

Understanding Generalization in Diffusion Models via Probability Flow Distance

Huijie Zhang, Zijian Huang, Siyi Chen

Jinfan Zhou, Zekai Zhang, Peng Wang, Qing Qu

{HUIJIEZH, NAMELESS, SIYICH, ZJF, ZZEKAI, PENGWA, QINGQU}@UMICH.EDU

University of Michigan, Ann Arbor, United States

Abstract

Diffusion models have emerged as a powerful class of generative models, capable of producing high-quality samples that generalize beyond the training data. However, evaluating this generalization remains challenging: theoretical metrics are often impractical for high-dimensional data, while no practical metrics rigorously measure generalization. In this work, we bridge this gap by introducing probability flow distance (PFD), a theoretically grounded and computationally efficient metric to measure distributional generalization. Specifically, PFD quantifies the distance between distributions by comparing their noise-to-data mappings induced by the probability flow ODE. Moreover, by using PFD under a teacher-student evaluation protocol, we empirically uncover several key generalization behaviors in diffusion models, including: (1) scaling behavior from memorization to generalization, (2) early learning and double descent training dynamics, and (3) bias-variance decomposition. Beyond these insights, our work lays a foundation for future empirical and theoretical studies on generalization in diffusion models.

1. Introduction

In recent years, diffusion models and their variants have revolutionized generative AI, achieving state-of-the-art performance in tasks ranging from image synthesis [21, 59] to molecular design [1, 31]. These models, including score-based generative models [66] and flow matching techniques [47, 50], learn the underlying data distribution through forward and reverse processes that gradually inject and remove noise. Their success raises a fundamental question: how can we rigorously evaluate the generalization ability of these models? A good evaluation framework is essential not only for deepening our understanding of the underlying mechanisms of generative modeling but also for providing principled guidance in designing more effective architectures, training strategies, and benchmarking methods.

However, existing metrics for evaluating the generalizability of diffusion models face significant limitations. Empirically, common metrics like Fréchet inception distance (FID) [28], Inception Score (IS) [64] focus on generation quality, but they cannot distinguish between memorization and generalization, as both can yield high-quality outputs. While metrics such as Kullback-Leibler (KL) divergence [16, 44, 54], total variation (TV) [14, 42, 43, 71], and Wasserstein distance [11, 12, 23, 24] are theoretically appealing, they are often computationally expensive and thus impractical for diffusion models. Furthermore, since the true data distribution is typically unknown, it makes such comparisons inherently challenging. In summary, existing metrics are neither accurate nor efficient for evaluating diffusion models in practice, highlighting the need for a generalization metric that is both theoretically grounded and practically tractable.

Our Contribution. In this work, we introduce a systematic framework for evaluating the generalizability of diffusion models through a novel metric, the probability flow distance (PFD). This metric quantifies distributional differences by leveraging the backward probability flow ODE (PF-ODE) [66], which is widely used in the sampling process of diffusion models. Unlike practical metrics such as FID, PFD provides a theoretically grounded measure of distance between distributions, offering a more reliable assessment of generalization. Compared to theoretical metrics like the Wasserstein distance, PFD is computationally efficient by leveraging the benign properties of PF-ODE. Moreover, under a distillation-based setting, we use this metric to study generalization error by comparing the PFD between the student and teacher models. Our analysis reveals several intriguing generalization phenomena that offer new insights into the learning behavior of diffusion models, including: (1) Scaling behavior from memorization to generalization. (2) Early learning and double descent of generalization in learning dynamics. (3) Bias and variance trade-off of the generalization error. Due to space limitations, we defer this to Appendix B.

2. Measuring Distribution Distance via Probability Flow Distance

In this section, we propose a new metric called probability flow distance (PFD), which is designed to quantify the distance between two arbitrary probability distributions. The design of PFD is motivated by the PF-ODE, which we first review in Section 2.1. We then formally define PFD in Section 2.2 and present its empirical estimation with theoretical guarantees in Section 2.3.

2.1. A Mapping from Noise to Target Distribution Spaces Induced by PF-ODE

In general, PF-ODE is a class of ordinary differential equations (ODE) that aim to reverse a forward process, where Gaussian noise is progressively added to samples drawn from an underlying distribution, denoted as p_{data} ¹. The forward process and the PF-ODE can be described as follows:

- *Forward process.* Given a sample $\mathbf{x}_0 \stackrel{i.i.d.}{\sim} p_{\text{data}}(\mathbf{x})$, the forward process progressively corrupts it by adding Gaussian noise. This process can be characterized by the stochastic differential equation (SDE) $d\mathbf{x}_t = f(t)\mathbf{x}_t dt + g(t)d\mathbf{w}_t$, where $t \in [0, T]$ is the time index, $\{\mathbf{w}_t\}_{t \in [0, T]}$ is a standard Wiener process, and $f(t), g(t) : \mathbb{R}_+ \rightarrow \mathbb{R}$ are drift and diffusion function functions that control the noise schedule. In this work, we adopt the noise schedule proposed by elucidated diffusion models (EDM) [36], where $f(t) = 0$ and $g(t) = \sqrt{2t}$. Substituting this into the SDE and integrating both sides, we obtain

$$\mathbf{x}_t = \mathbf{x}_0 + \int_0^t \sqrt{2\tau} d\mathbf{w}_\tau. \quad (1)$$

For ease of exposition, we use $p_t(\mathbf{x}_t)$ to denote the distribution of the noisy image \mathbf{x}_t for each $t \in [0, T]$. In particular, it is worth noting that $p_0(\mathbf{x}) = p_{\text{data}}(\mathbf{x})$ and $p_T(\mathbf{x}) \rightarrow \mathcal{N}(\mathbf{0}, T^2 \mathbf{I}_n)$ as $T \rightarrow +\infty$.

- *Probability flow ODE.* According to [66], the PF-ODE can transform a noise sample \mathbf{x}_T back into a clean data sample \mathbf{x}_0 . Specifically, under EDM noise scheduler, the PF-ODE admits the following form:

$$d\mathbf{x}_t = -t \nabla \log p_t(\mathbf{x}_t) dt, \quad (2)$$

1. This paper primarily focuses on image distribution.

where $\nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t)$ (or simply $\nabla \log p_t(\mathbf{x}_t)$) denotes the *score function* of the distribution $p_t(\mathbf{x}_t)$ at time $t \in [0, T]$. According to [66], the backward PF-ODE (2) and the forward SDE (1) have the same distribution at each timestep t . In practice, since the score function $\log p_t(\mathbf{x}_t)$ is unknown, in diffusion models we approximate it using a neural network $\mathbf{s}_\theta(\mathbf{x}_t, t)$ and employ a numerical solver to generate samples from Equation (2). Additional details are provided in Appendix A.3.

Benign properties of PF-ODE. The backward PF-ODE introduces a mapping $\Phi_{p_{\text{data}}}$ from \mathbf{x}_T to \mathbf{x}_0 . By taking the integral on both sides of (2) from T to 0, the mapping $\Phi_{p_{\text{data}}}$ can be defined as:

$$\Phi_{p_{\text{data}}}(\mathbf{x}_T) := \mathbf{x}_T - \int_T^0 t \nabla \log p_t(\mathbf{x}_t) dt. \quad (3)$$

Previous work [66] demonstrates that $\Phi_{p_{\text{data}}}(\mathbf{x}_T) \sim p_{\text{data}}(\mathbf{x})$ when $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, T^2 \mathbf{I}_n)$ as $T \rightarrow +\infty$. This implies that when the underlying distribution p_{data} is known, the score function $\nabla \log p_t(\mathbf{x}_t)$ becomes explicitly available, and the backward PF-ODE induces a deterministic mapping from the Gaussian distribution to p_{data} .

2.2. Definition of Probability Flow Distance

Definition 1 (Probability flow distance (PFD)) For any two given distributions p and q of the same dimension, we define their distribution distance as

$$\text{PFD}(p, q) := \left(\mathbb{E}_{\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, T^2 \mathbf{I})} \left[\|\Psi \circ \Phi_p(\mathbf{x}_T) - \Psi \circ \Phi_q(\mathbf{x}_T)\|_2^2 \right] \right)^{1/2}. \quad (4)$$

Here, Φ_p and Φ_q denote the mappings between the noise and image spaces for distributions p and q , respectively, as defined in (3), and $\Psi(\cdot)$ represents an image descriptor.

Intuitively, PFD measures the distance between two distributions p and q by comparing their respective noise-to-image mappings $\Phi_p(\cdot)$ and $\Phi_q(\cdot)$ starting from the same Gaussian noise input \mathbf{x}_T .

Moreover, the comparison is conducted in a transformed feature space defined by an image descriptor $\Psi(\cdot)$, which is typically implemented using a pre-trained neural network to effectively capture perceptual differences.

Under Definition 1, we show that PFD satisfies the axioms of a metric (Definition 2.15 in [61]).

Theorem 1 For any two distributions p and q , the PFD satisfies the following properties:

- (Positivity) $\text{PFD}(p, q) > 0$ for any $p \neq q$.
- (Identity Property) $\text{PFD}(p, q) = 0$ if and only if $p = q$.
- (Symmetry) $\text{PFD}(p, q) = \text{PFD}(q, p)$.
- (Triangle Inequality) $\text{PFD}(p, q) \leq \text{PFD}(p, p') + \text{PFD}(p', q)$ for all p' .

We defer the proof to Appendix C. Note that Theorem 1 establishes the theoretical validity of PFD as a metric for measuring distance between any two probability distributions.

2.3. Empirical Estimation of PFD

In practice, the expectation in (4) is intractable due to the complexity of the underlying distributions. Thus, we approximate the PFD using finite samples:

$$\widehat{\text{PFD}}(p, q) = \left(\frac{1}{M} \sum_{i=1}^M \left\| \Phi_p(\mathbf{x}_T^{(i)}) - \Phi_q(\mathbf{x}_T^{(i)}) \right\|_2^2 \right)^{1/2}. \quad (5)$$

Here, $\widehat{\text{PFD}}(p, q)$ is the empirical version of $\text{PFD}(p, q)$ computed over M independent samples $\{\mathbf{x}_T^{(i)}\}_{i=1}^M \stackrel{i.i.d.}{\sim} \mathcal{N}(0, T^2 \mathbf{I}_n)$ with $T \rightarrow \infty$.

Specifically, our finite-sample approximation relies on two key assumptions: (i) the score functions are smooth at all timesteps, and (ii) the score functions of two distributions remain uniformly close within a bounded region of the input space, which can be described as follows.

Assumption 1 *Let p and q be two distributions with the same dimension, where we assume:*

(i) *There exists a constant $L > 0$ such that for all $\mathbf{x}_1, \mathbf{x}_2$ and $t \in [0, T]$, it holds that*

$$\|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}_1) - \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_2)\|_2 \leq L \|\mathbf{x}_1 - \mathbf{x}_2\|_2, \quad (6)$$

and similarly for q_t .

(ii) *For all $t \in [0, T]$, there exists a constant $\epsilon > 0$ such that*

$$\|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \nabla_{\mathbf{x}} \log q_t(\mathbf{x})\|_2 \leq \epsilon. \quad (7)$$

The Lipschitz continuity of the score function is a common assumption widely adopted in the theoretical analysis of score functions in diffusion models [10, 12, 13, 15, 41, 76]. More recently, this property has been rigorously established under the assumption that the data distribution is a mixture of Gaussians [46]. The uniform closeness assumption holds when p, q follow Assumption 1 (i) and have support on a compact domain, which is often the case for image distributions. Under Assumption 1, the concentration of the empirical estimate $\widehat{\text{PFD}}(p, q)$ to $\text{PFD}(p, q)$ can be characterized as follows.

