(\lambda))} \text{Tr}[A(\Sigma + \kappa(\lambda)I)^{-2}\Sigma] \text{Tr}[B(\Sigma + \kappa(\lambda)I)^{-2}\Sigma] \end{aligned} \quad (18)$$

2237 where $\kappa(\lambda)$ can be solved from self consistent equation above. Note given the unnormalized trace,
 2238 the trace equivalence \asymp shall be understood through convergence of ratio.

2239 **Two-point equivalence for resolvents of different argument** This can be further generalized to
 2240 equivalence with two variables,

$$2243 \quad \text{Tr}[A\hat{\Sigma}(\lambda + \hat{\Sigma})^{-1}B\hat{\Sigma}(\lambda' + \hat{\Sigma})^{-1}] \asymp \text{Tr}[AT_{\Sigma}BT'_{\Sigma}] + \quad (19)$$

$$2244 \quad \frac{\kappa\kappa'}{n - \text{df}_2(\kappa, \kappa')} \text{Tr}[AG_{\Sigma}\Sigma G'_{\Sigma}] \text{Tr}[G'_{\Sigma}\Sigma G_{\Sigma}B] \quad (20)$$

2247 where $T_{\Sigma} := \Sigma(\Sigma + \kappa)^{-1}$, $T'_{\Sigma} := \Sigma(\Sigma + \kappa')^{-1}$, $G_{\Sigma} := (\Sigma + \kappa)^{-1}$, $G'_{\Sigma} := (\Sigma + \kappa')^{-1}$. and
 2248 $\text{df}_2(\kappa, \kappa') := \text{Tr}[\Sigma^2 G_{\Sigma} G'_{\Sigma}]$. When $\kappa = \kappa'$ it recovers Eq.18.

2249 As a brief note for derivation, this follows from the deterministic equivalence for free product of
 2250 matrices $A * B$ stated in Appendix A of [Atanasov et al. \(2025\)](#). Set $A = \Sigma$ as population covariance,
 2251 $B = \frac{1}{n}ZZ^T$ as whitened data, then $A * B = \hat{\Sigma}$. Thus,

$$2253 \quad \hat{\Sigma}(\lambda + \hat{\Sigma})^{-1}M\hat{\Sigma}(\lambda' + \hat{\Sigma})^{-1} \asymp T_{\Sigma}MT'_{\Sigma} + \kappa\kappa'G_{\Sigma}\Sigma G'_{\Sigma} \frac{\text{Tr}[G'_{\Sigma}\Sigma G_{\Sigma}M]}{n - \text{df}_2(\kappa, \kappa')} \quad (21)$$

2255 Note that q in their convention correspond sto our γ and that their df definition is normalized trace.
 2256 We note in passing that this two-point equivalent was derived in [Atanasov et al. \(2025; 2024a\)](#) using
 2257 a diagrammatic moment-method argument; it could also be derived by extending the leave-one-out
 2258 arguments used by [Bach \(2024\)](#) to prove the deterministic equivalents with a single λ listed above.

2268 C.2 PROOF FOR DETERMINISTIC EQUIVALENCE OF DENOISER EXPECTATION (PROPOSITION 1)
22692270 **Proposition 5** (Main result, deterministic equivalence of the expectation of score and denoiser).
2271 *The optimal linear score and denoiser using empirical covariance has the following deterministic*
2272 *equivalence.*

2273
$$\begin{aligned} \mathbb{E}_{\hat{\Sigma}} \left[\mathbf{v}^\top \mathbf{D}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma) \right] &\asymp \mathbf{v}^\top \mathbf{x} + \mathbf{v}^\top \kappa(\sigma^2)(\Sigma + \kappa(\sigma^2)I)^{-1}(\boldsymbol{\mu} - \mathbf{x}) \\ 2274 &= \mathbf{v}^\top \boldsymbol{\mu} + \mathbf{v}^\top \Sigma(\Sigma + \kappa(\sigma^2)I)^{-1}(\mathbf{x} - \boldsymbol{\mu}) \\ 2275 \\ 2276 \mathbb{E}_{\hat{\Sigma}} \left[\mathbf{v}^\top \mathbf{s}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma) \right] &\asymp \frac{\kappa(\sigma^2)}{\sigma^2} \mathbf{v}^\top (\Sigma + \kappa(\sigma^2)I)^{-1}(\boldsymbol{\mu} - \mathbf{x}) \end{aligned}$$

2279 *Proof.* Per assumption, assume the sample mean $\hat{\boldsymbol{\mu}} = \boldsymbol{\mu}$, consider only the effect of empirical
2280 covariance $\hat{\Sigma}$,

2281
$$\mathbf{D}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma) = \mathbf{x} + \sigma^2(\hat{\Sigma} + \sigma^2 I)^{-1}(\boldsymbol{\mu} - \mathbf{x})$$

2282 Using the deterministic equivalence Eq. 15, 17, in the sense that the trace with any independent matrix
2283 converge in ratio at limit.

2284
$$\begin{aligned} (\hat{\Sigma} + \sigma^2 I)^{-1} &\asymp \frac{\kappa(\sigma^2)}{\sigma^2} (\Sigma + \kappa(\sigma^2)I)^{-1} \\ 2285 \hat{\Sigma}(\hat{\Sigma} + \sigma^2 I)^{-1} &\asymp \Sigma(\Sigma + \kappa(\sigma^2)I)^{-1} \end{aligned}$$

2286 Then, given the a fixed measurement vector \mathbf{v} , and a noised input \mathbf{x} , the projection of score onto a
2287 vector can be framed as trace. The equivalence reads,

2288
$$\begin{aligned} \mathbb{E}_{\hat{\Sigma}} \left[\mathbf{v}^\top \mathbf{s}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma) \right] &= \mathbb{E}_{\hat{\Sigma}} \left[\mathbf{v}^\top (\hat{\Sigma} + \sigma^2 I)^{-1}(\boldsymbol{\mu} - \mathbf{x}) \right] \\ 2289 &= \mathbb{E}_{\hat{\Sigma}} \text{Tr} \left[(\hat{\Sigma} + \sigma^2 I)^{-1}(\boldsymbol{\mu} - \mathbf{x}) \mathbf{v}^\top \right] \\ 2290 &\asymp \frac{\kappa(\sigma^2)}{\sigma^2} \text{Tr} \left[(\Sigma + \kappa(\sigma^2)I)^{-1}(\boldsymbol{\mu} - \mathbf{x}) \mathbf{v}^\top \right] \\ 2291 &= \frac{\kappa(\sigma^2)}{\sigma^2} \mathbf{v}^\top (\Sigma + \kappa(\sigma^2)I)^{-1}(\boldsymbol{\mu} - \mathbf{x}) \end{aligned}$$

2292 Similarly, use the other equivalence, the denoiser projection has equivalence,

2293
$$\begin{aligned} \mathbb{E}_{\hat{\Sigma}} \left[\mathbf{v}^\top \mathbf{D}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma) \right] &= \mathbf{v}^\top \boldsymbol{\mu} + \mathbb{E}_{\hat{\Sigma}} \left[\mathbf{v}^\top \hat{\Sigma}(\hat{\Sigma} + \sigma^2 I)^{-1}(\mathbf{x} - \boldsymbol{\mu}) \right] \\ 2294 &\asymp \mathbf{v}^\top \boldsymbol{\mu} + \mathbf{v}^\top \Sigma(\Sigma + \kappa(\sigma^2)I)^{-1}(\mathbf{x} - \boldsymbol{\mu}) \\ 2295 &= \mathbf{v}^\top \mathbf{D}_{\Sigma}^*(\mathbf{x}; \kappa^{1/2}) \end{aligned}$$

2296 Thus, in the expectation sense, the effect of empirical data covariance (finite data) on the denoiser, is
2297 equivalent to renormalizing and increasing the effective noise scale $\sigma^2 \rightarrow \kappa(\sigma^2)$, similar to adding
2298 an adaptive Ridge parameter. \square 2299 **Interpretation** Measuring the deviation of the empirical covariance denoiser from the population
2300 covariance denoiser, at the same noise scale,

2301
$$\begin{aligned} \mathbb{E}_{\hat{\Sigma}} \left[\mathbf{v}^\top (\mathbf{D}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma) - \mathbf{D}_{\Sigma}^*(\mathbf{x}; \sigma)) \right] \\ 2302 &\asymp \mathbf{v}^\top \left[\kappa(\sigma^2)(\Sigma + \kappa(\sigma^2)I)^{-1} - \sigma^2(\Sigma + \sigma^2 I)^{-1} \right] (\boldsymbol{\mu} - \mathbf{x}) \end{aligned}$$

2303 Using push through identity $A^{-1} - B^{-1} = A^{-1}(B - A)B^{-1}$,

2304
$$\begin{aligned} \kappa(\Sigma + \kappa I)^{-1} - \sigma^2(\Sigma + \sigma^2 I)^{-1} \\ 2305 &= \kappa\sigma^2(\Sigma + \kappa I)^{-1}(\Sigma + \sigma^2 I)^{-1} \left(\frac{1}{\sigma^2}(\Sigma + \sigma^2 I) - \frac{1}{\kappa}(\Sigma + \kappa I) \right) \\ 2306 &= (\kappa(\sigma^2) - \sigma^2)\Sigma(\Sigma + \kappa I)^{-1}(\Sigma + \sigma^2 I)^{-1} \end{aligned}$$

We can represent the deviation as resolvent product. This makes it clear that the deviation is proportional to the effect of renormalization ($\kappa(\sigma^2) - \sigma^2$).

$$\begin{aligned} & \mathbb{E}_{\hat{\Sigma}} \left[\mathbf{v}^\top (\mathbf{D}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma) - \mathbf{D}_{\Sigma}^*(\mathbf{x}; \sigma)) \right] \\ & \asymp (\kappa(\sigma^2) - \sigma^2) \mathbf{v}^\top \Sigma (\Sigma + \kappa I)^{-1} (\Sigma + \sigma^2 I)^{-1} (\boldsymbol{\mu} - \mathbf{x}) \end{aligned}$$

Setting the measurement vector along population eigenvector \mathbf{u}_k , with eigenvalue λ_k , then the deviation reads

$$\mathbb{E}_{\hat{\Sigma}} \left[\mathbf{u}_k^\top (\mathbf{D}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma) - \mathbf{D}_{\Sigma}^*(\mathbf{x}; \sigma)) \right] = \frac{\lambda_k(\kappa - \sigma^2)}{(\lambda_k + \sigma^2)(\lambda_k + \kappa)} \mathbf{u}_k^\top (\boldsymbol{\mu} - \mathbf{x})$$

It's easy to see the deviation affects lower eigenspace more.