Theorem 2 *Suppose we are given two distributions p and q that satisfy the L -Lipschitz condition and are uniformly close in Assumption 1. Let $\widehat{\text{PFD}}(p, q)$ denote the empirical estimate of $\text{PFD}(p, q)$, computed as the average over M independent samples, as introduced in (5). Then, for any $\gamma > 0$, the empirical estimate satisfies the following bound:*

$$|\widehat{\text{PFD}}(p, q) - \text{PFD}(p, q)| \leq \gamma, \quad \text{whenever } M \geq \frac{\kappa^4(L, \epsilon)}{2\gamma^4} \log \frac{2}{\eta}, \quad (8)$$

with probability at least $1 - \eta$. Here, $\kappa(L, \epsilon) := \exp\left(\frac{LT^2}{2\xi}\right) \xi + \frac{\epsilon}{L} \left(\exp\left(\frac{LT^2}{2\xi}\right) - 1\right)$ is a constant, with a numerical constant $\xi > 0$ and a finite timestep T_ξ depending only on ξ .

We defer the proof to Appendix C. Given the score functions of both distributions are smooth and uniformly close, our result in Theorem 2 guarantees that $\text{PFD}(p, q)$ can be approximated to arbitrary precision by its empirical estimate $\widehat{\text{PFD}}(p, q)$ with high probability, given a finite number of samples.

3. Quantifying Generalization Error of Diffusion Models

In this section, we leverage the PFD metric in Section 2 to rigorously define and evaluate the generalization error of diffusion models. Specifically, this metric enables us to distinguish between memorization and generalization behaviors for diffusion models, as well as analyze the transition from memorization to generation (MtoG).

Definition 2 (Generalization and Memorization Errors) *Consider a diffusion model s_θ trained on a finite dataset $\mathcal{D} = \{\mathbf{y}^{(i)}\}_{i=1}^N$, where each sample $\mathbf{y}^{(i)}$ is drawn i.i.d. from the underlying distribution $p_{\text{data}}(\mathbf{x})$. Denote the learned distribution induced by a diffusion model s_θ as $p_\theta(\mathbf{x})$. Using the PFD metric, we can formally define the generalization and memorization errors as follows:*

$$\mathcal{E}_{\text{gen}}(\theta) := \text{PFD}(p_\theta, p_{\text{data}}), \quad \mathcal{E}_{\text{mem}}(\theta) := \text{PFD}(p_\theta, p_{\text{emp}}), \quad (9)$$

where the empirical distribution is given by $p_{\text{emp}}(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N \delta(\mathbf{x} - \mathbf{y}^{(i)})$, with $\delta(\cdot)$ denoting the Dirac delta function.

Here, given access to $p_{\text{emp}}(\mathbf{x})$, the memorization error $\mathcal{E}_{\text{mem}}(\theta)$ can be exactly computed (see Appendix E). We further show that $\mathcal{E}_{\text{mem}}(\theta)$ coincides with metrics introduced in [73, 75]. However, since the underlying distribution $p_{\text{data}}(\mathbf{x})$ is typically unknown in practice, we introduce a teacher–student evaluation protocol to analyze the generalization error of diffusion models.

Evaluation protocol of generalization. To study the generalization behavior of diffusion models (see Appendix B), we adopt a teacher–student framework. We treat a large-scale pretrained diffusion model $s_{\theta_t}(\mathbf{x}_t)$ with parameters θ_t as the teacher, inducing a distribution p_{θ_t} , which we take as a proxy for the true data distribution, i.e., $p_{\text{data}} = p_{\theta_t}$. We then train a student model s_θ using samples drawn from p_{θ_t} , and evaluate its generalization by comparing p_θ to p_{θ_t} using the generalization errors defined in Definition 2.

In our experiments for the rest of the paper, both teacher and student models adopt the U-Net architecture [60]. The teacher model s_{θ_t} is trained on the CIFAR-10 dataset [39] with a fixed model architecture (UNet-10 introduced in Appendix D.1). The student model s_θ is trained on samples generated by the teacher, with the number of training samples varying from $N = 2^6$ to $N = 2^{16}$, using the same training hyperparameters but different model sizes. For evaluating the generalization error in (9), we compute the PFD between the teacher and student models using $M = 10^4$ samples drawn from shared initial noise, as defined in (5). Similar for the memorization error, we compute the PFD between the student model and the empirical distribution of the training data. Additional details for the evaluation protocol and ablation studies with different teacher models are provided in Appendix D.2 and Appendix F.4, respectively. Due to the space limitation, we defer the generalization behaviors in diffusion models experiment in Appendix B.

4. Conclusion

In this work, we introduced Probability Flow Distance, a theoretically grounded and computationally tractable metric for evaluating the generalization ability of diffusion models. Using a teacher–student evaluation protocol, we empirically reveal several key generalization behaviors in learning diffusion models, including: (1) the scaling transition from memorization to generalization, (2) early learning and double descent training dynamics, and (3) a bias–variance trade-off of generalization error.

References

- [1] Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630(8016):493–500, 2024.
- [2] Sumukh K Aithal, Pratyush Maini, Zachary Lipton, and J Zico Kolter. Understanding hallucinations in diffusion models through mode interpolation. In *Advances in Neural Information Processing Systems*, volume 37, pages 134614–134644, 2024.
- [3] Fan Bao, Shen Nie, Kaiwen Xue, Yue Cao, Chongxuan Li, Hang Su, and Jun Zhu. All are worth words: A vit backbone for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [4] Peter L Bartlett, Philip M Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.
- [5] Peter L Bartlett, Andrea Montanari, and Alexander Rakhlin. Deep learning: a statistical viewpoint. *Acta numerica*, 30:87–201, 2021.
- [6] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- [7] Mikhail Belkin, Daniel Hsu, and Ji Xu. Two models of double descent for weak features. *SIAM Journal on Mathematics of Data Science*, 2(4):1167–1180, 2020.
- [8] Rajendra Bhatia. *Matrix analysis*, volume 169. Springer Science & Business Media, 2013.
- [9] Mikołaj Bińkowski, Dougal J. Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=r1lUozWCW>.
- [10] Adam Block, Youssef Mroueh, and Alexander Rakhlin. Generative modeling with denoising auto-encoders and langevin sampling. *arXiv preprint arXiv:2002.00107*, 2020.
- [11] Valentin De Bortoli. Convergence of denoising diffusion models under the manifold hypothesis. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL <https://openreview.net/forum?id=MhK5aXo3gB>. Expert Certification.
- [12] Hongrui Chen, Holden Lee, and Jianfeng Lu. Improved analysis of score-based generative modeling: User-friendly bounds under minimal smoothness assumptions. In *International Conference on Machine Learning*, pages 4735–4763. PMLR, 2023.
- [13] Minshuo Chen, Kaixuan Huang, Tuo Zhao, and Mengdi Wang. Score approximation, estimation and distribution recovery of diffusion models on low-dimensional data. In *Proceedings of the 40th International Conference on Machine Learning*, 2023.

- [14] Sitan Chen, Sinho Chewi, Holden Lee, Yuanzhi Li, Jianfeng Lu, and Adil Salim. The probability flow ODE is provably fast. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=KD6MFeWSAd>.
- [15] Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=zyLVMgsZ0U_.
- [16] Sitan Chen, Giannis Daras, and Alex Dimakis. Restoration-degradation beyond linear diffusions: A non-asymptotic analysis for ddim-type samplers. In *International Conference on Machine Learning*, pages 4462–4484. PMLR, 2023.
- [17] Zhengdao Chen. On the interpolation effect of score smoothing. *arXiv preprint arXiv:2502.19499*, 2025.
- [18] Min Jin Chong and David Forsyth. Effectively unbiased fid and inception score and where to find them. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6070–6079, 2020.
- [19] Earl A Coddington and Norman Levinson. *Theory of ordinary differential equations*. McGraw-Hill New York, 1955.
- [20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- [21] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024.
- [22] Ruiqi Gao, Emiel Hoogeboom, Jonathan Heek, Valentin De Bortoli, Kevin Patrick Murphy, and Tim Salimans. Diffusion models and gaussian flow matching: Two sides of the same coin. In *The Fourth Blogpost Track at ICLR 2025*, 2025. URL <https://openreview.net/forum?id=C8Yyg9wy0s>.
- [23] Xuefeng Gao and Lingjiong Zhu. Convergence analysis for general probability flow ODEs of diffusion models in wasserstein distances. In *The 28th International Conference on Artificial Intelligence and Statistics*, 2025. URL <https://openreview.net/forum?id=EkO8rb3liX>.
- [24] Xuefeng Gao, Hoang M. Nguyen, and Lingjiong Zhu. Wasserstein convergence guarantees for a general class of score-based generative models. *Journal of Machine Learning Research*, 26(43):1–54, 2025. URL <http://jmlr.org/papers/v26/24-0902.html>.
- [25] Xiangming Gu, Chao Du, Tianyu Pang, Chongxuan Li, Min Lin, and Ye Wang. On memorization in diffusion models. *arXiv preprint arXiv:2310.02664*, 2023.

- [26] Trevor Hastie, Robert Tibshirani, Jerome Friedman, et al. The elements of statistical learning, 2009.
- [27] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. CLIPScore: a reference-free evaluation metric for image captioning. In *Empirical Methods in Natural Language Processing*, 2021.
- [28] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- [29] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851, 2020.
- [30] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *The Collected Works of Wassily Hoeffding*, pages 409–426, 1994.
- [31] Emiel Hoogeboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d. In *International Conference on Machine Learning*, pages 8867–8887. PMLR, 2022.
- [32] Sadeep Jayasumana, Srikumar Ramalingam, Andreas Veit, Daniel Glasner, Ayan Chakrabarti, and Sanjiv Kumar. Rethinking fid: Towards a better evaluation metric for image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9307–9315, 2024.
- [33] Zahra Kadkhodaie, Florentin Guth, Eero P Simoncelli, and Stéphane Mallat. Generalization in diffusion models arises from geometry-adaptive harmonic representations. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=ANvmVS2Yr0>.
- [34] Mason Kamb and Surya Ganguli. An analytic theory of creativity in convolutional diffusion models. *arXiv preprint arXiv:2412.20292*, 2024.
- [35] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- [36] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Advances in Neural Information Processing Systems*, volume 35, pages 26565–26577, 2022.
- [37] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. <https://github.com/NVlabs/edm/tree/main>, 2022.
- [38] Ron Kohavi, David H Wolpert, et al. Bias plus variance decomposition for zero-one loss functions. In *International Conference on Machine Learning*, volume 96, pages 275–283. Citeseer, 1996.