C.3 PROOF FOR DETERMINISTIC EQUIVALENCE OF DENOISER FLUCTUATION (PROPOSITION 2)

Proof. Next, we examine the covariance of denoiser due to dataset realization, the score variance reads,

$$\begin{aligned} \mathcal{S}_s := Cov_{\hat{\Sigma}}[\mathbf{s}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma)] &= \mathbb{E}_{\hat{\Sigma}} \mathbf{s}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma) \mathbf{s}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma)^\top - \left(\mathbb{E}_{\hat{\Sigma}} \mathbf{s}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma) \right) \left(\mathbb{E}_{\hat{\Sigma}} \mathbf{s}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma) \right)^\top \\ &= \mathbb{E}_{\hat{\Sigma}} \left[(\hat{\Sigma} + \sigma^2 I)^{-1} (\boldsymbol{\mu} - \mathbf{x}) (\boldsymbol{\mu} - \mathbf{x})^\top (\hat{\Sigma} + \sigma^2 I)^{-1} \right] - \\ &\quad \mathbb{E}_{\hat{\Sigma}} \left[(\hat{\Sigma} + \sigma^2 I)^{-1} (\boldsymbol{\mu} - \mathbf{x}) \right] \mathbb{E}_{\hat{\Sigma}} \left[(\boldsymbol{\mu} - \mathbf{x})^\top (\hat{\Sigma} + \sigma^2 I)^{-1} \right] \\ &= \mathbb{E}_{\hat{\Sigma}} \left[(\hat{\Sigma} + \sigma^2 I)^{-1} (\boldsymbol{\mu} - \mathbf{x}) (\boldsymbol{\mu} - \mathbf{x})^\top (\hat{\Sigma} + \sigma^2 I)^{-1} \right] - \\ &\quad \mathbb{E}_{\hat{\Sigma}} \left[(\hat{\Sigma} + \sigma^2 I)^{-1} \right] (\boldsymbol{\mu} - \mathbf{x}) (\boldsymbol{\mu} - \mathbf{x})^\top \mathbb{E}_{\hat{\Sigma}} \left[(\hat{\Sigma} + \sigma^2 I)^{-1} \right] \end{aligned}$$

Note that the variance of denoiser and that of score has the simple scaling relationship, so we just need to study the score.

$$\mathcal{S}_{\mathbf{D}} = \sigma^4 \mathcal{S}_s$$

We are interested in the variance of score vector along a fixed probe vector \mathbf{v} ,

$$\begin{aligned} \mathbf{v}^\top \mathcal{S}_s \mathbf{v} &= Var_{\hat{\Sigma}}[\mathbf{v}^\top \mathbf{s}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma)] \\ &= \mathbb{E}_{\hat{\Sigma}} \left[\mathbf{v}^\top (\hat{\Sigma} + \sigma^2 I)^{-1} (\boldsymbol{\mu} - \mathbf{x}) (\boldsymbol{\mu} - \mathbf{x})^\top (\hat{\Sigma} + \sigma^2 I)^{-1} \mathbf{v} \right] - \left(\mathbf{v}^\top \mathbb{E}_{\hat{\Sigma}} \left[(\hat{\Sigma} + \sigma^2 I)^{-1} \right] (\boldsymbol{\mu} - \mathbf{x}) \right)^2 \\ &= \mathbb{E}_{\hat{\Sigma}} \underbrace{\text{Tr} \left[\mathbf{v} \mathbf{v}^\top (\hat{\Sigma} + \sigma^2 I)^{-1} (\boldsymbol{\mu} - \mathbf{x}) (\boldsymbol{\mu} - \mathbf{x})^\top (\hat{\Sigma} + \sigma^2 I)^{-1} \right]}_{2\text{nd moment}} - \underbrace{\left(\mathbb{E}_{\hat{\Sigma}} \text{Tr} \left[(\hat{\Sigma} + \sigma^2 I)^{-1} (\boldsymbol{\mu} - \mathbf{x}) \mathbf{v}^\top \right] \right)^2}_{1\text{st moment}} \end{aligned}$$

The two terms can be tackled by one-point and two-point equivalence Eq. 16,15. Abbreviating $A := \mathbf{v} \mathbf{v}^\top$, $B := (\boldsymbol{\mu} - \mathbf{x}) (\boldsymbol{\mu} - \mathbf{x})^\top$, $z := \sigma^2$.

$$\text{Tr}[A (\hat{\Sigma} + zI)^{-1}] \sim \frac{\kappa(z)}{z} \text{Tr}[A (\Sigma + \kappa(z)I)^{-1}]$$

$$\begin{aligned} \text{Tr}[A (\hat{\Sigma} + zI)^{-1} B (\hat{\Sigma} + zI)^{-1}] &\sim \frac{\kappa(z)^2}{z^2} \text{Tr}[A (\Sigma + \kappa(z)I)^{-1} B (\Sigma + \kappa(z)I)^{-1}] \\ &\quad + \frac{\kappa(z)^2}{z^2} \text{Tr}[A (\Sigma + \kappa(z)I)^{-2} \Sigma] \text{Tr}[B (\Sigma + \kappa(z)I)^{-2} \Sigma] \frac{1}{n - \text{df}_2(\kappa(z))} \end{aligned}$$

2376 The 2nd moment term is equivalent to,
2377

$$\begin{aligned}
2378 \quad & \text{Tr}[A(\hat{\Sigma} + zI)^{-1}B(\hat{\Sigma} + zI)^{-1}] \\
2379 \quad & \sim \frac{\kappa(z)^2}{z^2} \text{Tr}[\mathbf{v}\mathbf{v}^\top (\Sigma + \kappa(z)I)^{-1}(\boldsymbol{\mu} - \mathbf{x})(\boldsymbol{\mu} - \mathbf{x})^\top (\Sigma + \kappa(z)I)^{-1}] \\
2380 \quad & + \frac{\kappa(z)^2}{z^2} \frac{1}{n - \text{df}_2(\kappa(z))} \text{Tr}[\mathbf{v}\mathbf{v}^\top (\Sigma + \kappa(z)I)^{-2}\Sigma] \text{Tr}[(\boldsymbol{\mu} - \mathbf{x})(\boldsymbol{\mu} - \mathbf{x})^\top (\Sigma + \kappa(z)I)^{-2}\Sigma] \\
2381 \quad & = \frac{\kappa(z)^2}{z^2} \left(\mathbf{v}^\top (\Sigma + \kappa(z)I)^{-1}(\boldsymbol{\mu} - \mathbf{x}) \right)^2 \\
2382 \quad & + \frac{\kappa(z)^2}{z^2} \frac{1}{n - \text{df}_2(\kappa(z))} \left(\mathbf{v}^\top (\Sigma + \kappa(z)I)^{-2}\Sigma\mathbf{v} \right) \left((\boldsymbol{\mu} - \mathbf{x})^\top (\Sigma + \kappa(z)I)^{-2}\Sigma(\boldsymbol{\mu} - \mathbf{x}) \right)
\end{aligned}$$

2388 The first moment term is equivalent to,
2389

$$\begin{aligned}
2390 \quad & \text{Tr}[(\hat{\Sigma} + zI)^{-1}(\boldsymbol{\mu} - \mathbf{x})\mathbf{v}^\top] \sim \frac{\kappa(z)}{z} \text{Tr}[(\Sigma + \kappa(z)I)^{-1}(\boldsymbol{\mu} - \mathbf{x})\mathbf{v}^\top] \\
2391 \quad & = \frac{\kappa(z)}{z} \mathbf{v}^\top (\Sigma + \kappa(z)I)^{-1}(\boldsymbol{\mu} - \mathbf{x})
\end{aligned}$$

2395 Thus, combining the two terms, we obtain the variance of score at noised datapoint \mathbf{x} , along direction
2396 \mathbf{v} ,

$$\begin{aligned}
2397 \quad & \mathbf{v}^\top \mathcal{S}_s(\mathbf{x})\mathbf{v} = \text{Var}_{\hat{\Sigma}}[\mathbf{v}^\top \mathbf{s}_{\hat{\Sigma}}^*(\mathbf{x}; \sigma)] \\
2398 \quad & = \mathbb{E}_{\hat{\Sigma}} \text{Tr}[\mathbf{v}\mathbf{v}^\top (\hat{\Sigma} + \sigma^2 I)^{-1}(\boldsymbol{\mu} - \mathbf{x})(\boldsymbol{\mu} - \mathbf{x})^\top (\hat{\Sigma} + \sigma^2 I)^{-1}] \\
2399 \quad & - \left(\mathbb{E}_{\hat{\Sigma}} \text{Tr}[(\hat{\Sigma} + \sigma^2 I)^{-1}(\boldsymbol{\mu} - \mathbf{x})\mathbf{v}^\top] \right)^2 \\
2400 \quad & \sim \frac{\kappa(z)^2}{z^2} \left(\mathbf{v}^\top (\Sigma + \kappa(z)I)^{-1}(\boldsymbol{\mu} - \mathbf{x}) \right)^2 \\
2401 \quad & + \frac{\kappa(z)^2}{z^2} \left(\mathbf{v}^\top (\Sigma + \kappa(z)I)^{-2}\Sigma\mathbf{v} \right) \left((\boldsymbol{\mu} - \mathbf{x})^\top (\Sigma + \kappa(z)I)^{-2}\Sigma(\boldsymbol{\mu} - \mathbf{x}) \right) \frac{1}{n - \text{df}_2(\kappa(z))} \\
2402 \quad & - \left(\frac{\kappa(z)}{z} \mathbf{v}^\top (\Sigma + \kappa(z)I)^{-1}(\boldsymbol{\mu} - \mathbf{x}) \right)^2 \\
2403 \quad & = \frac{1}{n - \text{df}_2(\kappa(z))} \frac{\kappa(z)^2}{z^2} \left(\mathbf{v}^\top (\Sigma + \kappa(z)I)^{-2}\Sigma\mathbf{v} \right) \left((\boldsymbol{\mu} - \mathbf{x})^\top (\Sigma + \kappa(z)I)^{-2}\Sigma(\boldsymbol{\mu} - \mathbf{x}) \right) \\
2404 \quad & (z \mapsto \sigma^2) = \frac{1}{n - \text{df}_2(\kappa(\sigma^2))} \frac{\kappa(\sigma^2)^2}{\sigma^4} \left(\mathbf{v}^\top (\Sigma + \kappa(\sigma^2)I)^{-2}\Sigma\mathbf{v} \right) \left((\boldsymbol{\mu} - \mathbf{x})^\top (\Sigma + \kappa(\sigma^2)I)^{-2}\Sigma(\boldsymbol{\mu} - \mathbf{x}) \right)
\end{aligned}$$

2414 Per simple scaling, the variance of denoisers reads,
2415

$$\begin{aligned}
2416 \quad & \mathbf{v}^\top \mathcal{S}_D(\mathbf{x})\mathbf{v} = \sigma^4 \mathbf{v}^\top \mathcal{S}_s(\mathbf{x})\mathbf{v} \\
2417 \quad & \sim \frac{\kappa(\sigma^2)^2}{n - \text{df}_2(\kappa(\sigma^2))} \underbrace{\left(\mathbf{v}^\top (\Sigma + \kappa(\sigma^2)I)^{-2}\Sigma\mathbf{v} \right)}_{\square(\mathbf{v}, \kappa, \Sigma)} \underbrace{\left((\boldsymbol{\mu} - \mathbf{x})^\top (\Sigma + \kappa(\sigma^2)I)^{-2}\Sigma(\boldsymbol{\mu} - \mathbf{x}) \right)}_{\square(\boldsymbol{\mu} - \mathbf{x}, \kappa, \Sigma)}
\end{aligned}$$

2421 \square
2422

2423 C.3.1 INTERPRETATION AND DERIVATIONS

2424 **Dependency on probe direction \mathbf{v}** This dependency on \mathbf{v} tells us about the *anisotropy of uncertainty*, or variance of the score / denoiser prediction on different directions.
2425

$$\begin{aligned}
2426 \quad & \square(\mathbf{v}, \kappa, \Sigma) := \mathbf{v}^\top (\Sigma + \kappa(\sigma^2)I)^{-2}\Sigma\mathbf{v} \\
2427 \quad & = \mathbf{v}^\top U \frac{\Lambda}{(\Lambda + \kappa(\sigma^2))^2} U^\top \mathbf{v}
\end{aligned}$$

2430 Per assumption the probe vector \mathbf{v} is unit vector. Then this dependency is decided by the diagonal
 2431 matrix $\frac{\Lambda}{(\Lambda + \kappa(\sigma^2))^2} = \frac{\lambda_k}{(\lambda_k + \kappa(\sigma^2))^2}$.
 2432

2433 Consider when the probing vector is aligned exactly with the k th eigenvector \mathbf{u}_k , this term reads
 2434

$$\begin{aligned} 2436 \quad \square(\mathbf{u}_k, \kappa, \Sigma) &= \mathbf{u}_k^\top (\Sigma + \kappa(\sigma^2)I)^{-2} \Sigma \mathbf{u}_k \\ 2437 &= \frac{\lambda_k}{(\lambda_k + \kappa(\sigma^2))^2} \\ 2438 &= \chi(\lambda_k, \kappa(\sigma^2)) \end{aligned}$$

2442 We can discuss the different regime of $\chi(\lambda_k, \kappa)$ depending on λ_k and $\kappa(\sigma^2)$
 2443

- 2445 • High noise regime $\lambda_k \ll \kappa$: $\chi(\lambda_k, \kappa) \approx \frac{\lambda_k}{\kappa^2}$, so $\chi(\lambda_k, \kappa)$ will increase with λ_k . Higher
 2446 variance directions have larger uncertainties.
 2447
- 2448 • Low noise regime $\lambda_k \gg \kappa$: $\chi(\lambda_k, \kappa) \approx \frac{1}{\lambda_k}$, so $\chi(\lambda_k, \kappa)$ will decrease with λ_k . Lower
 2449 variance directions have larger uncertainties!
 2450
- 2451 • Regarding κ , $\chi(\lambda_k, \kappa)$ is monotonic decreasing with κ , i.e. higher the noise scale, the
 2452 smaller the variance.
 2453
- 2454 • Regarding λ_k , $\chi(\lambda_k, \kappa)$ has one unique maximum, where $\arg \max_{\lambda} \chi(\lambda, \kappa) = \kappa$, and
 2455 $\max_{\lambda} \chi(\lambda, \kappa) = \frac{1}{4\kappa}$. So it's a bell shaped function of λ_k . (Proof below.)
 2456
- 2457 – This shows that at different noise level or $\kappa(\sigma^2)$, there is always some direction with
 2458 variance comparable to $\kappa(\sigma^2)$ which will have the largest variance!
 2459
- 2460 – Further the largest variance will be inverse proportional to $\kappa(\sigma^2)$, i.e. generally larger
 2461 variance at lower noise case.
 2462

2464 This result is definitely not obvious! It shows that the anisotropy of the uncertainty depends on the
 2465 renormalized noise scale $\kappa(\sigma^2)$, and the maximal uncertainty are focused around the PC dimensions
 2466 with variance similar to $\kappa(\sigma^2)$.
 2467

2468 **Proof of unique maximum of $\chi(\lambda, \kappa)$**

2469 Given

$$2470 \quad \chi(\lambda, \kappa) = \frac{\lambda}{(\lambda + \kappa)^2}$$

2471 Then

$$\begin{aligned} 2472 \quad \frac{d\chi(\lambda, \kappa)}{d\lambda} &= \frac{(\lambda + \kappa)^2 - 2(\lambda + \kappa)\lambda}{(\lambda + \kappa)^4} \\ 2473 &= \frac{\kappa - \lambda}{(\lambda + \kappa)^3} \end{aligned}$$

2474 Setting gradient to zero yield unique stationary point, $\kappa = \lambda$.
 2475

2476 Given $\kappa, \lambda > 0$, we have the unique maximum w.r.t. λ .
 2477

$$2478 \quad \arg \max_{\lambda} \chi(\lambda, \kappa) = \kappa$$

$$2479 \quad \max_{\lambda} \chi(\lambda, \kappa) = \frac{1}{4\kappa}$$

2484 **Dependency on the probe point \mathbf{x} .** The dependency on probe point \mathbf{x} tells us about the spatial
 2485 in-homogeneity of the uncertainty.
 2486

$$\begin{aligned}
 2487 \quad \square(\mathbf{x} - \boldsymbol{\mu}, \kappa, \boldsymbol{\Sigma}) &= (\mathbf{x} - \boldsymbol{\mu})^\top (\boldsymbol{\Sigma} + \kappa(\sigma^2)I)^{-2} \boldsymbol{\Sigma} (\mathbf{x} - \boldsymbol{\mu}) \\
 2488 &= (\mathbf{x} - \boldsymbol{\mu})^\top U \frac{\Lambda}{(\Lambda + \kappa(\sigma^2))^2} U^\top (\mathbf{x} - \boldsymbol{\mu}) \\
 2489 &= \sum_k \frac{\lambda_k}{(\lambda_k + \kappa(\sigma^2))^2} (\mathbf{u}_k^\top (\mathbf{x} - \boldsymbol{\mu}))^2 \\
 2490 &= \sum_k \chi(\lambda_k, \kappa(\sigma^2)) (\mathbf{u}_k^\top (\mathbf{x} - \boldsymbol{\mu}))^2
 \end{aligned}$$

2491
 2492
 2493
 2494
 2495
 2496
 2497
 2498
 2499 This is similar to the dependency above, except that now our argument $\mathbf{x} - \boldsymbol{\mu}$ is no longer unit
 2500 normed, but any probing direction in the sample space.

2501 Note, generally the noised sample \mathbf{x} from a certain realization of dataset is distributed like $\mathcal{N}(\boldsymbol{\mu}, \hat{\boldsymbol{\Sigma}} +$
 2502 $\sigma^2 I)$ (under Gaussian data assumption), so
 2503

$$2504 \quad \mathbf{v}^\top (\mathbf{x} - \boldsymbol{\mu}) \sim \mathcal{N}(0, \mathbf{v}^\top (\hat{\boldsymbol{\Sigma}} + \sigma^2 I) \mathbf{v})$$

2505
 2506
 2507
 2508 Consider a probe point on the hyper elliptical shell defined by $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma} + \sigma^2 I)$, then if the point falls
 2509 on the line $\mathbf{x} = \boldsymbol{\mu} + c\mathbf{u}_k$. $\|\mathbf{x} - \boldsymbol{\mu}\|^2 = c^2 \approx \sigma^2 + \lambda_k$
 2510

2511 Then

$$\begin{aligned}
 2512 \quad \square(\mathbf{x} - \boldsymbol{\mu}, \kappa, \boldsymbol{\Sigma}) &= \square(c\mathbf{u}_k, \kappa, \boldsymbol{\Sigma}) \\
 2513 &= c^2 \frac{\lambda_k}{(\lambda_k + \kappa(\sigma^2))^2} \\
 2514 &\approx \frac{(\lambda_k + \sigma^2)\lambda_k}{(\lambda_k + \kappa(\sigma^2))^2} \\
 2515 &= (\sigma^2 + \lambda_k) \chi(\lambda_k, \kappa(\sigma^2)) \\
 2516 &= \xi(\lambda_k, \sigma^2)
 \end{aligned}$$

2517
 2518
 2519
 2520
 2521
 2522
 2523
 2524 • **High noise regime**, $\kappa > \sigma^2 \gg \lambda$, then $\xi(\lambda, \sigma^2) \approx \frac{\sigma^2 \lambda}{\kappa^2(\sigma^2)} < \frac{\lambda}{\kappa(\sigma^2)} \ll 1$, which scale
 2525 linearly with PC variance λ , higher the PC, larger the variance.
 2526

2527
 2528 • **Low noise regime**, $\lambda \gg \kappa > \sigma^2$, then $\xi(\lambda, \sigma^2) \approx 1$. Then all points on the ellipsoid have
 2529 large variance.
 2530

2531 • At any fixed σ^2 , this function monotonically increase with λ_k .
 2532

2533
 2534 – The score or denoiser variance is larger when the probing point $\boldsymbol{\mu} + c\mathbf{u}_k$ is deviating
 2535 along those higher variance directions \mathbf{u}_k .
 2536
 2537 – When the probing point is deviating along low variance directions, the variance is
 2538 lower.

2538
2539**Derivation of properties of $\xi(\lambda, \sigma^2)$** 2540
2541
25422543
Derivative2544
2545
2546
2547
2548
2549
2550
2551
2552
2553
2554

$$\begin{aligned}\frac{d\xi(\lambda, \sigma^2)}{d\lambda} &= \frac{(\sigma^2 + 2\lambda)(\lambda + \kappa(\sigma^2))^2 - 2(\lambda + \kappa(\sigma^2))(\sigma^2 + \lambda)\lambda}{(\lambda + \kappa(\sigma^2))^4} \\ &= \frac{(\sigma^2 + 2\lambda)(\lambda + \kappa(\sigma^2)) - 2(\sigma^2 + \lambda)\lambda}{(\lambda + \kappa(\sigma^2))^3} \\ &= \frac{(\sigma^2 + 2\lambda)\kappa(\sigma^2) - \lambda\sigma^2}{(\lambda + \kappa(\sigma^2))^3} \\ &= \frac{\sigma^2\kappa(\sigma^2) + (2\kappa(\sigma^2) - \sigma^2)\lambda}{(\lambda + \kappa(\sigma^2))^3}\end{aligned}$$

2555 Note that through the self consistent equation $\kappa(\sigma^2) - \sigma^2 > 0$, thus $\frac{d\xi(\lambda, \sigma^2)}{d\lambda} > 0, \forall \lambda$. The function
2556 is monotonically increasing for λ .2557 Given that $\kappa(\sigma^2) > \sigma^2 > 0$, we have bounds2558
2559
2560

$$\xi(\lambda, \sigma^2) = \frac{(\sigma^2 + \lambda)\lambda}{(\lambda + \kappa(\sigma^2))^2} < \frac{\lambda}{\lambda + \kappa(\sigma^2)} < 1$$

2561
2562
2563**Overall scaling with sample** Finally, we marginalize over space and direction, obtaining an overall quantification of consistency of denoiser, and study its scaling property.2564
2565First, *marginalizing (summing) all directions*, we have2566
2567
2568
2569
2570

$$\begin{aligned}\sum_k \square(\mathbf{u}_k, \kappa, \Sigma) &= \sum_k \mathbf{u}_k^\top (\Sigma + \kappa(\sigma^2)I)^{-2} \Sigma \mathbf{u}_k \\ &= \text{Tr}[(\Sigma + \kappa(\sigma^2)I)^{-2} \Sigma]\end{aligned}$$

2571
2572
2573
2574
2575
2576
2577
2578
2579
2580
2581

This can be further abbreviated as following,

$$\begin{aligned}\sum_k \square(\mathbf{u}_k, \kappa, \Sigma) &= \text{Tr}[(\Sigma + \kappa(\sigma^2)I)^{-2} \Sigma I] \\ &= \text{Tr}[(\Sigma + \kappa(\sigma^2)I)^{-2} \Sigma \frac{1}{\kappa(\sigma^2)} (\Sigma + \kappa(\sigma^2)I - \Sigma)] \\ &= \frac{1}{\kappa(\sigma^2)} (\text{Tr}[(\Sigma + \kappa(\sigma^2)I)^{-1} \Sigma] - \text{Tr}[(\Sigma + \kappa(\sigma^2)I)^{-2} \Sigma^2]) \\ &= \frac{\text{df}_1(\kappa) - \text{df}_2(\kappa)}{\kappa}\end{aligned}$$

2582
2583
2584Next, *marginalize (averaging) over space*. Here we consider the noised distribution starting from the true target distribution $\mathbf{x} \sim p(\mathbf{x}; \sigma) = p_0(\mathbf{x}) * \mathcal{N}(0, \sigma^2 I)$. For us, the only thing matter is the 2nd moment, so for arbitrary distribution we have,2585
2586
2587

2588 Thus,

2589
2590
2591

$$\begin{aligned}\mathbb{E}_{\mathbf{x}} \square(\mathbf{x} - \mu, \kappa, \Sigma) &= (\mathbf{x} - \mu)^\top (\Sigma + \kappa(\sigma^2)I)^{-2} \Sigma (\mathbf{x} - \mu) \\ &= \text{Tr}[(\Sigma + \sigma^2 I)(\Sigma + \kappa(\sigma^2)I)^{-2} \Sigma]\end{aligned}$$

This can also be abbreviated using degree of freedom,

$$\begin{aligned}
\mathbb{E}_{\mathbf{x}} \square(\mathbf{x} - \boldsymbol{\mu}, \kappa, \boldsymbol{\Sigma}) &= \text{Tr} \left[(\boldsymbol{\Sigma} + \sigma^2 I) (\boldsymbol{\Sigma} + \kappa(\sigma^2) I)^{-2} \boldsymbol{\Sigma} \right] \\
&= \text{Tr} \left[(\boldsymbol{\Sigma} + \kappa(\sigma^2) I)^{-2} \boldsymbol{\Sigma}^2 \right] + \sigma^2 \text{Tr} \left[(\boldsymbol{\Sigma} + \kappa(\sigma^2) I)^{-2} \boldsymbol{\Sigma} \right] \\
&= \text{df}_2(\kappa) + \frac{\sigma^2}{\kappa} (\text{df}_1(\kappa) - \text{df}_2(\kappa)) \\
&= \frac{\sigma^2}{\kappa} \text{df}_1(\kappa) + \left(1 - \frac{\sigma^2}{\kappa}\right) \text{df}_2(\kappa)
\end{aligned}$$