- [39] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [40] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- [41] Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence for score-based generative modeling with polynomial complexity. In *Advances in Neural Information Processing Systems*, volume 35, pages 22870–22882, 2022.
- [42] Gen Li, Yu Huang, Timofey Efimov, Yuting Wei, Yuejie Chi, and Yuxin Chen. Accelerating convergence of score-based diffusion models, provably. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.
- [43] Gen Li, Yuting Wei, Yuxin Chen, and Yuejie Chi. Towards non-asymptotic convergence for diffusion-based generative models. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=4VGEeER6W9>.
- [44] Puheng Li, Zhong Li, Huishuai Zhang, and Jiang Bian. On the generalization properties of diffusion models. In *Advances in Neural Information Processing Systems*, volume 36, pages 2097–2127, 2023.
- [45] Xiang Li, Yixiang Dai, and Qing Qu. Understanding generalizability of diffusion models requires rethinking the hidden gaussian structure. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=Sk2duBGvrK>.
- [46] Yingyu Liang, Zhenmei Shi, Zhao Song, and Yufa Zhou. Unraveling the smoothness properties of diffusion models: A gaussian mixture perspective. *arXiv preprint arXiv:2405.16418*, 2024.
- [47] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.
- [48] Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. In *Advances in Neural Information Processing Systems*, volume 33, pages 20331–20342, 2020.
- [49] Sheng Liu, Zhihui Zhu, Qing Qu, and Chong You. Robust training under label noise by over-parameterization. In *International Conference on Machine Learning*, pages 14153–14172. PMLR, 2022.
- [50] Xingchao Liu, Chengyue Gong, and qiang liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=XVjTT1nw5z>.
- [51] Calvin Luo. Understanding diffusion models: A unified perspective. *arXiv preprint arXiv:2208.11970*, 2022.

- [52] Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz Barak, and Ilya Sutskever. Deep double descent: Where bigger models and more data hurt. *Journal of Statistical Mechanics: Theory and Experiment*, 2021(12):124003, 2021.
- [53] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021.
- [54] Shen Nie, Hanzhong Allan Guo, Cheng Lu, Yuhao Zhou, Chenyu Zheng, and Chongxuan Li. The blessing of randomness: SDE beats ODE in general diffusion-based image editing. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=DesYwmUG00>.
- [55] Matthew Niedoba, Berend Zwartsenberg, Kevin Murphy, and Frank Wood. Towards a mechanistic explanation of diffusion model generalization. *arXiv preprint arXiv:2411.19339*, 2024.
- [56] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=a68SUt6zFt>. Featured Certification.
- [57] Ed Pizzi, Sreya Dutta Roy, Sugosh Nagavara Ravindra, Priya Goyal, and Matthijs Douze. A self-supervised descriptor for image copy detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14532–14542, 2022.
- [58] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PmLR, 2021.
- [59] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [60] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015.
- [61] Walter Rudin. Principles of mathematical analysis. 2021.
- [62] Mehdi SM Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

- [63] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=TIIdIXIpzhoI>.
- [64] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, volume 29, 2016.
- [65] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=St1gjarCHLP>.
- [66] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=PXTIG12RRHS>.
- [67] George Stein, Jesse Cresswell, Rasa Hosseinzadeh, Yi Sui, Brendan Ross, Valentin Vilecroze, Zhaoyan Liu, Anthony L Caterini, Eric Taylor, and Gabriel Loaiza-Ganem. Exposing flaws of generative model evaluation metrics and their unfair treatment of diffusion models. In *Advances in Neural Information Processing Systems*, volume 36, 2024.
- [68] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016.
- [69] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9446–9454, 2018.
- [70] Binxu Wang and John J Vastola. The unreasonable effectiveness of gaussian score approximation for diffusion models and its applications. *arXiv preprint arXiv:2412.09726*, 2024.
- [71] Ruofeng Yang, Zhijie Wang, Bo Jiang, and Shuai Li. Leveraging drift to improve sample complexity of variance exploding diffusion models. In *Advances in Neural Information Processing Systems*, volume 37, pages 107662–107702, 2024.
- [72] Zitong Yang, Yaodong Yu, Chong You, Jacob Steinhardt, and Yi Ma. Rethinking bias-variance trade-off for generalization of neural networks. In *International Conference on Machine Learning*, pages 10767–10777. PMLR, 2020.
- [73] TaeHo Yoon, Joo Young Choi, Sehyun Kwon, and Ernest K Ryu. Diffusion probabilistic models generalize when they fail to memorize. In *ICML 2023 Workshop on Structured Probabilistic Inference & Generative Modeling*, 2023.
- [74] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Sy8gdB9xx>.

- [75] Huijie Zhang, Jinfan Zhou, Yifu Lu, Minzhe Guo, Peng Wang, Liyue Shen, and Qing Qu. The emergence of reproducibility and consistency in diffusion models. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 60558–60590. PMLR, 21–27 Jul 2024. URL <https://proceedings.mlr.press/v235/zhang24cn.html>.
- [76] Zhenyu Zhu, Francesco Locatello, and Volkan Cevher. Sample complexity bounds for score-matching: Causal discovery and generative modeling. In *Advances in Neural Information Processing Systems*, volume 36, pages 3325–3337, 2023.

Appendix A. Related Works

In this section, we briefly review related work on generalization metrics for diffusion models, discuss diffusion model generalizability, and cover the fundamentals of training diffusion models.

A.1. Generalization Metrics for Diffusion Models

Generalization metrics quantify the distance between the learned distribution and the underlying data distribution in diffusion models. To measure this distributional gap, theoretical works commonly employ metrics such as Kullback-Leibler (KL) divergence [16, 44, 54], total variation (TV) [14, 42, 43, 71], and Wasserstein distance [11, 12, 23, 24]. However, these metrics are practically inefficient for diffusion models. Practical metrics focus on various perspective, including negative log-likelihood (NLL) [66], image generation quality: Fréchet inception distance (FID) [28], inception score (IS) [64], FD_{dinov2} [67], maximum mean discrepancy (MMD) [9], CLIP maximum mean discrepancy (CMMD) [32]; alignment: `CLIPscore` [27], and `precision, recall` [40, 62]. However, these practical metrics are not explicitly designed to evaluate the generalizability of diffusion models. Thus, there is a need for a generalization metric that are both theoretical grounded and practically efficient for diffusion models. To address this gap, we propose PFD, a novel generalization metric that is theoretically proven to be a valid distributional distance and can be efficiently approximated by its empirical version using a polynomial number of samples. In practice, PFD requires fewer samples for estimation and is the only existing metric that explicitly quantifies generalization in diffusion models.

A.2. Diffusion Model Generalizability

Recent works have shown that diffusion models transition from memorization to generalization as the number of training samples increases [73, 75]. With sufficient data, models trained with different architectures, loss functions, and even disjoint datasets can reproduce each other’s outputs, indicating a strong convergence toward the underlying data distribution [33, 75]. To explain this strong generalization, [33] attribute it to the emergence of a geometric-adaptive harmonic basis, while others argue that generalization arises from interpolation across the data manifold [2, 17]. Theoretical insights by [44] provide generalization bounds using KL-divergence under simplified models. More recent efforts focus on characterizing the learned noise-to-image mapping for generalized diffusion models, either through Gaussian parameterizations [45, 70] or patch-wise optimal score functions [34, 55]. However, despite these theoretical analyses and qualitative insights, prior work lacks a quantitative framework for measuring generalizability. In this paper, we propose PFD, a metric that enables such quantitative evaluation. Using this measure, we uncover further insights into the generalization behavior of diffusion models, as discussed in Appendix B.

A.3. Training diffusion models

To enable sampling via the PF-ODE (2), we train a neural network $s_{\theta}(\mathbf{x}_t, t)$ to approximate the score function $\nabla \log p_t(\mathbf{x}_t)$ using denoising score matching loss [66]:

$$\min_{\theta} \ell(\theta) = \frac{1}{N} \sum_{i=1}^N \int_0^T \lambda_t \mathbb{E}_{\epsilon \sim \mathcal{N}(0, T^2 \mathbf{I}_n)} \left[\left\| s_{\theta}(\mathbf{x}^{(i)} + t\epsilon, t) + \epsilon/t \right\|_2^2 \right] dt, \quad (10)$$

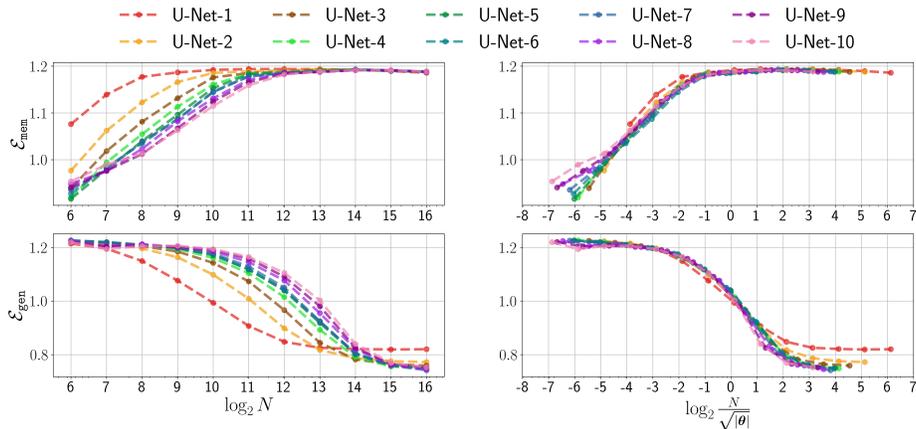


Figure 1: **Scaling behavior in the MtoG transition.** \mathcal{E}_{mem} and \mathcal{E}_{gen} plotted against $\log_2(N)$ for a range of U-Net architectures (U-Net-1 to U-Net-10). Right: the same metrics plotted against $\log_2(N/\sqrt{|\theta|})$, where $|\theta|$ is the number of model parameters.

λ_t denotes a scalar weight for the loss at t . Given the learned score function, the corresponding noise-to-image mapping is:

$$\Phi_{p_\theta}(x_T) = x_T - \int_T^0 t s_\theta(x_t, t) dt. \quad (11)$$

Although alternative training objectives exist, such as predicting noise x_T [29], clean image x_0 [36], rectified flow $x_T - x_0$ [50] or other linear combinations of x_0 and x_T [63], prior works [22, 51] have shown that it is still possible to recover an approximate score function $s_\theta(x_t, t)$ from these methods.

Appendix B. Measuring Key Generalization Behaviors in Diffusion Models

Based on the evaluation protocol in Section 3, this section reveals several key generalization behaviors in diffusion models: (i) MtoG scaling behaviors with model capacity and training size (Appendix B.1), (ii) early learning and double descent in learning dynamics (Appendix B.2), and (iii) bias-variance trade-off of generalization error (Appendix B.3).

B.1. Scaling Behaviors of the MtoG Transition

First, we investigate the scaling behavior of the MtoG transition with respect to both model capacity $|\theta|$ and training data size N , using the metrics \mathcal{E}_{gen} and \mathcal{E}_{mem} . We evaluate ten U-Net architectures on the CIFAR-10 dataset, with model sizes ranging from 0.9M to 55.7M parameters (U-Net-1 to U-Net-10). For each model, we compute \mathcal{E}_{mem} and \mathcal{E}_{gen} across varying training dataset sizes, following the evaluation protocol outlined in Section 3. We report our results in Figure 1 with additional experimental details provided in Appendix D.4, where we observe the following:

Finding I.1: Scaling training data N induces MtoG transition under fixed model capacity $|\theta|$.