Thus, we have

$$\begin{aligned}
\mathbb{E}_{\mathbf{x}} \sum_k \mathbf{u}_k^\top \mathcal{S}_D(\mathbf{x}) \mathbf{u}_k &\asymp \frac{\kappa(\sigma^2)^2}{n - \text{df}_2(\kappa(\sigma^2))} \sum_k \underbrace{\left(\mathbf{u}_k^\top (\Sigma + \kappa(\sigma^2)I)^{-2} \Sigma \mathbf{u}_k \right)}_{\square(\mathbf{v}, \kappa, \Sigma)} \underbrace{\mathbb{E}_{\mathbf{x}} \left((\mu - \mathbf{x})^\top (\Sigma + \kappa(\sigma^2)I)^{-2} \Sigma (\mu - \mathbf{x}) \right)}_{\square(\mu - \mathbf{x}, \kappa, \Sigma)} \\
&= \frac{\kappa(\sigma^2)^2}{n - \text{df}_2(\kappa(\sigma^2))} \text{Tr} \left[(\Sigma + \kappa(\sigma^2)I)^{-2} \Sigma \right] \text{Tr} \left[(\Sigma + \sigma^2 I) (\Sigma + \kappa(\sigma^2)I)^{-2} \Sigma \right] \\
&= \frac{\kappa(\sigma^2)^2}{n - \text{df}_2(\kappa(\sigma^2))} \times \frac{\text{df}_1(\kappa) - \text{df}_2(\kappa)}{\kappa} \times \left(\frac{\sigma^2}{\kappa} \text{df}_1(\kappa) + \left(1 - \frac{\sigma^2}{\kappa}\right) \text{df}_2(\kappa) \right) \\
&= \frac{(\text{df}_1(\kappa) - \text{df}_2(\kappa)) \times (\sigma^2 \text{df}_1(\kappa) + (\kappa - \sigma^2) \text{df}_2(\kappa))}{n - \text{df}_2(\kappa(\sigma^2))} \\
&=: \Delta(n, \sigma^2, \Lambda)
\end{aligned}$$

Now, marginalized over space and direction, this is only a function of the population spectrum, sample number and noise scale. Note n is the sample number, so it makes sense when n goes to infinity, then $\hat{\Sigma} \rightarrow \Sigma$ and $\kappa \rightarrow \sigma^2$, the variance reduce to zero.

Basically the higher the κ , the smaller the $df_2(\kappa)$, so $n - df_2(\kappa(\sigma^2))$ will be larger, which scale down $\frac{1}{n - df_2(\kappa(\sigma^2))} \frac{\kappa(\sigma^2)^2}{\sigma^4}$.

Note, when we compare our theory with the empirical measurement of deviation of denoiser or samples between the two splits, we used the following lemma to use the variance to predict the expected MSE deviation.

Lemma 1 (Expected MSE between two i.i.d. samples doubles the variance). *Let X, Y be i.i.d. random variables with variance $S = \text{Var}(X)$. Then their mean squared error (MSE) is double the variance.*

$$\mathbb{E}[(X - Y)^2] = 2S.$$

Proof. Expanding and using independence,

$$\mathbb{E}[(X - Y)^2] = \mathbb{E}[X^2] + \mathbb{E}[Y^2] - 2\mathbb{E}[XY] = 2\mathbb{E}[X^2] - 2\mathbb{E}[X]\mathbb{E}[Y].$$

Since $\text{Var}(X) = \mathbb{E}[X^2] - (\mathbb{E}[X])^2 = S$, this simplifies to $2S$.

C.4 INTEGRAL REPRESENTATION OF MATRIX FRACTIONAL POWER (BALAKRISHNAN FORMULA)

Lemma 2 (Scalar beta integral identity). *The integral identity*

$$\int_0^\infty \frac{t^{-\alpha} dt}{\lambda + t} = \frac{\pi}{\sin(\pi\alpha)} \lambda^{-\alpha}, \quad \alpha \in (0, 1)$$

2646 *Proof.* Recall the definition of Beta function,
 2647

$$\begin{aligned} 2648 \quad B(p, q) &= \int_0^1 u^{p-1} (1-u)^{q-1} du \\ 2649 \\ 2650 \quad &= \frac{\Gamma(p)\Gamma(q)}{\Gamma(p+q)} \\ 2651 \\ 2652 \end{aligned}$$

2653 We can turn it into beta function via change of variable $u = \frac{t}{\lambda+t}$, then $t \in [0, \infty)$ maps to $u \in [0, 1)$.
 2654

$$\begin{aligned} 2655 \quad t &= \frac{u\lambda}{1-u} \\ 2656 \\ 2657 \quad dt &= \frac{\lambda}{(1-u)^2} du \\ 2658 \\ 2659 \end{aligned}$$

$$\begin{aligned} 2660 \quad \int_0^\infty \frac{t^{-\alpha} dt}{\lambda+t} &= \int_0^\infty \left(\frac{t}{t+\lambda}\right) t^{-1-\alpha} dt \\ 2661 \\ 2662 \quad &= \int_0^1 u \left(\frac{u\lambda}{1-u}\right)^{-1-\alpha} \frac{\lambda}{(1-u)^2} du \\ 2663 \\ 2664 \quad &= \lambda^{-\alpha} \int_0^1 u^{-\alpha} (1-u)^{\alpha-1} du \\ 2665 \\ 2666 \quad &= \lambda^{-\alpha} B(1-\alpha, \alpha) \\ 2667 \\ 2668 \end{aligned}$$

2669 and using Euler's reflection formula, we have

$$2670 \quad B(1-\alpha, \alpha) = \frac{\pi}{\sin(\pi\alpha)}$$

2673 Thus,

$$2674 \quad \int_0^\infty \frac{t^{-\alpha} dt}{\lambda+t} = \frac{\pi}{\sin(\pi\alpha)} \lambda^{-\alpha}$$

2676 \square

2677 **Corollary 1** (Integral formula for power one half). *In the special case of $\alpha = 1/2$*

$$2679 \quad \pi \lambda^{-1/2} = \int_0^\infty \frac{t^{-1/2} dt}{\lambda+t} = 2 \int_0^\infty \frac{ds}{\lambda+s^2}$$

2682 *Proof.* Use simple change of variable $t \rightarrow s^2$,

$$\begin{aligned} 2684 \quad \pi \lambda^{-1/2} &= \int_0^\infty \frac{t^{-1/2} dt}{\lambda+t} \\ 2685 \\ 2686 \quad &= \int_0^\infty \frac{s^{-1} ds^2}{\lambda+s^2} \\ 2687 \\ 2688 \quad &= \int_0^\infty \frac{2ds}{\lambda+s^2} \\ 2689 \\ 2690 \end{aligned}$$

2691 \square

2692 **Corollary 2** (Integral representation of fractional matrix power). *The matrix version of such identity,
 2693 for self-adjoint, positive semi definite matrix $A \succeq 0$,*

$$2695 \quad \int_0^\infty (A+tI)^{-1} t^{-\alpha} dt = \frac{\pi}{\sin(\pi\alpha)} A^{-\alpha}, \quad \alpha \in (0, 1)$$

2697 *Similarly, for $z \geq 0, z \in \mathbb{R}$,*

$$2699 \quad \int_0^\infty (A+(z+t)I)^{-1} t^{-\alpha} dt = \frac{\pi}{\sin(\pi\alpha)} (A+zI)^{-\alpha}, \quad \alpha \in (0, 1)$$

2700 **Corollary 3** (Integral representation of matrix one half). *The matrix version of such identity, for*
 2701 *self-adjoint, positive semi definite matrix $A \succeq 0$,*
 2702

$$2703 A^{-1/2} = \frac{1}{\pi} \int_0^\infty (A + tI)^{-1} t^{-1/2} dt = \frac{2}{\pi} \int_0^\infty (A + s^2 I)^{-1} ds$$

2705 **Lemma 3** (Resolvent Identity). *When $u \neq s$, we have identity*
 2706

$$2707 (A + sI)^{-1}(A + uI)^{-1} = \frac{1}{s - u} \left((A + uI)^{-1} - (A + sI)^{-1} \right)$$

$$2710 A(A + sI)^{-1}(A + uI)^{-1} = \frac{1}{s - u} \left(A(A + uI)^{-1} - A(A + sI)^{-1} \right)$$

$$2711 = \frac{s(A + sI)^{-1} - u(A + uI)^{-1}}{s - u}$$

2715 *Proof.* Note that

$$2717 \begin{aligned} & \left((A + sI) - (A + uI) \right) (A + sI)^{-1}(A + uI)^{-1} \\ &= (A + uI)^{-1} - (A + sI)^{-1} \\ &= (s - u)(A + sI)^{-1}(A + uI)^{-1} \end{aligned}$$

2722 Thus,

$$2724 (A + sI)^{-1}(A + uI)^{-1} = \frac{1}{(s - u)} \left((A + uI)^{-1} - (A + sI)^{-1} \right)$$

2726 as corollary

$$2728 A(A + sI)^{-1}(A + uI)^{-1} = \frac{1}{s - u} \left(A(A + uI)^{-1} - A(A + sI)^{-1} \right)$$

$$2730 = \frac{1}{s - u} \left(I - u(A + uI)^{-1} - I + s(A + sI)^{-1} \right)$$

$$2732 = \frac{s(A + sI)^{-1} - u(A + uI)^{-1}}{s - u}$$

2735 Note that this formula has no real pole, and it behaves nicely when denominator vanishes, and the
 2736 RHS becomes a derivative.

$$2738 \lim_{s \rightarrow u} \frac{s(A + sI)^{-1} - u(A + uI)^{-1}}{s - u} = \frac{d}{ds} s(A + sI)^{-1}$$

$$2740 = (A + sI)^{-1} - s(A + sI)^{-2}$$

$$2741 = A(A + sI)^{-2}$$

$$2743 \lim_{s \rightarrow u} \frac{1}{(s - u)} \left((A + uI)^{-1} - (A + sI)^{-1} \right) = -\frac{d}{du} (A + uI)^{-1}$$

$$2744 = (A + uI)^{-2}$$

2747 \square

2749 **C.5 PROOF FOR EXPECTATION OF THE SAMPLING MAPPING (APPROXIMATE VERSION,
 2750 INFINITE σ_T , PROPOSITION 3)**

2752 Using empirical covariance and mean to realize the sampling, we have

$$2753 \mathbf{x}(\mathbf{x}_{\sigma_T}, \sigma_0) = \hat{\boldsymbol{\mu}} + (\hat{\boldsymbol{\Sigma}} + \sigma_0^2 I)^{1/2} (\hat{\boldsymbol{\Sigma}} + \sigma_T^2 I)^{-1/2} (\mathbf{x}_{\sigma_T} - \hat{\boldsymbol{\mu}})$$

2754 For the final sampling outcome $\sigma_0 \rightarrow 0$, this reads
 2755

$$2756 \mathbf{x}(\mathbf{x}_{\sigma_T}, 0) = \hat{\boldsymbol{\mu}} + \hat{\boldsymbol{\Sigma}}^{1/2}(\hat{\boldsymbol{\Sigma}} + \sigma_T^2 I)^{-1/2}(\mathbf{x}_{\sigma_T} - \hat{\boldsymbol{\mu}})$$

2758 As before, assume the sample mean equals the population one, then the finite sample effect comes
 2759 from the matrix $\hat{\boldsymbol{\Sigma}}^{1/2}(\hat{\boldsymbol{\Sigma}} + \sigma_T^2 I)^{-1/2}$

$$2761 \mathbf{x}(\mathbf{x}_{\sigma_T}, 0) = \boldsymbol{\mu} + \hat{\boldsymbol{\Sigma}}^{1/2}(\hat{\boldsymbol{\Sigma}} + \sigma_T^2 I)^{-1/2}(\mathbf{x}_{\sigma_T} - \boldsymbol{\mu})$$

2762 Note that for sampling, under EDM convention, the initial noise are sampled with variance $\sigma_T^2 I$,
 2763 $\mathbf{x}_{\sigma_T} \sim \mathcal{N}(0, \sigma_T^2 I)$, notably for practical diffusion models, initial noise variances are large, $\sigma_T^2 \sim$
 2765 6000. Thus we can define a normalized initial noise $\bar{\mathbf{x}} = (\mathbf{x}_{\sigma_T} - \boldsymbol{\mu})/\sigma_T$.

2766 As a large initial noise limit, given that $\boldsymbol{\Sigma}$ has finite spectral norm,

$$2768 \lim_{\sigma \rightarrow \infty} \sigma \boldsymbol{\Sigma}^{1/2}(\boldsymbol{\Sigma} + \sigma^2 I)^{-1/2} = \boldsymbol{\Sigma}^{1/2}$$

2770 and when $\sigma_T \rightarrow \infty$ the normalized initial noise are sampled from standard Gaussian, $\bar{\mathbf{x}} \sim \mathcal{N}(0, I)$.

2771 Equivalently, we can consider expansion as orders of $1/\sigma$,

$$2773 \sigma \boldsymbol{\Sigma}^{1/2}(\boldsymbol{\Sigma} + \sigma^2 I)^{-1/2} = \boldsymbol{\Sigma}^{1/2}(I + \frac{1}{\sigma^2} \boldsymbol{\Sigma})^{-1/2} \\ 2774 \approx \boldsymbol{\Sigma}^{1/2}(I - \frac{1}{2} \frac{1}{\sigma^2} \boldsymbol{\Sigma} + \dots) \\ 2775 \approx \boldsymbol{\Sigma}^{1/2} - \frac{1}{2} \frac{1}{\sigma^2} \boldsymbol{\Sigma}^{3/2} + \dots$$

2779 If we keep the zeroth-order term, then we get the approximation

$$2781 \sigma \boldsymbol{\Sigma}^{1/2}(\boldsymbol{\Sigma} + \sigma^2 I)^{-1/2} \approx \boldsymbol{\Sigma}^{1/2}$$

2783 Consider approximation,

$$2785 \mathbf{x}(\mathbf{x}_{\sigma_T}, 0) = \boldsymbol{\mu} + \hat{\boldsymbol{\Sigma}}^{1/2}(\hat{\boldsymbol{\Sigma}} + \sigma_T^2 I)^{-1/2}(\mathbf{x}_{\sigma_T} - \boldsymbol{\mu}) \\ 2786 \approx \boldsymbol{\mu} + \hat{\boldsymbol{\Sigma}}^{1/2}(\frac{\mathbf{x}_{\sigma_T} - \boldsymbol{\mu}}{\sigma_T}) \\ 2787 = \boldsymbol{\mu} + \hat{\boldsymbol{\Sigma}}^{1/2} \bar{\mathbf{x}}$$

2790 then we can study the effect of finite sample on sampling mapping via the matrix $\hat{\boldsymbol{\Sigma}}^{1/2}$.

2792 **Proposition 6.** *Deterministic equivalence of empirical covariance matrix one half*

$$2794 \hat{\boldsymbol{\Sigma}}^{1/2} = \frac{2}{\pi} \int_0^\infty \hat{\boldsymbol{\Sigma}}(\hat{\boldsymbol{\Sigma}} + u^2 I)^{-1} du \quad (23)$$

$$2796 \hat{\boldsymbol{\Sigma}}^{1/2} \approx \frac{2}{\pi} \int_0^\infty \boldsymbol{\Sigma}(\boldsymbol{\Sigma} + \kappa(u^2) I)^{-1} du \quad (24)$$

2799 *Proof.* Combining Lemma 3 with deterministic equivalence of one point \ref{ref}

2800 \square

2801 This result can be compared to population covariance half, when renormalization effect vanish
 2802 $\kappa(u^2) \rightarrow u^2$.

$$2803 \boldsymbol{\Sigma}^{1/2} = \frac{2}{\pi} \int_0^\infty \boldsymbol{\Sigma}(\boldsymbol{\Sigma} + u^2 I)^{-1} du$$

2806 Since $\kappa(u^2) > u^2$ point by point in the integral, the sample version leads to larger shrinkage.

$$2807 \mathbf{v}^\top \hat{\boldsymbol{\Sigma}}^{1/2} \mathbf{v} < \mathbf{v}^\top \boldsymbol{\Sigma}^{1/2} \mathbf{v}$$

2808 Concretely, if we measure along spectral modes \mathbf{u}_k of population covariance,
 2809

$$\begin{aligned} 2810 \quad \mathbf{u}_k^\top \hat{\Sigma}^{1/2} \mathbf{u}_k &\asymp \frac{2}{\pi} \int_0^\infty \mathbf{u}_k^\top \Sigma \left(\Sigma + \kappa(u^2)I \right)^{-1} \mathbf{u}_k du \\ 2811 \\ 2812 &= \frac{2}{\pi} \int_0^\infty \frac{\lambda_k}{\lambda_k + \kappa(u^2)} du \\ 2813 \\ 2814 &< \frac{2}{\pi} \int_0^\infty \frac{\lambda_k}{\lambda_k + u^2} du \\ 2815 \\ 2816 &= \lambda_k^{1/2} \\ 2817 \end{aligned}$$

2818 **C.6 PROOF FOR EXPECTATION OF THE SAMPLING MAPPING (FULL VERSION, FINITE σ_T)**

2820 Next, we consider the finite σ_T case, which involves two matrix half and their equivalence. To prove
 2821 this, we proceed in two steps 1) use integral identity to represent matrix of this form $A^{1/2}(A+zI)^{-1/2}$,
 2822 2) apply one point deterministic equivalence.

2823 **Proposition 7.** *Integral representation, for self-adjoint, positive semi definite matrix $A \succeq 0$,*

$$\begin{aligned} 2825 \quad A^{1/2}(A+zI)^{-1/2} &= \frac{4}{\pi^2} \int_0^\infty \int_0^\infty A(A+u^2I)^{-1}(A+(z+v^2)I)^{-1} dudv \\ 2826 \\ 2827 &= \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \frac{A(A+(z+u^2)I)^{-1} - A(A+v^2I)^{-1}}{v^2 - u^2 - z} dudv \\ 2828 \\ 2829 \end{aligned}$$

2830 *Proof.* Next, we can study matrix of this form,

$$2831 \quad A^{1/2}(A+zI)^{-1/2}$$

2832 using the integral representation above twice, we have

$$\begin{aligned} 2834 \quad A^{1/2}(A+zI)^{-1/2} &= AA^{-1/2}(A+zI)^{-1/2} \\ 2835 \\ 2836 &= \frac{1}{\pi^2} A \int_0^\infty (A+sI)^{-1} s^{-1/2} ds \int_0^\infty (A+(z+t)I)^{-1} t^{-1/2} dt \\ 2837 \\ 2838 &= \frac{1}{\pi^2} \int_0^\infty \int_0^\infty A(A+sI)^{-1}(A+(z+t)I)^{-1} t^{-1/2} s^{-1/2} ds dt \\ 2839 \\ 2840 &= \frac{4}{\pi^2} \int_0^\infty \int_0^\infty A(A+u^2I)^{-1}(A+(z+v^2)I)^{-1} dudv \\ 2841 \\ 2842 \\ 2843 \end{aligned}$$

2844 To deal with this *product of resolvent*, we can turn it into *difference of resolvent* via Lemma 3,

$$2846 \quad (A+sI)^{-1}(A+tI)^{-1} = \frac{1}{(s-t)} \left((A+tI)^{-1} - (A+sI)^{-1} \right)$$

2848 Now using the identity, we have

$$\begin{aligned} 2849 \quad A^{1/2}(A+zI)^{-1/2} &= \frac{1}{\pi^2} \int_0^\infty \int_0^\infty A(A+sI)^{-1}(A+(z+t)I)^{-1} t^{-1/2} s^{-1/2} ds dt \\ 2850 \\ 2851 &= \frac{1}{\pi^2} \int_0^\infty \int_0^\infty \frac{A(A+(z+t)I)^{-1} - A(A+sI)^{-1}}{s - z - t} t^{-1/2} s^{-1/2} ds dt \\ 2852 \\ 2853 \\ 2854 \\ 2855 \end{aligned}$$

2856 Putting it together,

$$\begin{aligned} 2857 \quad A^{1/2}(A+zI)^{-1/2} &= \frac{1}{\pi^2} \int_0^\infty \int_0^\infty \frac{A(A+(z+t)I)^{-1} - A(A+sI)^{-1}}{s - t - z} t^{-1/2} s^{-1/2} ds dt \\ 2858 \\ 2859 &= \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \frac{A(A+(z+u^2)I)^{-1} - A(A+v^2I)^{-1}}{v^2 - u^2 - z} dudv \\ 2860 \\ 2861 \end{aligned}$$

□

2862 Next we are ready to use the one-point deterministic equivalence.
2863
2864 **Proposition 8.** *For sample covariance matrix $\hat{\Sigma}$, the following expression has deterministic equiva-*
2865 *lent to the double integral of population covariance,*
2866
$$\hat{\Sigma}^{1/2}(\hat{\Sigma} + \sigma^2 I)^{-1/2} \asymp \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \frac{\kappa(\sigma^2 + u^2) - \kappa(v^2)}{(\sigma^2 + u^2) - v^2} \Sigma(\Sigma + \kappa(\sigma^2 + u^2)I)^{-1}(\Sigma + \kappa(v^2)I)^{-1} dudv$$
2867
2868

2869 *Proof.* Using Proposition 7, set $A \rightarrow \hat{\Sigma}$ we can apply the deterministic equivalence \ref for resolvents
2870

$$\begin{aligned} 2871 \hat{\Sigma}^{1/2}(\hat{\Sigma} + \sigma^2 I)^{-1/2} &= \hat{\Sigma} \hat{\Sigma}^{-1/2}(\hat{\Sigma} + \sigma^2 I)^{-1/2} \\ 2872 &= \frac{1}{\pi^2} \int_0^\infty \int_0^\infty \hat{\Sigma}(\hat{\Sigma} + sI)^{-1}(\hat{\Sigma} + (\sigma^2 + t)I)^{-1} t^{-1/2} s^{-1/2} ds dt \\ 2873 &= \frac{1}{\pi^2} \int_0^\infty \int_0^\infty \frac{\hat{\Sigma}(\hat{\Sigma} + (\sigma^2 + t)I)^{-1} - \hat{\Sigma}(\hat{\Sigma} + sI)^{-1}}{s - t - \sigma^2} t^{-1/2} s^{-1/2} ds dt \\ 2874 &\asymp \frac{1}{\pi^2} \int_0^\infty \int_0^\infty \frac{\Sigma(\Sigma + \kappa(\sigma^2 + t)I)^{-1} - \Sigma(\Sigma + \kappa(s)I)^{-1}}{s - t - \sigma^2} t^{-1/2} s^{-1/2} ds dt \\ 2875 & \\ 2876 & \\ 2877 & \\ 2878 & \\ 2879 & \end{aligned}$$

2880 Note there is no pole in this double integral, i.e. when $s = t + \sigma^2$, $\Sigma(\Sigma + \kappa(\sigma^2 + t)I)^{-1} =$
2881 $\Sigma(\Sigma + \kappa(s)I)^{-1}$, thus both numerator and denominator vanish, and the limit is well defined as a
2882 derivative!
2883

$$\begin{aligned} 2884 RHS &= \frac{1}{\pi^2} \int_0^\infty \int_0^\infty \frac{\Sigma(\Sigma + \kappa(\sigma^2 + t)I)^{-1} - \Sigma(\Sigma + \kappa(s)I)^{-1}}{s - t - \sigma^2} t^{-1/2} s^{-1/2} ds dt \\ 2885 &= \frac{1}{\pi^2} \int_0^\infty \int_0^\infty \frac{(\kappa(s) - \kappa(\sigma^2 + t))\Sigma(\Sigma + \kappa(\sigma^2 + t)I)^{-1}(\Sigma + \kappa(s)I)^{-1}}{s - t - \sigma^2} t^{-1/2} s^{-1/2} ds dt \\ 2886 &= \frac{1}{\pi^2} \int_0^\infty \int_0^\infty \frac{\kappa(s) - \kappa(\sigma^2 + t)}{s - (\sigma^2 + t)} \Sigma(\Sigma + \kappa(\sigma^2 + t)I)^{-1}(\Sigma + \kappa(s)I)^{-1} t^{-1/2} s^{-1/2} ds dt \\ 2887 & \\ 2888 & \\ 2889 & \\ 2890 & \end{aligned}$$