As shown in Figure 1 (left), for a fixed model capacity $|\theta|$, our metrics reveal a clear transition from memorization to generalization as the number of training samples N increases. Notably,

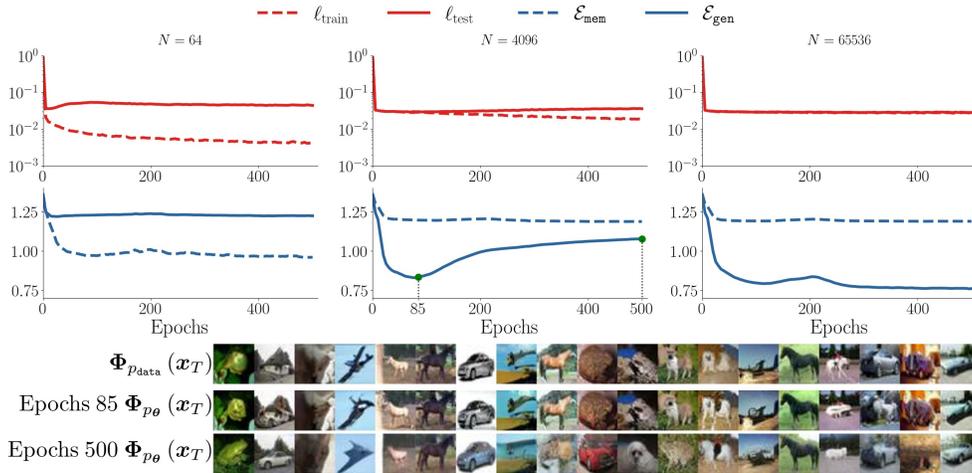


Figure 2: **Training dynamics of diffusion models in different regimes.** The top figure plots $\mathcal{E}_{\text{mem}}, \mathcal{E}_{\text{gen}}, \ell_{\text{train}}, \ell_{\text{test}}$ over training epochs for different dataset sizes: $N = 2^6$ (left), 2^{12} (middle), 2^{16} (right). The bottom figure visualizes the generation when $N = 2^{12}$. The top row shows samples from the underlying distribution $\Phi_{\text{data}}(\mathbf{x}_T)$, while the middle and bottom rows display outputs from the trained diffusion model $\Phi_{p_\theta}(\mathbf{x}_T)$ at epoch 85 and 500, respectively.

larger models transition more slowly to generalization, as their greater capacity allows them to memorize more of the training data. Compared to prior studies of this transition [73, 75], our results more accurately capture the underlying behavior by directly measuring the distributional distance between the learned and ground-truth distributions. In contrast, earlier approaches [73, 75] assess generalization based on the deviation of generated samples from the training data, which does not reliably reflect true generalization.

Finding I.2: MtoG transition governed consistently by the ratio $N/\sqrt{|\theta|}$. Moreover, in contrast to prior work that focuses solely on the effect of training sample size N , our results in Figure 1 (right) reveal a consistent scaling behavior when using our metric, governed by the ratio $N/\sqrt{|\theta|}$ between data size and model capacity. Remarkably, both \mathcal{E}_{gen} and \mathcal{E}_{mem} metrics exhibit near-identical MtoG transition curves across models of varying sizes when plotted against this ratio. As such, analogous to the empirical scaling laws observed in large language models [35], this predictable trend provides practical guidance for the development of diffusion models, particularly when scaling up model size, data, or compute to achieve optimal performance gains.

B.2. Early Learning and Double Descent in Learning Dynamics

Building on the findings in Appendix B.1, we further examine the generalization behavior across different training regimes. Under the evaluation protocol in Section 3, we analyze the learning dynamics of a U-Net model with fixed model capacity (UNet-10 introduced in Appendix D.1) trained with the number of data samples $N = 2^6, 2^{12},$ and 2^{16} , corresponding to the memorization, transition, and generalization regimes in Appendix B.1, respectively. The model is trained using stochastic gradient descent (SGD) for 500 epochs, during which we track $\mathcal{E}_{\text{mem}}, \mathcal{E}_{\text{gen}}, \ell_{\text{train}},$ and ℓ_{test} at each epoch. The results in Figure 2 reveal several notable generalization behaviors that align with phenomena previously observed in the training of overparameterized deep models [52, 74]:

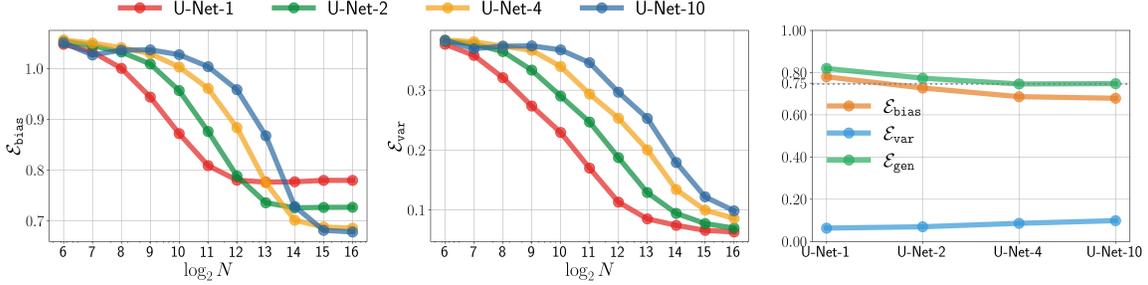


Figure 3: **Bias–Variance Trade-off.** Left figure shows $\mathcal{E}_{\text{bias}}$ and \mathcal{E}_{var} as functions of the number of training samples N for various network architectures. Right figure plots the generalization error \mathcal{E}_{gen} , bias $\mathcal{E}_{\text{bias}}$, and variance \mathcal{E}_{var} across different network architectures with a fixed training sample size of $N = 2^{16}$.

Finding II.1: Early learning behavior in memorization and transition regimes. As shown in Figure 2 (left & middle), in both the memorization ($N = 2^6$) and transition ($N = 2^{12}$) regimes, the generalization error initially decreases during training but reaches its minimum at an early epoch, after which it begins to increase again. This *early learning* (or early generalization) phenomenon becomes more salient as the training sample size increases from the memorization to the transition regime. As shown in the visualization at the bottom of Section 3, the model at Epoch 85 clearly exhibits generalization, whereas the model at Epoch 500 fails to generalize. This is also corroborated by the divergence of training loss ℓ_{train} and test loss ℓ_{test} at the top of the figure. It is worth mentioning that, although early learning behavior has been theoretically and visually demonstrated in previous works [44, 45], PFD is the first metric to provide empirical evidence of this phenomenon.

Finding II.2: Double descent of the generalization error in the generalization regime. In contrast, as shown in Figure 2 (right), training in the generalization regime ($N = 2^{16}$) reveals a clear instance of the *double descent* phenomenon [52] in the generalization error. Specifically, the error initially decreases, then increases during intermediate training epochs, and finally decreases again as training approaches convergence. Notably, this non-monotonic behavior is not captured by the standard training and test losses ℓ_{train} and ℓ_{test} , both of which decrease monotonically throughout training. This implies that extended training can improve generalization performance in the generalization regime.

Remarks. For both cases, it should be noted that these generalization phenomena observed through our metrics are not unique to diffusion models. Similar surprising behaviors have been previously reported in training overparameterized deep learning models with extensive theoretical investigations [4, 5, 7, 52, 74]. For example, the early learning phenomenon has been widely observed when training models with limited or noisy data, such as in deep image priors [69] and learning with label noise [48, 49]. Similarly, the double descent phenomenon has been reported in the training dynamics of overparameterized models [52]. These observations challenge the traditional view of generalization and highlight the critical role of inductive bias and training time in the learning process. Similarly, our findings imply that such factors should also be carefully considered when training diffusion models.

B.3. Bias-variance Trade-off of the Generalization Error

In statistical learning theory, bias-variance trade-off is a classical yet fundamental concept in supervised learning which helps us understand and analyze the sources of prediction error in the model [6, 26, 38, 72]. Specifically, bias–variance decomposition expresses the expected generalization error as the sum of two components: (i) the *bias term*, which quantifies the discrepancy between the expected model prediction and the true function—high bias indicates systematic error or underfitting; and (ii) the *variance term*, which measures the prediction variability of the model across different training sets—high variance reflects sensitivity to data fluctuations or overfitting.

However, in unsupervised learning settings such as diffusion models, the notion of generalization error was not well-defined prior to our work, in contrast to the well-established definitions in supervised learning. As a result, bias–variance decomposition in this context remains largely unexplored. In this work, we address this gap through the generalization error measure \mathcal{E}_{gen} (see Equation (9)), which admits a bias–variance decomposition analogous to that in the supervised setting, as we detail below.

Definition 3 (Bias-Variance Decomposition of \mathcal{E}_{gen}) *Based on the same setup as Definition 2, we can decompose \mathcal{E}_{gen} in Equation (9) as*

$$\mathbb{E}_{\mathcal{D}} [\mathcal{E}_{\text{gen}}^2 (p_{\theta(\mathcal{D})})] = \mathcal{E}_{\text{bias}}^2 + \mathcal{E}_{\text{var}} \quad (12)$$

where $p_{\theta(\mathcal{D})}$ denotes the distribution induced by a diffusion model θ (\mathcal{D}) trained on a given training dataset \mathcal{D} sampled from p_{data} . Specifically, the bias and variance terms are defined as:

$$\mathcal{E}_{\text{bias}} := \mathbb{E}_{x_T} (\|\Psi \circ \Phi_{p_{\text{data}}}(x_T) - \overline{\Psi \circ \Phi_{p_{\theta}}}(x_T)\|_2^2)^{1/2}, \mathcal{E}_{\text{var}} := \mathbb{E}_{\mathcal{D}} \mathbb{E}_{x_T} [\|\Psi \circ \Phi_{p_{\theta(\mathcal{D})}}(x_T) - \overline{\Psi \circ \Phi_{p_{\theta}}}(x_T)\|_2^2],$$

with $\overline{\Psi \circ \Phi_{p_{\theta}}}(\cdot) := \mathbb{E}_{\mathcal{D}} [\Psi \circ \Phi_{p_{\theta(\mathcal{D})}}(\cdot)]$.

Intuitively, our definitions of the bias term $\mathcal{E}_{\text{bias}}$ and the variance term \mathcal{E}_{var} are both well-justified: (i) $\mathcal{E}_{\text{bias}}$ quantifies the systematic error between the learned distribution p_{θ} and the ground-truth distribution p_{data} ; and (ii) \mathcal{E}_{var} captures the variability of model predictions across different training sets by measuring the distance between p_{θ} and the mean $\overline{p_{\theta}}$ which can be empirically estimated by averaging over multiple datasets \mathcal{D} sampled from p_{data} . Experimental results, following the protocol in Section 3, are shown in Figure 3, with detailed settings in Appendix D.6.

In Figure 3 (a), when diffusion models are trained in the generalization regime, the resulting generalization decomposition aligns with classical bias–variance theory from supervised learning: as model complexity increases, the bias term $\mathcal{E}_{\text{bias}}$ decreases while the variance term \mathcal{E}_{var} increases, resulting in a U-shaped generalization error curve. Additionally, Figure 3 (b) further illustrates the effect of the training sample size N and number of parameters $|\theta|$: increasing N reduces both $\mathcal{E}_{\text{bias}}$ and \mathcal{E}_{var} , thereby lowering the generalization error \mathcal{E}_{gen} , as expected; In contrast, increasing $|\theta|$ consistently increases \mathcal{E}_{var} , and its effect on $\mathcal{E}_{\text{bias}}$ depends on the size of N : it decreases $\mathcal{E}_{\text{bias}}$ when $N \geq 2^{15}$ but increases it when $N \leq 2^{11}$.

Appendix C. Proof in Section 2

Proof [Proof of Theorem 1] It is trivial to show $\text{PFD}(p, q) > 0$ for any $p \neq q$ and $\text{PFD}(p, q) = \text{PFD}(q, p)$, and thus we omit the proof.