2891 This formulation shows that there is no real poles.
2892

2893 We can remove the singularity at 0 via $t \rightarrow u^2, s \rightarrow v^2$ change of variables
2894

$$RHS = \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \frac{\kappa(\sigma^2 + u^2) - \kappa(v^2)}{(\sigma^2 + u^2) - v^2} \Sigma(\Sigma + \kappa(\sigma^2 + u^2)I)^{-1}(\Sigma + \kappa(v^2)I)^{-1} dudv$$
2895
2896

2897 Thus we obtain the desired equivalence,
2898

$$\hat{\Sigma}^{1/2}(\hat{\Sigma} + \sigma^2 I)^{-1/2} \asymp \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \frac{\kappa(\sigma^2 + u^2) - \kappa(v^2)}{(\sigma^2 + u^2) - v^2} \Sigma(\Sigma + \kappa(\sigma^2 + u^2)I)^{-1}(\Sigma + \kappa(v^2)I)^{-1} dudv$$
2899
2900
2901

2902 Note that the coefficient $\frac{\kappa(\sigma^2 + u^2) - \kappa(v^2)}{(\sigma^2 + u^2) - v^2}$ has nice behavior when $(\sigma^2 + u^2) - v^2 \rightarrow 0$, i.e. it becomes
2903 a derivative of κ (Lemma 3). So there is no singularity in the integrand. \square
2904

2905 **Interpretation** We can compare it to sampling mapping with the population covariance, i.e. infinite
2906 data limit. Using Prop 7, setting $A \rightarrow \Sigma$, the double integral representation of the denoiser mapping
2907 reads,
2908

$$\begin{aligned} 2909 \Sigma^{1/2}(\Sigma + \sigma^2 I)^{-1/2} &= \frac{1}{\pi^2} \int_0^\infty \int_0^\infty \Sigma(\Sigma + (\sigma^2 + t)I)^{-1}(\Sigma + sI)^{-1} t^{-1/2} s^{-1/2} ds dt \\ 2910 &= \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \Sigma(\Sigma + (\sigma^2 + u^2)I)^{-1}(\Sigma + v^2 I)^{-1} dudv \\ 2911 & \\ 2912 & \\ 2913 & \end{aligned}$$

2914 Indeed, since $\kappa(\sigma^2 + u^2) > (\sigma^2 + u^2)$ and $\kappa(v^2) > v^2$, this creates a larger shrinkage, especially at
2915 small eigen dimensions.

2916 C.7 PROOF FOR FLUCTUATION OF THE SAMPLING MAPPING (APPROXIMATE VERSION,
 2917 INFINITE σ_T , PROPOSITION 4)

2920 Now let's consider the variance of the generated outcome with the infinite σ_T approximation, ignoring
 2921 estimation error in μ ,

$$\begin{aligned} \mathbf{x}_{\sigma_0} &= \mu + \hat{\Sigma}^{1/2}(\hat{\Sigma} + \sigma_T^2 I)^{-1/2}(\mathbf{x}_{\sigma_T} - \mu) \\ &\approx \mu + \hat{\Sigma}^{1/2}\left(\frac{\mathbf{x}_{\sigma_T} - \mu}{\sigma_T}\right) \\ &= \mu + \hat{\Sigma}^{1/2}\bar{\mathbf{x}} \end{aligned}$$

2933 So the variance coming from estimation of the covariance , let $\bar{\mathbf{x}} := \frac{\mathbf{x}_{\sigma_T} - \mu}{\sigma_T}$ i.e. normalized deviation
 2934 from center.

2936 **Proposition 9** (Main result, variance of generated sample under empirical data covariance.).

$$\begin{aligned} \text{Var}_{\hat{\Sigma}}[\mathbf{v}^\top \hat{\Sigma}^{1/2} \bar{\mathbf{x}}] &\asymp \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \frac{\kappa \kappa'}{n - \text{df}_2(\kappa, \kappa')} [\mathbf{v}^\top \Sigma(\Sigma + \kappa I)^{-1}(\Sigma + \kappa' I)^{-1} \mathbf{v}] \right. \\ &\quad \left. \times [\bar{\mathbf{x}}^\top \Sigma(\Sigma + \kappa I)^{-1}(\Sigma + \kappa' I)^{-1} \bar{\mathbf{x}}] \right\} du dv \end{aligned}$$

2947 where $\kappa := \kappa(u^2)$, $\kappa' := \kappa(v^2)$ are variables needing to be integrated over.

2958 *Proof.* Represent variance by moments,

$$\begin{aligned} \text{Var}_{\hat{\Sigma}}[\mathbf{v}^\top \hat{\Sigma}^{1/2} \bar{\mathbf{x}}] &= \mathbb{E}_{\hat{\Sigma}}[(\mathbf{v}^\top \hat{\Sigma}^{1/2} \bar{\mathbf{x}})^2] - \mathbb{E}_{\hat{\Sigma}}[\mathbf{v}^\top \hat{\Sigma}^{1/2} \bar{\mathbf{x}}]^2 \\ &= \mathbb{E}_{\hat{\Sigma}}[\mathbf{v}^\top \hat{\Sigma}^{1/2} \bar{\mathbf{x}} \bar{\mathbf{x}}^\top \hat{\Sigma}^{1/2} \mathbf{v}] - \mathbb{E}_{\hat{\Sigma}}[\mathbf{v}^\top \hat{\Sigma}^{1/2} \bar{\mathbf{x}}] \mathbb{E}_{\hat{\Sigma}}[\bar{\mathbf{x}}^\top \hat{\Sigma}^{1/2} \mathbf{v}] \quad \text{using Eq. 23} \\ &= \mathbb{E}_{\hat{\Sigma}}\left\{ \mathbf{v}^\top \left[\frac{2}{\pi} \int_0^\infty \hat{\Sigma}(\hat{\Sigma} + u^2 I)^{-1} du \right] \bar{\mathbf{x}} \bar{\mathbf{x}}^\top \left[\frac{2}{\pi} \int_0^\infty \hat{\Sigma}(\hat{\Sigma} + v^2 I)^{-1} dv \right] \mathbf{v} \right\} \\ &\quad - \mathbb{E}_{\hat{\Sigma}}\left\{ \mathbf{v}^\top \left[\frac{2}{\pi} \int_0^\infty \hat{\Sigma}(\hat{\Sigma} + u^2 I)^{-1} du \right] \bar{\mathbf{x}} \right\} \mathbb{E}_{\hat{\Sigma}}\left\{ \bar{\mathbf{x}}^\top \left[\frac{2}{\pi} \int_0^\infty \hat{\Sigma}(\hat{\Sigma} + v^2 I)^{-1} dv \right] \mathbf{v} \right\} \end{aligned}$$

2970 Using the integral representation and exchanging the integral with expectation,
 2971

$$\begin{aligned}
 2972 \text{RHS} &= \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \mathbb{E}_{\hat{\Sigma}} \left\{ \mathbf{v}^\top \hat{\Sigma} (\hat{\Sigma} + u^2 I)^{-1} \bar{\mathbf{x}} \bar{\mathbf{x}}^\top \hat{\Sigma} (\hat{\Sigma} + v^2 I)^{-1} \mathbf{v} \right\} \right. \\
 2973 &\quad \left. - \mathbb{E}_{\hat{\Sigma}} \left[\mathbf{v}^\top \hat{\Sigma} (\hat{\Sigma} + u^2 I)^{-1} \bar{\mathbf{x}} \right] \mathbb{E}_{\hat{\Sigma}} \left[\bar{\mathbf{x}}^\top \hat{\Sigma} (\hat{\Sigma} + v^2 I)^{-1} \mathbf{v} \right] \right\} dudv \\
 2974 &\qquad\qquad\qquad \text{integral representation of matrix half} \\
 2975 &\asymp \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \mathbb{E}_{\hat{\Sigma}} \left\{ \mathbf{v}^\top \hat{\Sigma} (\hat{\Sigma} + u^2 I)^{-1} \bar{\mathbf{x}} \bar{\mathbf{x}}^\top \hat{\Sigma} (\hat{\Sigma} + v^2 I)^{-1} \mathbf{v} \right\} \right. \\
 2976 &\quad \left. - \left[\mathbf{v}^\top \Sigma (\Sigma + \kappa(u^2)I)^{-1} \bar{\mathbf{x}} \right] \left[\bar{\mathbf{x}}^\top \Sigma (\Sigma + \kappa(v^2)I)^{-1} \mathbf{v} \right] \right\} dudv \\
 2977 &\qquad\qquad\qquad \text{using one point equivalence} \\
 2978 &\asymp \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \text{Tr} [\mathbf{v} \mathbf{v}^\top T_\Sigma \bar{\mathbf{x}} \bar{\mathbf{x}}^\top T_\Sigma'] + \frac{\kappa \kappa'}{n - \text{df}_2(\kappa, \kappa')} \text{Tr} [\mathbf{v} \mathbf{v}^\top G_\Sigma \Sigma G_\Sigma'] \text{Tr} [\bar{\mathbf{x}} \bar{\mathbf{x}}^\top G_\Sigma' \Sigma G_\Sigma] \right. \\
 2979 &\quad \left. - \left[\mathbf{v}^\top \Sigma (\Sigma + \kappa(u^2)I)^{-1} \bar{\mathbf{x}} \right] \left[\bar{\mathbf{x}}^\top \Sigma (\Sigma + \kappa(v^2)I)^{-1} \mathbf{v} \right] \right\} dudv \\
 2980 &\qquad\qquad\qquad \text{using two point equivalence} \\
 2981 &= \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \frac{\kappa \kappa'}{n - \text{df}_2(\kappa, \kappa')} \text{Tr} [\mathbf{v} \mathbf{v}^\top G_\Sigma \Sigma G_\Sigma'] \text{Tr} [\bar{\mathbf{x}} \bar{\mathbf{x}}^\top G_\Sigma' \Sigma G_\Sigma] \right\} dudv \\
 2982 &\qquad\qquad\qquad \text{first trace cancels out.} \\
 2983 &= \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \frac{\kappa \kappa'}{n - \text{df}_2(\kappa, \kappa')} [\mathbf{v}^\top G_\Sigma \Sigma G_\Sigma' \mathbf{v}] [\bar{\mathbf{x}}^\top G_\Sigma' \Sigma G_\Sigma \bar{\mathbf{x}}] \right\} dudv \\
 2984 &= \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \frac{\kappa \kappa'}{n - \text{df}_2(\kappa, \kappa')} [\mathbf{v}^\top \Sigma (\Sigma + \kappa(v^2)I)^{-1} (\Sigma + \kappa(u^2)I)^{-1} \mathbf{v}] \right. \\
 2985 &\quad \left. \times [\bar{\mathbf{x}}^\top \Sigma (\Sigma + \kappa(v^2)I)^{-1} (\Sigma + \kappa(u^2)I)^{-1} \bar{\mathbf{x}}] \right\} dudv \\
 2986 &= \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \frac{\kappa \kappa'}{n - \text{df}_2(\kappa, \kappa')} [\mathbf{v}^\top \Sigma (\Sigma + \kappa I)^{-1} (\Sigma + \kappa' I)^{-1} \mathbf{v}] \right. \\
 2987 &\quad \left. \times [\bar{\mathbf{x}}^\top \Sigma (\Sigma + \kappa I)^{-1} (\Sigma + \kappa' I)^{-1} \bar{\mathbf{x}}] \right\} dudv
 \end{aligned}$$