- Proof of $p = q \Leftrightarrow \text{PFD}(p, q) = 0$:

– (\Rightarrow) If $p = q$, $\nabla \log p_t(\mathbf{x}_t) = \nabla \log q_t(\mathbf{x}_t)$, thus:

$$d\mathbf{x}_t = -t(\nabla \log p_t(\mathbf{x}_t) - \nabla \log q_t(\mathbf{x}_t)) dt = 0 \quad (13)$$

Thus, $\Phi_p(\mathbf{x}_T) - \Phi_q(\mathbf{x}_T)$ is the solution of the ODE function Equation (13) with initial $\mathbf{x}_T = \mathbf{0}$. Thus $\Phi_p(\mathbf{x}_T) - \Phi_q(\mathbf{x}_T) = \mathbf{0}$ for all \mathbf{x}_T . Thus $\text{PFD}(p, q) = 0$

– (\Leftarrow) If $\text{PFD}(p, q) = 0$ and Φ_p, Φ_q are continuous function w.r.t \mathbf{x}_T , then we have $\Phi_p(\mathbf{x}_T) = \Phi_q(\mathbf{x}_T)$ for all \mathbf{x}_T . If $\mathbf{x}_0 = \Phi(\mathbf{x}_T)$, from the transformation of probability identities, we have:

$$p(\mathbf{x}_0) = \frac{\partial}{\partial[\mathbf{x}_0]_1} \cdots \frac{\partial}{\partial[\mathbf{x}_0]_n} \int_{\{\epsilon | \Phi(\epsilon) \leq \mathbf{x}_0\}} p_{\mathcal{N}}(\epsilon) d^n \epsilon, \quad (14)$$

where $[\mathbf{x}_0]_i$ denotes the i -th element of \mathbf{x}_0 , $\mathbf{f}(\epsilon) \leq \mathbf{x}_0$ denotes the element wise less or equal. $p_{\mathcal{N}}(\cdot)$ is the probability density function (PDF) of Gaussian distribution $\mathcal{N}(\mathbf{0}, T^2 \mathbf{I}_n)$. Thus, for all \mathbf{x}_0 we have:

$$\begin{aligned} p(\mathbf{x}_0) - q(\mathbf{x}_0) &= \frac{\partial}{\partial[\mathbf{x}_0]_1} \cdots \frac{\partial}{\partial[\mathbf{x}_0]_n} \int_{\{\epsilon | \Phi_p(\epsilon) \leq \mathbf{x}_0\}} p_{\mathcal{N}}(\epsilon) d^n \epsilon \\ &\quad - \frac{\partial}{\partial[\mathbf{x}_0]_1} \cdots \frac{\partial}{\partial[\mathbf{x}_0]_n} \int_{\{\epsilon | \Phi_q(\epsilon) \leq \mathbf{x}_0\}} p_{\mathcal{N}}(\epsilon) d^n \epsilon, \\ &= \frac{\partial}{\partial[\mathbf{x}_0]_1} \cdots \frac{\partial}{\partial[\mathbf{x}_0]_n} \int_{\{\epsilon | \Phi_p(\epsilon) \leq \mathbf{x}_0\}} p_{\mathcal{N}}(\epsilon) d^n \epsilon \\ &\quad - \frac{\partial}{\partial[\mathbf{x}_0]_1} \cdots \frac{\partial}{\partial[\mathbf{x}_0]_n} \int_{\{\epsilon | \Phi_p(\epsilon) \leq \mathbf{x}_0\}} p_{\mathcal{N}}(\epsilon) d^n \epsilon, \\ &= 0, \end{aligned} \quad (15)$$

so $p = q$.

- Proof of $\text{PFD}(p, q) \leq \text{PFD}(p, p') + \text{PFD}(p', q)$:

$$\begin{aligned} &\text{PFD}(p, q) \\ &= \left(\mathbb{E}_{\mathbf{x}_T \sim \mathcal{N}(0, T^2 \mathbf{I})} \left[\|\Phi_p(\mathbf{x}_T) - \Phi_q(\mathbf{x}_T)\|_2^2 \right] \right)^{1/2} \\ &\leq \left(\mathbb{E}_{\mathbf{x}_T \sim \mathcal{N}(0, T^2 \mathbf{I})} \left[(\|\Phi_p(\mathbf{x}_T) - \Phi_{p'}(\mathbf{x}_T)\|_2 + \|\Phi_{p'}(\mathbf{x}_T) - \Phi_q(\mathbf{x}_T)\|_2)^2 \right] \right)^{1/2} \\ &\leq \left(\mathbb{E}_{\mathbf{x}_T \sim \mathcal{N}(0, T^2 \mathbf{I})} \left[\|\Phi_p(\mathbf{x}_T) - \Phi_q(\mathbf{x}_T)\|_2^2 \right] \right)^{1/2} \\ &\quad + \left(\mathbb{E}_{\mathbf{x}_T \sim \mathcal{N}(0, T^2 \mathbf{I})} \left[\|\Phi_p(\mathbf{x}_T) - \Phi_{p'}(\mathbf{x}_T)\|_2^2 \right] \right)^{1/2} \\ &= \text{PFD}(p, p') + \text{PFD}(p', q) \end{aligned} \quad (16)$$

■

Lemma 1 Under Assumption 1, for all $\mathbf{x}_T \in \mathcal{N}(\mathbf{0}, T^2 \mathbf{I}_n)$, as $T \rightarrow \infty$, we have:

$$\|\Phi_p(\mathbf{x}_T) - \Phi_q(\mathbf{x}_T)\|_2 \leq \exp\left(\frac{LT_\xi^2}{2}\right) \xi + \frac{\epsilon}{L} \left(\exp\left(\frac{LT_\xi^2}{2}\right) - 1 \right), \quad (17)$$

where ξ is a numerical constant and a finite timestep T_ξ depending only on ξ .

Proof [Proof of Theorem 1] Let $\phi_t, t \in [0, T]$ denotes the ODE trajectory:

$$\begin{aligned} \phi_t &= \mathbf{x}_t^p - \mathbf{x}_t^q, \\ \mathbf{x}_t^p &= \mathbf{x}_T - \int_T^t \tau \nabla_{\mathbf{x}} \log p_\tau(\mathbf{x}_\tau^p) d\tau, \\ \mathbf{x}_t^q &= \mathbf{x}_T - \int_T^t \tau \nabla_{\mathbf{x}} \log q_\tau(\mathbf{x}_\tau^q) d\tau, \end{aligned} \quad (18)$$

From the definition, $\phi_0 = \Phi_p(\mathbf{x}_T) - \Phi_q(\mathbf{x}_T)$. Because $\lim_{T \rightarrow \infty} \phi_t = \mathbf{x}_T - \mathbf{x}_T = \mathbf{0}$, from the $\epsilon - \delta$ definition of the limit, given \mathbf{x}_T , and a constant ξ , there exists a finite T_ξ related to ξ such that:

$$\|\phi_t\|_2 \leq \xi \quad \text{for all } t \geq T_\xi. \quad (19)$$

As $t \leq T_\xi$, we have:

$$\begin{aligned} \frac{d\phi_t}{dt} &= -t (\nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t^p) - \nabla_{\mathbf{x}} \log q_t(\mathbf{x}_t^q)), \\ \|\phi_{T_0}\|_2 &\leq \xi. \end{aligned} \quad (20)$$

Apply Assumption 1 to Equation (20), we could obtain the following integral inequality w.r.t $\|\phi_t\|_2$:

$$\begin{aligned} \frac{d\|\phi_t\|_2}{dt} &\leq \left\| \frac{d\phi_t}{dt} \right\|_2 \\ &\leq t \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t^p) - \nabla_{\mathbf{x}} \log q_t(\mathbf{x}_t^q)\|_2 \\ &\leq t(\epsilon + L \|\phi_t\|_2), \\ \|\phi_{T_\xi}\|_2 &\leq \xi, \quad 0 \leq t \leq T_\xi, \end{aligned} \quad (21)$$

where the first inequality comes from the fact that $\frac{d\|\phi_t\|_2}{dt} \leq \left\| \frac{d\phi_t}{dt} \right\|_2$. From Grönwall's inequality [19], we could solve $\|\Phi_p(\mathbf{x}_T) - \Phi_q(\mathbf{x}_T)\|_2 = \|\phi_0\|_2 \leq \exp\left(\frac{LT_\xi^2}{2}\right) \xi + \frac{\epsilon}{L} \left(\exp\left(\frac{LT_\xi^2}{2}\right) - 1 \right)$. ■

Proof [Proof of Theorem 2] Let $\mathbf{X} := \|\Phi_p(\mathbf{x}_T) - \Phi_q(\mathbf{x}_T)\|_2^2$. From Theorem 1,

$$0 \leq \mathbf{X} \leq \kappa^2(L, \epsilon),$$

with $\kappa(L, \epsilon) := \exp\left(\frac{LT_\xi^2}{2}\right) \xi + \frac{\epsilon}{L} \left(\exp\left(\frac{LT_\xi^2}{2}\right) - 1 \right)$. From Hoeffding's inequality [30], we have:

$$\mathbb{P} \left(\left| \mathbb{E}[\mathbf{X}] - \frac{1}{M} \sum_{i=1}^M \mathbf{X}_i \right| \geq \gamma \right) \leq 2 \exp \left(-\frac{2M\gamma^2}{\kappa^4(L, \epsilon)} \right), \quad (22)$$

with M samples to achieve γ accuracy. Thus, we could guarantee $\mathbb{P}\left(\left|\mathbb{E}[\mathbf{X}] - \frac{1}{M} \sum_{i=1}^M \mathbf{X}_i\right| \leq \gamma\right)$ with probability η , when:

$$M \geq \frac{\kappa^4(L, \epsilon)}{2\gamma^2} \log \frac{2}{\eta}. \quad (23)$$

Because

$$\left| \text{PFD}(p, q) - \widehat{\text{PFD}}(p, q) \right| = \left| \sqrt{\mathbb{E}[\mathbf{X}]} - \sqrt{\frac{1}{M} \sum_{i=1}^M \mathbf{X}_i} \right| \quad (24)$$

$$\leq \sqrt{\left| \mathbb{E}[\mathbf{X}] - \frac{1}{M} \sum_{i=1}^M \mathbf{X}_i \right|}. \quad (25)$$

We could guarantee that $\mathbb{P}\left(\left| \text{PFD}(p, q) - \widehat{\text{PFD}}(p, q) \right| \leq \gamma\right)$ with probability η , when:

$$M \geq \frac{\kappa^4(L, \epsilon)}{2\gamma^4} \log \frac{2}{\eta}. \quad (26)$$

■

Example 1 *The Wasserstein-2 distance $W_2(\cdot, \cdot)$ is the lower bound of the probability flow distance, i.e.,*

$$W_2(p, q) \leq \text{PFD}(p, q), \quad (27)$$

Specifically, let p and q be multivariate Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$, $\mathcal{N}(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)$, respectively, where $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2 \in \mathbb{R}^n$ and $\boldsymbol{\Sigma}_1, \boldsymbol{\Sigma}_2 \in \mathbb{R}^{n \times n}$. The PFD is given by

$$\text{PFD}(p, q) = \left(\|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2 + \left\| \boldsymbol{\Sigma}_1^{1/2} - \boldsymbol{\Sigma}_2^{1/2} \right\|_F \right)^{1/2}, \quad (28)$$

under this case, the equality in Equation (27) holds when $\boldsymbol{\Sigma}_1 \boldsymbol{\Sigma}_2 = \boldsymbol{\Sigma}_2 \boldsymbol{\Sigma}_1$.