3002 Thus we arrive at our result
 3003

$$\begin{aligned}
 3004 \text{Var}[\mathbf{v}^\top \hat{\Sigma}^{1/2} \bar{\mathbf{x}}] &\asymp \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \frac{\kappa \kappa'}{n - \text{df}_2(\kappa, \kappa')} [\mathbf{v}^\top \Sigma (\Sigma + \kappa I)^{-1} (\Sigma + \kappa' I)^{-1} \mathbf{v}] \right. \\
 3005 &\quad \left. \times [\bar{\mathbf{x}}^\top \Sigma (\Sigma + \kappa I)^{-1} (\Sigma + \kappa' I)^{-1} \bar{\mathbf{x}}] \right\} dudv
 \end{aligned}$$

3009 \square

3012 C.7.1 INTERPRETATION

3014 **Anisotropy: effect of probe vector** If we marginalize over the $\bar{\mathbf{x}}$, assuming $\bar{\mathbf{x}} \sim \mathcal{N}(0, I)$ from
 3015 white noise, and consider only the effect of probe direction \mathbf{v} ,

$$\begin{aligned}
 3017 \mathbb{E}_{\bar{\mathbf{x}}} \text{Var}[\mathbf{v}^\top \hat{\Sigma}^{1/2} \bar{\mathbf{x}}] &\asymp \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \frac{\kappa \kappa'}{n - \text{df}_2(\kappa, \kappa')} [\mathbf{v}^\top \Sigma (\Sigma + \kappa I)^{-1} (\Sigma + \kappa' I)^{-1} \mathbf{v}] \right. \\
 3018 &\quad \left. \times \mathbb{E}_{\bar{\mathbf{x}}} [\bar{\mathbf{x}}^\top \Sigma (\Sigma + \kappa I)^{-1} (\Sigma + \kappa' I)^{-1} \bar{\mathbf{x}}] \right\} dudv \\
 3019 &\asymp \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \frac{\kappa \kappa' \text{Tr}[\Sigma (\Sigma + \kappa I)^{-1} (\Sigma + \kappa' I)^{-1}]}{n - \text{df}_2(\kappa, \kappa')} [\mathbf{v}^\top \Sigma (\Sigma + \kappa I)^{-1} (\Sigma + \kappa' I)^{-1} \mathbf{v}] \right\} dudv
 \end{aligned}$$

3024

3025

$$\begin{aligned}
3026 \quad \text{Tr}[\Sigma(\Sigma + \kappa I)^{-1}(\Sigma + \kappa' I)^{-1}] &= \frac{1}{\kappa} \text{Tr}[(\Sigma + \kappa I - \Sigma)\Sigma(\Sigma + \kappa I)^{-1}(\Sigma + \kappa' I)^{-1}] \\
3027 &= \frac{1}{\kappa} \text{Tr}[\Sigma(\Sigma + \kappa' I)^{-1}] - \frac{1}{\kappa} \text{Tr}[\Sigma^2(\Sigma + \kappa I)^{-1}(\Sigma + \kappa' I)^{-1}] \\
3028 &= \frac{1}{\kappa} \text{df}_1(\kappa') - \frac{1}{\kappa} \text{df}_2(\kappa, \kappa') \\
3029 &= \frac{1}{\kappa'} \text{df}_1(\kappa) - \frac{1}{\kappa'} \text{df}_2(\kappa, \kappa')
\end{aligned}$$

3030

Using this identity

3031

3032

$$3033 \quad \mathbb{E}_{\bar{x}} \text{Var}[\mathbf{v}^\top \hat{\Sigma}^{1/2} \bar{x}] \asymp \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \frac{\kappa'(\text{df}_1(\kappa') - \text{df}_2(\kappa, \kappa'))}{n - \text{df}_2(\kappa, \kappa')} [\mathbf{v}^\top \Sigma(\Sigma + \kappa I)^{-1}(\Sigma + \kappa' I)^{-1} \mathbf{v}] \right\} dudv$$

3034

3035

3036

Let's set the direction as the eigenvector \mathbf{u}_k , and the corresponding eigenvalue λ_k

3037

3038

3039

$$\mathbb{E}_{\bar{x}} \text{Var}[\mathbf{u}_k^\top \hat{\Sigma}^{1/2} \bar{x}] \asymp \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \frac{\kappa'(\text{df}_1(\kappa') - \text{df}_2(\kappa, \kappa'))}{n - \text{df}_2(\kappa, \kappa')} \frac{\lambda_k}{(\lambda_k + \kappa)(\lambda_k + \kappa')} \right\} dudv$$

3040

3041

Inhomogeneity: effect of initial noise Since the variance is symmetric in \bar{x} and \mathbf{v} , so we can marginalize over \mathbf{v} while keeping the \bar{x} dependency. Note that we assume \mathbf{v} is unit norm, so summation over \mathbf{u}_k eigenvectors (instead of expectation) is equivalent to trace.

3042

3043

3044

3045

$$\sum_k \text{Var}[\mathbf{u}_k^\top \hat{\Sigma}^{1/2} \bar{x}] = \text{Tr Var}[\hat{\Sigma}^{1/2} \bar{x}] \asymp \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \frac{\kappa'(\text{df}_1(\kappa') - \text{df}_2(\kappa, \kappa'))}{n - \text{df}_2(\kappa, \kappa')} [\bar{x}^\top \Sigma(\Sigma + \kappa I)^{-1}(\Sigma + \kappa' I)^{-1} \bar{x}] \right\} dudv \quad (25)$$

3046

3047

Scaling: effect of sample number and scaling Finally marginalizing over both factors, we have the overall scaling.

3048

3049

3050

3051

3052

$$\mathbb{E}_{\bar{x}} \sum_k \text{Var}[\mathbf{u}_k^\top \hat{\Sigma}^{1/2} \bar{x}] = \mathbb{E}_{\bar{x}} \text{Tr Var}[\hat{\Sigma}^{1/2} \bar{x}] \asymp \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \left\{ \frac{(\text{df}_1(\kappa') - \text{df}_2(\kappa, \kappa'))(\text{df}_1(\kappa) - \text{df}_2(\kappa, \kappa'))}{n - \text{df}_2(\kappa, \kappa')} \right\} dudv \quad (26)$$

3053

3054

3055

3056

3057

3058

3059

3060

3061

3062

3063

3064

3065

3066

3067

3068

3069

3070

3071

3072

3073

3074

3075

3076

3077

3078 **D EXPERIMENTAL DETAILS**3079 **D.1 NUMERICAL METHODS**

3080 **Numerical evaluation of renormalized Ridge $\kappa(z)$.** We computed $\kappa(z)$ as the solution to the
 3081 self-consistent Silverstein equation

$$3085 \quad \kappa(z) - z = \gamma \sum_{k=1}^p w_k \frac{\kappa(z) \lambda_k}{\kappa(z) + \lambda_k}, \quad (27)$$

3086 where $\{\lambda_k\}$ are the eigenvalues of Σ and $\{w_k\}$ are their normalized weights. For scalar z , we solved
 3087 this nonlinear equation using Newton’s method with analytical derivative

$$3088 \quad \kappa'(z) = 1 - \gamma \sum_{k=1}^p w_k \frac{\lambda_k^2}{(\kappa(z) + \lambda_k)^2},$$

3089 falling back to a robust root-finder for purely real inputs. For a sequence of z values along a path,
 3090 we used an “analytic continuation” procedure in which the solution at the previous z served as the
 3091 initial guess for the next, ensuring branch continuity and numerical stability, particularly for small z .
 3092 Further, we generally start the path from z with high norm and solve with continuation back to small
 3093 z . A caching mechanism stored previously computed (z, κ) pairs, with nearest-neighbor retrieval for
 3094 initial guesses, further accelerating repeated evaluations. This approach yields accurate and smooth
 3095 $\kappa(z)$ profiles suitable for downstream quadrature-based integration.

3102 **Numerical evaluation of the integral over deterministic equivalence** The analytical results in
 3103 Eqs. 9,10 involving integral to infinity are not trivial to evaluate. To avoid truncation error, we used
 3104 the following scheme by translating the integration onto a finite domain.

3105 We approximated the double integral

$$3106 \quad \frac{4}{\pi^2} \int_0^\infty \int_0^\infty \frac{\kappa \kappa' \operatorname{Tr}[\Sigma(\Sigma + \kappa I)^{-1}(\Sigma + \kappa' I)^{-1}]}{n - df_2(\kappa, \kappa')} [\mathbf{v}^\top \Sigma(\Sigma + \kappa I)^{-1}(\Sigma + \kappa' I)^{-1} \mathbf{v}] du dv \quad (28)$$

3111 using a Gauss–Legendre quadrature scheme combined with the tangent mapping $u = \tan \theta$ to
 3112 transform the semi-infinite domain $[0, \infty)$ to a finite interval $[0, \pi/2)$.

3113 We first generated n_{nodes} Gauss–Legendre nodes θ_i and weights w_i on $[0, \pi/2]$, then applied the
 3114 transformation $u = \tan \theta$ with Jacobian $J(\theta) = 1/\cos^2 \theta$ to obtain quadrature points on $[0, \infty)$. This
 3115 was performed independently for u and v , and their 2D tensor product provided the integration grid.

3116 The κ values were computed at each u^2 and v^2 using a numerically stable, vectorized evaluation
 3117 of the spectral mapping function $\kappa(z)$ derived from the eigenspectrum of Σ . The integrand was
 3118 then assembled by evaluating the trace term, the scalar bilinear form $\mathbf{v}^\top (\cdot) \mathbf{v}$, and the denominator
 3119 $n - df_2(\kappa, \kappa')$ on the full 2D grid. Quadrature weights and Jacobians were applied multiplicatively,
 3120 and the sum over all grid points yielded the numerical approximation to the integral.

3121 Similar quadrature is used for the single integral equivalence Eq.9, where we integrate over 1d grid.

3123 This approach yields high accuracy while avoiding explicit truncation of the infinite domain, as the
 3124 nonlinear mapping concentrates quadrature nodes where the integrand varies most rapidly.

3126 **D.2 LINEAR DENOISER EXPERIMENTS**

3128 To cross validate against our theory and numerical scheme, we performed extensive validation via
 3129 linear denoiser set up using empirical denoiser.

3131 We compute the empirical covariance of a dataset and then used the following functions implementing
 3132 the linear one-step denoiser and the full sampling map (Wiener filter).

```

3132
3133     def dnoised_X(x, Xmean, sample_cov, sigma2,):
3134         # single step denoiser
3135         return x + sigma2 * (Xmean - x) @ torch.inverse(sample_cov +
3136             torch.eye(sample_cov.shape[0], device=x.device) * sigma2)
3137
3138     def wiener_gen_X(x, Xmean, wiener_matrix, sigmaT,):
3139         if x.dim() == 1:
3140             # Single vector case
3141             return Xmean + wiener_matrix @ (x * sigmaT - Xmean)
3142         else:
3143             # Batched vector case - x should be shape (batch_size, ndim)
3144             return Xmean[None, :] + (x * sigmaT - Xmean[None, :]) @ wiener_matrix.T
3145
3146     def build_wiener_matrix(eigvals, eigvecs, sigmaT=80.0, sigma0=0.0, EPS=1E-16,
3147         clip=True):
3148         if clip:
3149             eigvals = torch.clamp(eigvals, min=EPS)
3150         scaling = ((eigvals + sigma0**2) / (eigvals + sigmaT**2)).sqrt()
3151         return eigvecs @ torch.diag(scaling) @ eigvecs.T
3152
3153
3154
3155
3156
3157
3158
3159
3160
3161
3162
3163
3164
3165
3166
3167
3168
3169
3170
3171
3172
3173
3174
3175
3176
3177
3178
3179
3180
3181
3182
3183
3184
3185

```

We keep the $\sigma_0 = 0$ in theory, in reality, it's usually set to a small positive number e.g. 0.002. So in a few cases, we tested this and reported the results in appendix. Generally, it acts as a floor for generated variance, thus remedy the overshrinking effect.

We found when the dataset size is not enough, e.g. rank deficient $\hat{\Sigma}$, the eigendecomposition is not stable, sometimes generating negative eigenvalues, which affects the matrix square root operation in Wiener matrix. Even if we clip them, there is often numerical artifacts at small eigenspaces. One solution is, we use higher precision float64 number to yield similar results with the theory.