Proof [Proof of Example 1]

Proof of $W_2(p, q) \leq \text{PFD}(p, q)$. From the definition of Wasserstein-2 distance:

$$W_2(p, q) = \inf_{\gamma \in \Gamma(p, q)} \left(\mathbb{E}_{(\mathbf{x}_p, \mathbf{x}_q) \sim \gamma} \|\mathbf{x}_p - \mathbf{x}_q\|_2^2 \right)^{1/2}, \quad (29)$$

where $\Gamma(p, q)$ is the set of all couplings of p and q . As proved by [66], the noise-to-image mapping Φ_p and Φ_q pushes the Gaussian distribution $\mathcal{N}(\mathbf{0}, T^2 \mathbf{I}_n)$ to the p and q distribution respectively. Thus we could find the coupling $\gamma_{\text{PFD}} := (\Phi_p, \Phi_q)_{\#} \mathcal{N}(\mathbf{0}, T^2 \mathbf{I}_n)$, i.e., the pushforward of $\mathcal{N}(\mathbf{0}, T^2 \mathbf{I}_n)$ by (Φ_p, Φ_q) , such that

$$\text{PFD}(p, q) = \left(\mathbb{E}_{(\mathbf{x}_p, \mathbf{x}_q) \sim \gamma_{\text{PFD}}} \|\mathbf{x}_p - \mathbf{x}_q\|_2^2 \right)^{1/2} \geq W_2(p, q) \quad (30)$$

When distribution $p(\mathbf{x})$ is Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with $\boldsymbol{\mu} \in \mathbb{R}^n$ and $\boldsymbol{\Sigma} \in \mathbb{R}^{n \times n}$, from Equation (1), we have $p_t(\mathbf{x})$ is $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma} + \sigma_t^2 \mathbf{I}_n)$, thus the score function could be calculated as,

$$\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) = (\boldsymbol{\Sigma} + t^2 \mathbf{I}_n)^{-1} (\boldsymbol{\mu} - \mathbf{x}). \quad (31)$$

By plugging in Equation (31) to Equation (3), we could obtain the ODE equation w.r.t \mathbf{x} :

$$d\mathbf{x} = -t (\boldsymbol{\Sigma} + t^2 \mathbf{I}_n)^{-1} (\boldsymbol{\mu} - \mathbf{x}) dt, . \quad (32)$$

The above ODE equation has a close form solution:

$$\mathbf{x}_t = \boldsymbol{\mu} + \mathbf{U} \text{diag} \left(\left[\sqrt{\frac{\lambda_1 + t^2}{\lambda_1 + T^2}}, \dots, \sqrt{\frac{\lambda_n + t^2}{\lambda_n + T^2}} \right] \right) \mathbf{U}^\top (\mathbf{x}_T - \boldsymbol{\mu}) \quad (33)$$

where $\mathbf{U}, \lambda_k, k \in [n]$ are singlar value decomposition of $\boldsymbol{\Sigma}$, $\boldsymbol{\Sigma} = \mathbf{U} \text{diag}([\lambda_1, \dots, \lambda_n]) \mathbf{U}^\top$. $\text{diag}(\cdot)$ convert a vector in \mathbb{R}^n into diagonal matrix $\mathbb{R}^{n \times n}$, and $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, T^2 \mathbf{I}_n)$. Let $\mathbf{x}_T = T\boldsymbol{\epsilon}$ with $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$. As $t = 0$ and $T \rightarrow \infty$, we have:

$$\mathbf{x}_t = \left(\mathbf{I}_n - \mathbf{U} \text{diag} \left(\left[\sqrt{\frac{\lambda_1 + t^2}{\lambda_1 + T^2}}, \dots, \sqrt{\frac{\lambda_n + t^2}{\lambda_n + T^2}} \right] \right) \mathbf{U}^\top \right) \boldsymbol{\mu}, \quad (34)$$

$$+ \mathbf{U} \text{diag} \left(\left[T \sqrt{\frac{\lambda_1 + t^2}{\lambda_1 + T^2}}, \dots, T \sqrt{\frac{\lambda_n + t^2}{\lambda_n + T^2}} \right] \right) \mathbf{U}^\top \mathbf{x}_T, \quad (35)$$

$$= \boldsymbol{\mu} + \mathbf{U} \text{diag} \left(\left[\sqrt{\lambda_1}, \dots, \sqrt{\lambda_n} \right] \right) \mathbf{U}^\top \mathbf{x}_T, \quad (36)$$

$$= \boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2} \mathbf{x}_T = \boldsymbol{\Phi}(\mathbf{x}_T). \quad (37)$$

Thus, plugging in Definition 1, we have:

$$\text{PFD}(p, q) = \left(\mathbb{E}_{\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, T^2 \mathbf{I})} \left[\left\| \boldsymbol{\Phi}_1(\mathbf{x}_T) - \boldsymbol{\Phi}_2(\mathbf{x}_T) \right\|_2^2 \right] \right)^{1/2} \quad (38)$$

$$= \left(\mathbb{E}_{\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, T^2 \mathbf{I})} \left[\left\| \boldsymbol{\mu}_1 + \boldsymbol{\Sigma}_1^{1/2} \mathbf{x}_T - \boldsymbol{\mu}_2 - \boldsymbol{\Sigma}_2^{1/2} \mathbf{x}_T \right\|_2^2 \right] \right)^{1/2} \quad (39)$$

$$= \left(\left\| \boldsymbol{\mu}_1 - \boldsymbol{\mu}_2 \right\|_2^2 + \left\| \boldsymbol{\Sigma}_1^{1/2} - \boldsymbol{\Sigma}_2^{1/2} \right\|_F^2 \right)^{1/2} \quad (40)$$

$$= \left(\left\| \boldsymbol{\mu}_1 - \boldsymbol{\mu}_2 \right\|_2^2 + \text{Tr} \left(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2 - 2\boldsymbol{\Sigma}_1^{1/2} \boldsymbol{\Sigma}_2^{1/2} \right) \right)^{1/2} \quad (41)$$

From Wasserstein-2 distance for Gaussian distribution p, q has close form solution:

$$W_2(p, q) = \left(\left\| \boldsymbol{\mu}_1 - \boldsymbol{\mu}_2 \right\|_2^2 + \text{Tr} \left(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2 - 2 \left(\boldsymbol{\Sigma}_1^{1/2} \boldsymbol{\Sigma}_2 \boldsymbol{\Sigma}_1^{1/2} \right)^{1/2} \right) \right)^{1/2}. \quad (42)$$

From Theorem 2, we have $W_2(p, q) \leq \text{PFD}(p, q)$. And specifically, $W_2(p, q) = \text{PFD}(p, q)$ when $\boldsymbol{\Sigma}_1 \boldsymbol{\Sigma}_2 = \boldsymbol{\Sigma}_2 \boldsymbol{\Sigma}_1$. \blacksquare

Lemma 2 Given two positive semi-definite matrix $\boldsymbol{\Sigma}_1, \boldsymbol{\Sigma}_2 \in \mathbb{R}^{n \times n}$,

$$0 \leq \text{Tr} \left(\boldsymbol{\Sigma}_1^{1/2} \boldsymbol{\Sigma}_2^{1/2} \right) \leq \text{Tr} \left(\left(\boldsymbol{\Sigma}_1^{1/2} \boldsymbol{\Sigma}_2 \boldsymbol{\Sigma}_1^{1/2} \right)^{1/2} \right). \quad (43)$$

Table 1: U-Net architectures details.

Name	Dimensions for encoder blocks	Number of residual blocks	Number of parameters $ \theta $
U-Net-1	[32, 32, 32]	4	0.9M
U-Net-2	[64, 64, 64]	4	3.5M
U-Net-3	[96, 96, 96]	4	7.9M
U-Net-4	[128, 128, 128]	4	14.0M
U-Net-5	[80, 160, 160]	4	17.1M
U-Net-6	[160, 160, 160]	3	17.8M
U-Net-7	[160, 160, 160]	4	21.8M
U-Net-8	[192, 192, 192]	4	31.3M
U-Net-9	[224, 224, 224]	4	42.7M
U-Net-10	[256, 256, 256]	4	55.7M

Proof [Proof of Theorem 2] Because Σ_1, Σ_2 are positive semi-definite matrix, $\text{Tr} \left(\Sigma_1^{1/2} \Sigma_2^{1/2} \right) \geq 0$ and

$$\text{Tr} \left(\left(\Sigma_1^{1/2} \Sigma_2 \Sigma_1^{1/2} \right)^{1/2} \right) = \text{Tr} \left(\sqrt{\left(\Sigma_1^{1/2} \Sigma_2^{1/2} \right) \left(\Sigma_1^{1/2} \Sigma_2^{1/2} \right)^\top} \right) = \left\| \Sigma_1^{1/2} \Sigma_2^{1/2} \right\|_*, \quad (44)$$

where $\|\cdot\|_*$ is the nuclear norm (trace norm). From trace norm inequality ([8] Chapter IV, Section 2), for a random matrix M , $\text{Tr}(M) \leq \|M\|_*$. Thus, we have:

$$\text{Tr} \left(\Sigma_1^{1/2} \Sigma_2^{1/2} \right) \leq \left\| \Sigma_1^{1/2} \Sigma_2^{1/2} \right\|_*. \quad (45)$$

■

Appendix D. Experiments

In this section, we provide experimental details and additional discussion of the main results presented in Section 3 and Appendix B.

D.1. Network architecture details

In this subsection, we provide details of the U-Net architectures, as summarized in Table 1. The U-Net follows an encoder-decoder design, where the encoder comprises multiple encoder blocks. The column "**Dimensions for encoder blocks**" indicates the feature dimensions of each encoder block, while "**Number of residual blocks**" specifies how many residual blocks are used within each encoder block. The decoder is symmetric to the encoder. For further architectural details, please refer to [37]. By varying the encoder block dimensions and the number of residual blocks, we scale the U-Net model from 0.9M to 55.7M parameters.

D.2. Evaluation protocol

In this subsection, we provide details of the evaluation protocol introduced in Section 3, as well as the comparison between the synthetic data from the teacher model and the real dataset.

Experiment settings for evaluation protocol. The teacher model θ_t and the student model θ share a similar U-Net architecture [60] with different numbers of parameters, as introduced in Appendix D.1. The teacher model, with UNet-10 architecture, is trained on the CIFAR-10 dataset [39] using the EDM noise scheduler [36], with a batch size of 128 for 1,000 epochs. The student model ² is trained using the variance-preserving (VP) noise scheduler [29], under the same training hyperparameters. We use one A40 GPU with 48 GB video random access memory (VRAM) for all experiments. We generated three subsets of initial noise $\{\mathbf{x}_{\text{train},T}^{(i)}\}_{i=1}^N, \{\mathbf{x}_{\text{gen},T}^{(i)}\}_{i=1}^M, \{\mathbf{x}_{\text{test},T}^{(i)}\}_{i=1}^M \stackrel{\text{iid}}{\sim} \mathcal{N}(0, T^2 \mathbf{I}_n)$. The training and test datasets are produced using the teacher model:

$$\mathcal{D} := \{\mathbf{x}_{\text{train}}^{(i)}\}_{i=1}^N = \{\Phi_{p_{\theta_t}}(\mathbf{x}_{\text{train},T}^{(i)})\}_{i=1}^N, \quad \mathcal{D}_{\text{test}} := \{\mathbf{x}_{\text{test}}^{(i)}\}_{i=1}^M = \{\Phi_{p_{\theta_t}}(\mathbf{x}_{\text{test},T}^{(i)})\}_{i=1}^M.$$

To evaluate the student model, we generate an evaluation dataset from itself:

$$\mathcal{D}_{\text{gen}} := \{\mathbf{x}_{\text{gen}}^{(i)}\}_{i=1}^M = \{\Phi_{p_{\theta}}(\mathbf{x}_{\text{gen},T}^{(i)})\}_{i=1}^M.$$

All samples are generated using the second-order Heun solver [36] with 18 sampling steps. We vary the number of training samples N from 2^6 to 2^{16} in powers of two. M is set to 50,000 for the experiments in Appendix D.3, and 10,000 for the rest.

Experiment settings for validating the synthetic dataset with real real-world dataset. We evaluate FID and \mathcal{E}_{mem} for diffusion models with UNet-4 architecture, trained separately on the synthetic dataset \mathcal{D} and CIFAR-10 training dataset. We keep the number of training dataset N the same for these two settings, ranging from 2^6 to 2^{15} , with a power of 2. Then we evaluate the FID between \mathcal{D}_{gen} and $\mathcal{D}_{\text{test}}$ (CIFAR-10 test dataset) for the synthetic (real-world) setting, with $M = 10000$. To evaluate \mathcal{E}_{mem} , we use the initial noise $\{\mathbf{x}_{\text{gen}}^{(i)}\}_{i=1}^M$.

D.3. Comparison with practical metrics for generalization evaluation

In this subsection, we expand upon the experiment presented in Section 3, which compares our proposed metric with practical metrics for evaluating generalization. We compare \mathcal{E}_{gen} and \mathcal{E}_{mem} with well-used generative model metrics, including FID, $\text{FD}_{\text{DINOv2}}$, KID, CMMD, Precision, Recall, NND, IS. We also including the training and testing loss $\ell_{\text{train}}, \ell_{\text{test}}$ (Equation (10)) as comparison. We evaluating their ability in capturing the MtoG transition, under the evaluation protocol proposed in Section 3.

We use UNet-10 for the student model in this experiment. We summarized datasets used by these metrics in Table 2. Results are shown in Figure 4, summarized into one sentence, only \mathcal{E}_{gen} and \mathcal{E}_{mem} could quantitatively capture this transition. We include detailed discussions below:

Results discussions. Figure 4 (bottom) is consistent with prior empirical observations [73, 75]: In the memorization regimes ($N = 2^6$), p_{θ} tends to memorize the empirical distribution p_{emp} , resulting in similar generation between $\Phi_{p_{\text{emp}}}(\mathbf{x}_T)$ and $\Phi_{p_{\theta}}(\mathbf{x}_T)$; in the transition regime ($N = 2^{12}$), the model lacks sufficient capacity to memorize and the sample complexity is inadequate for generalization, leading to poor-quality generations $\Phi_{p_{\theta}}(\mathbf{x}_T)$; in the generalization regimes ($N = 2^{16}$), p_{θ} captures the underlying distribution p_{data} , and the generations $\Phi_{p_{\text{data}}}(\mathbf{x}_T)$ and $\Phi_{p_{\theta}}(\mathbf{x}_T)$ are closely aligned.

As shown in Figure 4 (top), when N increases, \mathcal{E}_{mem} consistently increases and \mathcal{E}_{gen} consistently decreases. This aligns with our intuition: as sample complexity grows, models tend to generalize and memorize less. In contrast, all other metrics fail to capture this transition effectively. The reasons can be summarized as follows:

2. The architecture of the student model varies across experiments and will be described in detail for each specific case.

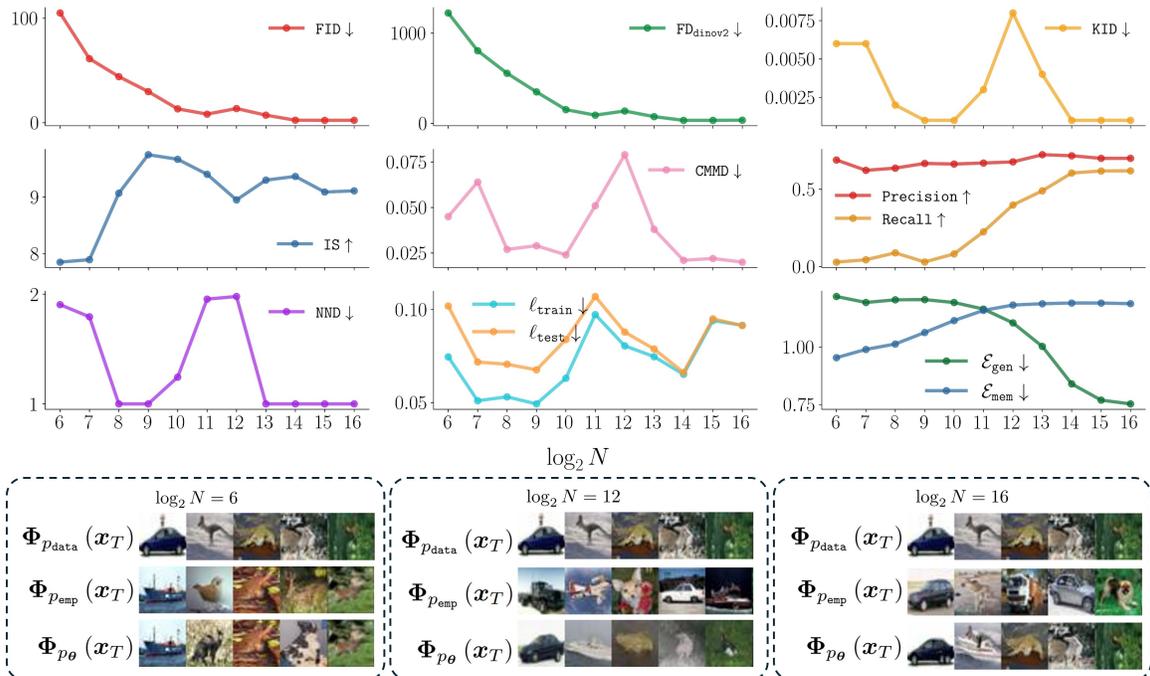


Figure 4: **Comparison of practical metrics on the MtoG transition.** The top figure plots multiple evaluation metrics as functions of $\log_2 N$. The bottom figure visualizes the generation under three numbers of training samples ($2^6, 2^{12}, 2^{16}$). For each setting, the figure shows generations from the underlying distribution (top row), empirical data distribution (middle row), and the learned distribution from the diffusion model (bottom row). Each column corresponds to the same initial noise.

Metric	Dataset(s)
FID, FD _{DINOv2} , KID, CMMD, Precision, Recall, NND	\mathcal{D}_{gen} vs. $\mathcal{D}_{\text{test}}$
FID _{train} , FD _{DINOv2,train}	\mathcal{D} vs. $\mathcal{D}_{\text{test}}$
IS	\mathcal{D}_{gen}
ℓ_{train}	\mathcal{D}
ℓ_{test}	$\mathcal{D}_{\text{test}}$
$\mathcal{E}_{\text{mem}}, \mathcal{E}_{\text{gen}}$	$\{\mathbf{x}_{\text{gen},T}^{(i)}\}_{i=1}^M$

Table 2: Datasets used to evaluate each metric.

- **FID, FD_{DINOv2}, KID, IS, and CMMD are sensitive to generation quality.** Image quality metrics, including FID, FD_{DINOv2}, KID, IS, and CMMD, show degradation in performance at $N = 2^{12}$. This drop is primarily due to degraded visual quality in the generated samples, as visualize in Figure 4 (bottom-middle). However, at this sample complexity, the generated data still captures low-level features such as colors and structures from the underlying distribution. This is evident

from the visual similarity between $\Phi_{p_{\text{data}}}(x_T)$ and $\Phi_{p_{\theta}}(x_T)$, suggesting the model have some generalizability. In comparison, only \mathcal{E}_{gen} decreases consistently around $N = 2^{12}$, indicating it captures generalizability better than others despite visual degradation.

- **FID, $\text{FD}_{\text{DINOv2}}$ and Recall are sensitive to diversity.** The monotonic trends for FID, $\text{FD}_{\text{DINOv2}}$ and Recall are due to their sensitivity to the diversity of \mathcal{D}_{gen} , rather than their ability to measure generalizability. At small N , the model memorizes the training samples, resulting in \mathcal{D}_{gen} closely resembling \mathcal{D} and exhibiting significantly lower diversity than $\mathcal{D}_{\text{test}}$, since $N \ll M$. Under these conditions, FID, $\text{FD}_{\text{DINOv2}}$ are large because they are biased towards the diversity of the evaluation samples (as proved in [18]). Meanwhile, Recall is low because the support of $\mathcal{D}_{\text{test}}$ is limited, reducing the probability that samples drawn from \mathcal{D}_{gen} lie within the support of $\mathcal{D}_{\text{test}}$. In contrast, \mathcal{E}_{gen} measures generalizability by directly quantifying the distance between the generation from the learned distribution and the underlying distribution and is less affected by the diversity of the generated samples.
- **NND and ℓ fail to capture the generalizability.** The NND, originally designed for assessing the generalization of GANs, is sensitive to image quality and increases during the transition regime. Additionally, it produces identical values across a wide range of sample sizes (e.g., $N = 2^8, 2^9, 2^{13}, 2^{14}, 2^{15}, 2^{16}$), making it unreliable for evaluating generalization in diffusion models. Similarly, neither the training loss ℓ_{train} nor the test loss ℓ_{test} exhibits a consistent decreasing trend as N increases, indicating that these losses do not directly reflect either memorization or generalization. While the loss gap $\ell_{\text{test}} - \ell_{\text{train}}$ does tend to decrease with larger N , it cannot serve as a robust generalization metric either. This is because even a randomly initialized model θ can exhibit a small loss gap.

In conclusion, \mathcal{E}_{mem} and \mathcal{E}_{gen} are the only metrics that could capture the MtoG transition for diffusion models. They evaluate the generalization (memorization) by directly measuring the distance between the learned distribution by the diffusion model and the underlying (empirical) distribution. Unlike other metrics, they are less affected by the quality or diversity of the evaluating samples.

D.4. Scaling Behaviors of the MtoG Transition

In this subsection, we provide detailed experimental settings for Appendix B.1, along with additional experiments to further investigate the MtoG transition across more architectures (e.g., Transformer-based models [3]).

Experiment settings The detailed architectures of the student models, from U-Net-1 to U-Net-10, are provided in Appendix D.1, with model sizes ranging from 0.9M to 55.7M parameters. We scale up the architectures by increasing the dimensionality of the encoder blocks and the number of residual blocks.

MtoG Transition between U-Net and Transformer architecture To further investigate the impact of network architecture, we compare the U-Net architecture with the transformer-based UViT [3]. Specifically, we use the U-Net-9 from Table 1, containing 42.7M parameters, and design the UViT model with comparable parameters of 44.2M. Both models are trained for 1000 epochs. Using the same experimental setup described in Appendix B.1, we plot the MtoG transition curves for both U-Net and UViT, as shown in Figure 5.