D.3 DEEP NEURAL NETWORK EXPERIMENTS

We used following preconditioning scheme inspired by [Karras et al. \(2022\)](#), for all our architectures for comparison.

```

3166
3167     class EDMPrecondWrapper(nn.Module):
3168         def __init__(self, model, sigma_data=0.5, sigma_min=0.002, sigma_max=80,
3169             rho=7.0):
3170             super().__init__()
3171             self.model = model
3172             self.sigma_data = sigma_data
3173             self.sigma_min = sigma_min
3174             self.sigma_max = sigma_max
3175             self.rho = rho
3176
3177         def forward(self, X, sigma, cond=None, ):
3178             sigma[sigma == 0] = self.sigma_min
3179             ## edm preconditioning for input and output
3180             ## https://github.com/NVlabs/edm/blob/main/training/networks.py#L632
3181             # unsqueeze sigma to have same dimension as X (which may have 2-4 dim)
3182             sigma_vec = sigma.view([-1, ] + [1, ] * (X.ndim - 1))
3183             c_skip = self.sigma_data ** 2 / (sigma_vec ** 2 + self.sigma_data ** 2)
3184             c_out = sigma_vec * self.sigma_data / (sigma_vec ** 2 + self.sigma_data ** 2).sqrt()
3185             c_in = 1 / (self.sigma_data ** 2 + sigma_vec ** 2).sqrt()
3186             c_noise = sigma.log() / 4
3187             model_out = self.model(c_in * X, c_noise, cond=cond)
3188             return c_skip * X + c_out * model_out
3189
3190
3191
3192
3193
3194
3195
3196
3197
3198
3199
3200
3201
3202
3203
3204
3205
3206
3207
3208
3209
3210
3211
3212
3213
3214
3215
3216
3217
3218
3219
3220
3221
3222
3223
3224
3225
3226
3227
3228
3229
3230
3231
3232
3233
3234
3235
3236
3237
3238
3239
3240
3241
3242
3243
3244
3245
3246
3247
3248
3249
3250
3251
3252
3253
3254
3255
3256
3257
3258
3259
3260
3261
3262
3263
3264
3265
3266
3267
3268
3269
3270
3271
3272
3273
3274
3275
3276
3277
3278
3279
3280
3281
3282
3283
3284
3285
3286
3287
3288
3289
3290
3291
3292
3293
3294
3295
3296
3297
3298
3299
3300
3301
3302
3303
3304
3305
3306
3307
3308
3309
3310
3311
3312
3313
3314
3315
3316
3317
3318
3319
3320
3321
3322
3323
3324
3325
3326
3327
3328
3329
3330
3331
3332
3333
3334
3335
3336
3337
3338
3339
3340
3341
3342
3343
3344
3345
3346
3347
3348
3349
3350
3351
3352
3353
3354
3355
3356
3357
3358
3359
3360
3361
3362
3363
3364
3365
3366
3367
3368
3369
3370
3371
3372
3373
3374
3375
3376
3377
3378
3379
3380
3381
3382
3383
3384
3385
3386
3387
3388
3389
3390
3391
3392
3393
3394
3395
3396
3397
3398
3399
3400
3401
3402
3403
3404
3405
3406
3407
3408
3409
3410
3411
3412
3413
3414
3415
3416
3417
3418
3419
3420
3421
3422
3423
3424
3425
3426
3427
3428
3429
3430
3431
3432
3433
3434
3435
3436
3437
3438
3439
3440
3441
3442
3443
3444
3445
3446
3447
3448
3449
3450
3451
3452
3453
3454
3455
3456
3457
3458
3459
3460
3461
3462
3463
3464
3465
3466
3467
3468
3469
3470
3471
3472
3473
3474
3475
3476
3477
3478
3479
3480
3481
3482
3483
3484
3485
3486
3487
3488
3489
3490
3491
3492
3493
3494
3495
3496
3497
3498
3499
3500
3501
3502
3503
3504
3505
3506
3507
3508
3509
3510
3511
3512
3513
3514
3515
3516
3517
3518
3519
3520
3521
3522
3523
3524
3525
3526
3527
3528
3529
3530
3531
3532
3533
3534
3535
3536
3537
3538
3539
3540
3541
3542
3543
3544
3545
3546
3547
3548
3549
3550
3551
3552
3553
3554
3555
3556
3557
3558
3559
3560
3561
3562
3563
3564
3565
3566
3567
3568
3569
3570
3571
3572
3573
3574
3575
3576
3577
3578
3579
3580
3581
3582
3583
3584
3585
3586
3587
3588
3589
3590
3591
3592
3593
3594
3595
3596
3597
3598
3599
3600
3601
3602
3603
3604
3605
3606
3607
3608
3609
3610
3611
3612
3613
3614
3615
3616
3617
3618
3619
3620
3621
3622
3623
3624
3625
3626
3627
3628
3629
3630
3631
3632
3633
3634
3635
3636
3637
3638
3639
3640
3641
3642
3643
3644
3645
3646
3647
3648
3649
3650
3651
3652
3653
3654
3655
3656
3657
3658
3659
3660
3661
3662
3663
3664
3665
3666
3667
3668
3669
3670
3671
3672
3673
3674
3675
3676
3677
3678
3679
3680
3681
3682
3683
3684
3685
3686
3687
3688
3689
3690
3691
3692
3693
3694
3695
3696
3697
3698
3699
3700
3701
3702
3703
3704
3705
3706
3707
3708
3709
3710
3711
3712
3713
3714
3715
3716
3717
3718
3719
3720
3721
3722
3723
3724
3725
3726
3727
3728
3729
3730
3731
3732
3733
3734
3735
3736
3737
3738
3739
3740
3741
3742
3743
3744
3745
3746
3747
3748
3749
3750
3751
3752
3753
3754
3755
3756
3757
3758
3759
3760
3761
3762
3763
3764
3765
3766
3767
3768
3769
3770
3771
3772
3773
3774
3775
3776
3777
3778
3779
3780
3781
3782
3783
3784
3785
3786
3787
3788
3789
3790
3791
3792
3793
3794
3795
3796
3797
3798
3799
3800
3801
3802
3803
3804
3805
3806
3807
3808
3809
3810
3811
3812
3813
3814
3815
3816
3817
3818
3819
3820
3821
3822
3823
3824
3825
3826
3827
3828
3829
3830
3831
3832
3833
3834
3835
3836
3837
3838
3839
3840
3841
3842
3843
3844
3845
3846
3847
3848
3849
3850
3851
3852
3853
3854
3855
3856
3857
3858
3859
3860
3861
3862
3863
3864
3865
3866
3867
3868
3869
3870
3871
3872
3873
3874
3875
3876
3877
3878
3879
3880
3881
3882
3883
3884
3885
3886
3887
3888
3889
3890
3891
3892
3893
3894
3895
3896
3897
3898
3899
3900
3901
3902
3903
3904
3905
3906
3907
3908
3909
3910
3911
3912
3913
3914
3915
3916
3917
3918
3919
3920
3921
3922
3923
3924
3925
3926
3927
3928
3929
3930
3931
3932
3933
3934
3935
3936
3937
3938
3939
3940
3941
3942
3943
3944
3945
3946
3947
3948
3949
3950
3951
3952
3953
3954
3955
3956
3957
3958
3959
3960
3961
3962
3963
3964
3965
3966
3967
3968
3969
3970
3971
3972
3973
3974
3975
3976
3977
3978
3979
3980
3981
3982
3983
3984
3985
3986
3987
3988
3989
3990
3991
3992
3993
3994
3995
3996
3997
3998
3999
3999

```

EDM Loss Function We employ the loss function \mathcal{L}_{EDM} introduced in the Elucidated Diffusion Model (EDM) paper [Karras et al. \(2022\)](#), which is one specific weighting scheme for training diffusion models.

For each data point $\mathbf{x} \in \mathbb{R}^d$, the loss is computed as follows. The noise level for each data point is sampled from a log-normal distribution with hyperparameters P_{mean} and P_{std} (e.g., $P_{\text{mean}} = -1.2$ and $P_{\text{std}} = 1.2$). Specifically, the noise level σ is sampled via

$$\sigma = \exp(P_{\text{mean}} + P_{\text{std}} \epsilon), \quad \epsilon \sim \mathcal{N}(0, 1).$$

The weighting function per noise scale is defined as:

$$w(\sigma) = \frac{\sigma^2 + \sigma_{\text{data}}^2}{(\sigma \sigma_{\text{data}})^2},$$

with hyperparameter σ_{data} (e.g., $\sigma_{\text{data}} = 0.5$). The noisy input \mathbf{y} is created by the following,

$$\mathbf{y} = \mathbf{x} + \sigma \mathbf{n}, \quad \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d),$$

Let $D_{\theta}(\mathbf{y}, \sigma, \text{labels})$ denote the output of the denoising network when given the noisy input \mathbf{y} , the noise level σ , and optional conditioning labels. The EDM loss per data point can be computed as:

$$\mathcal{L}(\mathbf{x}) = w(\sigma) \|D_{\theta}(\mathbf{x} + \sigma \mathbf{n}, \sigma, \text{labels}) - \mathbf{x}\|^2.$$

Taking expectation over the data points and noise scales, the overall loss reads

$$\mathcal{L}_{\text{EDM}} = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} \mathbb{E}_{\mathbf{n} \sim \mathcal{N}(0, \mathbf{I}_d)} \mathbb{E}_{\sigma} \left[w(\sigma) \|D_{\theta}(\mathbf{x} + \sigma \mathbf{n}, \sigma, \text{labels}) - \mathbf{x}\|^2 \right] \quad (29)$$

Hyperparameter Settings: DiT All experiments use DiT backbones with consistent architectural and optimization settings unless otherwise specified. Key hyperparameters:

- **Model architecture:** patch size 2 or 4 (used once for FFHQ64, discarded for worse performance), hidden size 384, depth 6 layers, 6 attention heads, MLP ratio 4.
- **Datasets:** FFHQ-32, AFHQ-32, CIFAR-32, and FFHQ-64; subsampled at varying sizes (300, 1k, 3k, 10k, 30k) with two non-overlapping splits per size.
- **Training objective:** Denoising Score Matching (DSM) under EDM parametrization.
- **Training schedule:** 50000 steps with batch size 256, Adam optimizer with learning rate 1×10^{-4} .
- **Evaluation:** fixed-noise seed, sampling with 35 steps with Heun sampler; evaluation sample size 1000, batch size 512.

Hyperparameter Settings: UNet All CNN-UNet experiments follow consistent architectural and optimization settings unless noted. Key hyperparameters:

- **Model architecture:** UNet with base channels 128 ; channel multipliers $\{1, 2, 2, 2\}$; self-attention at resolution 8.
- **Datasets:** FFHQ-32, AFHQ-32, CIFAR-32, and FFHQ-64; subsampled at varying sizes (300, 1k, 3k, 10k, 30k) with two non-overlapping splits per size.
- **Training objective:** Denoising Score Matching (DSM) under EDM parametrization.
- **Training schedule:** 50000 steps, batch size 256, Adam with learning rate 1×10^{-4} .
- **Evaluation:** fixed-noise seed, sampling with 35 steps with Heun sampler; evaluation sample size 1000, batch size 512.

Computation Cost All experiments were conducted on NVIDIA A100 or H100 GPUs. Training DiT and CNN models on 32×32 resolution datasets typically required 5–8 hours to complete. In contrast, DiT models trained on FFHQ64 were substantially more expensive, taking approximately 24 hours per run.

3240 **E USAGE OF LLMS**
32413242 We used LLMs in three ways. First, as a research assistant, to look up tools related to deterministic
3243 equivalence and to point us toward integral identities for fractional matrix powers, which we then
3244 verified and derived independently. Second, as a coding agent to help us generate plotting and analysis
3245 code for our results. Third, as a writing aid, for polishing technical text and providing feedback on
3246 clarity and presentation of the whole paper.
3247
3248
3249
3250
3251
3252
3253
3254
3255
3256
3257
3258
3259
3260
3261
3262
3263
3264
3265
3266
3267
3268
3269
3270
3271
3272
3273
3274
3275
3276
3277
3278
3279
3280
3281
3282
3283
3284
3285
3286
3287
3288
3289
3290
3291
3292
3293