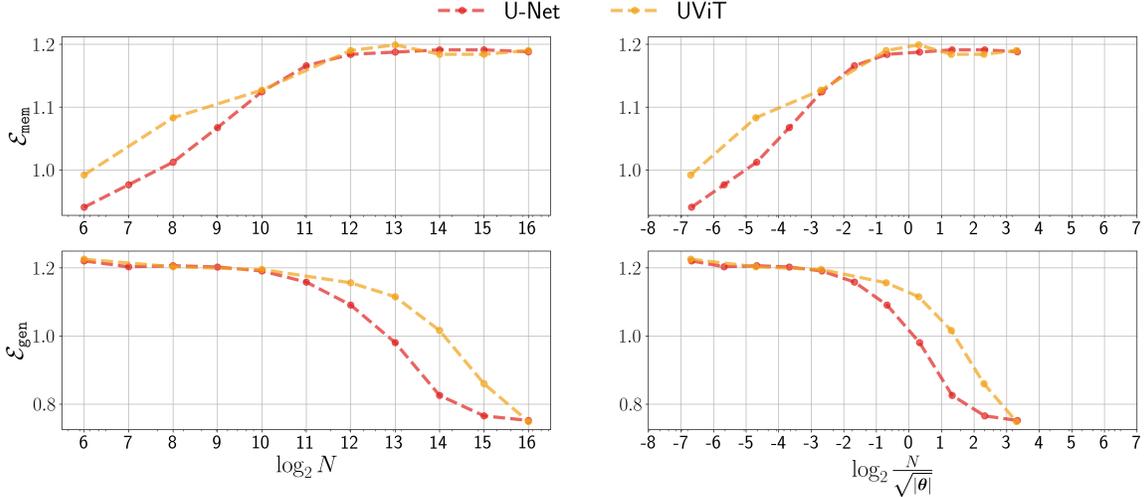


Figure 5: **Comparison of scaling behavior between U-Net and Transformer architectures in the MtoG transition.** \mathcal{E}_{mem} and \mathcal{E}_{gen} plotted against $\log_2(N)$ for U-Net architecture (U-Net-9) and UViT architecture. Right: the same metrics plotted against $\log_2(N/\sqrt{|\theta|})$, where $|\theta|$ is the number of model parameters.

As illustrated in Figure 5, with a similar number of parameters and the same training data sizes, UViT exhibits a higher \mathcal{E}_{mem} in the memorization regime ($2^6 \leq N \leq 2^{10}$) and a higher \mathcal{E}_{gen} in the generalization regime ($2^{11} \leq N \leq 2^{15}$), suggesting a lower model capacity compared to U-Net under these conditions. However, when provided with sufficient training data ($N = 2^{16}$), UViT achieves a lower \mathcal{E}_{gen} , demonstrating better generalization performance. This observation is consistent with prior findings on transformer architectures in classification tasks: transformer-based models, lacking the inductive biases inherent to CNNs, tend to generalize poorly when trained on limited data [20].

D.5. Early Learning and Double Descent in Learning Dynamics

In this subsection, we build on the discussion from Appendix B.2. In Figure 2, we evaluate ℓ_{train} and ℓ_{test} across the three training regimes. Notably, the gap $\ell_{\text{test}} - \ell_{\text{train}}$ emerges as a practical heuristic for identifying the training regime: In the memorization regime, the gap increases steadily with training; In the transition regime, the gap remains near zero during early training (when generalization improves) and increases for further training (when generalization degrades); in the generalization regime, the gap remains close to zero throughout training. While $\ell_{\text{test}} - \ell_{\text{train}}$ is not a strict measure of generalization, it proves to be a useful empirical indicator of training regimes for diffusion models. Practically, by setting aside a test dataset to estimate this gap, we can more effectively identify the training regime for diffusion models.

D.6. Bias-variance Decomposition of Generalization Error

To approximate $\overline{\Psi \circ \Phi_{p_\theta}(\cdot)}$, we independently sample two training datasets, \mathcal{D}_1 and \mathcal{D}_2 , for each specified number of training samples N . We then train two student models, $\theta(\mathcal{D}_1)$ and $\theta(\mathcal{D}_2)$, using these datasets. The quantity $\overline{\Psi \circ \Phi_{p_\theta}(\cdot)}$ is approximated as follows:

$$\overline{\Psi \circ \Phi_{p_\theta}}(\cdot) \approx \frac{1}{2}(\Psi \circ \Phi_{p_{\theta(\mathcal{D}_1)}}(\cdot) + \Psi \circ \Phi_{p_{\theta(\mathcal{D}_2)}}(\cdot)). \quad (46)$$

Appendix E. Further Discussions of \mathcal{E}_{mem}

In this section, we present the mathematical formulation for estimating \mathcal{E}_{mem} and compare it with existing memorization metric.

Empirically estimate \mathcal{E}_{mem} . As described in Definition 1 and Definition 2, estimating \mathcal{E}_{mem} requires access to the mapping $\Phi_{p_{\text{emp}}}(\cdot)$. According to Equation (3), this mapping is determined by the score function of the empirical distribution, denoted as $\nabla \log \hat{p}_t(\mathbf{x}_t)$. Based on prior works [25, 36, 74], the score function of the empirical distribution has a closed-form expression:

$$\nabla \log \hat{p}_t(\mathbf{x}_t) = \frac{1}{T^2} \left(\frac{\mathbb{E}_{\mathbf{x} \sim p_{\text{emp}}}[\mathcal{N}(\mathbf{x}_t; \mathbf{x}, T^2 \mathbf{I}_n) \cdot \mathbf{x}]}{\mathbb{E}_{\mathbf{x} \sim p_{\text{emp}}}[\mathcal{N}(\mathbf{x}_t; \mathbf{x}, T^2 \mathbf{I}_n)]} - \mathbf{x}_t \right), \quad (47)$$

where $p_{\text{emp}}(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N \delta(\mathbf{x} - \mathbf{y}^{(i)})$ corresponds to the empirical distribution over the training dataset $\mathbf{y}^{(i)}_{i=1}^N$. This formulation allows us to numerically compute $\nabla \log \hat{p}_t(\mathbf{x}_t)$ for any given t . Subsequently, we can use a numerical solver to estimate the integral in Equation (3), thereby enabling the estimation of \mathcal{E}_{mem} .

Comparison between existing memorization metric and \mathcal{E}_{mem} . Previous works [73, 75] define memorization metric as:

$$\text{M Distance}(p_\theta) := \mathbb{E}_{\mathbf{x}_T} \left[\min_{\mathbf{x} \sim p_{\text{emp}}} \|\Psi(\mathbf{x}) - \Psi \circ \Phi_{p_\theta}(\mathbf{x}_T)\|_2 \right], \quad (48)$$

A generated sample $\Phi_{p_\theta}(\mathbf{x}_T)$ is a memorized sample if it is close enough to one of the sample \mathbf{x} from p_{emp} . It is easy to show that \mathcal{E}_{mem} is a more strict metric than M Distance, i.e. " $\mathcal{E}_{\text{mem}}(p^\theta) = 0$ " is a sufficient but not necessary condition for " $\text{M Distance}(p^\theta) = 0$ ".

Appendix F. Ablation Study

In this section, we present ablation studies on the evaluation protocol, examining the effects of different noise schedulers and sampling methods (Appendix F.1), image descriptors (Appendix F.2), sample sizes for evaluation (Appendix F.3), and teacher models (Appendix F.4).

F.1. Sampling Methods

In this subsection, we present ablation studies on various noise schedulers and sampling strategies. Specifically, we evaluate the performance of the following methods: Variance Preserving (VP) [66], Variance Exploding (VE) [66], iDDPM [53] + DDIM [65], and EDM [36]. The specific form of $f(t), g(t)$ used in each approach are detailed in Table 1 of [36]. Additionally, each method also differs in its choice of ODE solver and timestep discretization strategy. For sampling, we use 256 steps for VP, 1000 for VE, 100 for iDDPM + DDIM, and 18 for EDM. All experiments are conducted under the evaluation protocol described in Section 3, where we estimate the \mathcal{E}_{gen} under different

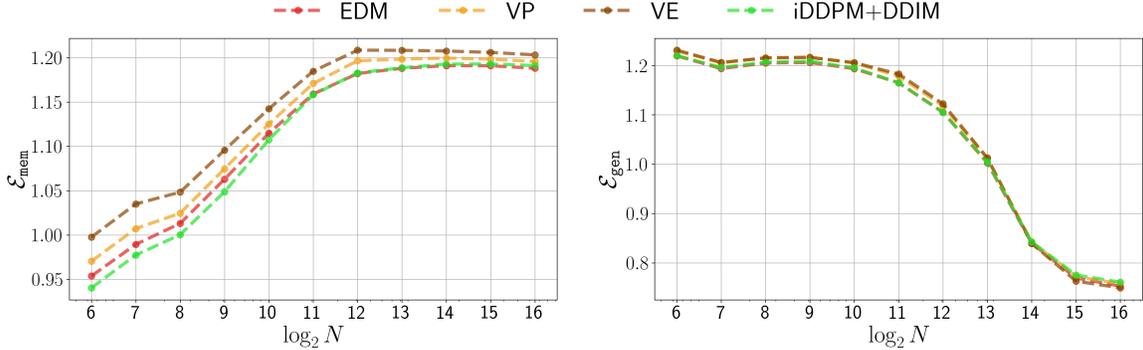


Figure 6: **Comparison of different sampling methods.** \mathcal{E}_{mem} and \mathcal{E}_{gen} plotted against $\log_2(N)$ for different sampling methods, including: EDM, VP, VE, iDDPM+DDIM.

training samples N . The student models use the UNet-10 architecture. During the ablation study, both the teacher and student models use the same sampling method³ as specified above.

As shown in Figure 6, different samplers yield highly consistent results, demonstrating that PFD can be extended to various noise schedules, i.e., different choices of $f(t)$ and $g(t)$.

F.2. Image Descriptors

In this subsection, we present ablation studies on the image descriptor Ψ used in Equation (4). The descriptors evaluated include DINOv2 [56], InceptionV3 [68], CLIP [58], SSCD [57], and the identity function. All experiments follow the evaluation protocol described in Section 3, where we estimate both \mathcal{E}_{mem} and \mathcal{E}_{gen} across varying training sample sizes N and different student model architectures: U-Net-1, U-Net-2, U-Net-4, and U-Net-10.

As shown in Figure 7, different feature embeddings reveal a consistent trend in the memorization-to-generalization (MtoG) transition across various U-Net architectures. With limited training samples, smaller models exhibit lower generalization scores. Conversely, with sufficient training data, larger models tend to have lower generalization scores. When comparing with \mathcal{E}_{gen} measured in pixel space (i.e., using the identity function as the descriptor), we observe that \mathcal{E}_{gen} values are nearly identical across diffusion architectures when $N \geq 2^{15}$. In this regime, all models have learned low-level image features such as color and structure; however, only the larger models capture high-level perceptual details. Because pixel-space measurements fail to reflect these high-level features, they yield similar \mathcal{E}_{gen} values regardless of model size. Therefore, it is better to evaluate \mathcal{E}_{gen} in a feature space, which better captures perceptual differences between models.

F.3. Evaluation Sample Number

In this subsection, we present ablation studies on the number of samples M used by PFD to approximate PFD, as defined in Equation (5). All experiments follow the evaluation protocol described in Section 3, where we estimate \mathcal{E}_{gen} across varying training sample sizes N and different student model architectures: U-Net-1, U-Net-2, U-Net-4, and U-Net-10. We vary $M \in$

3. Note that the noise scheduler used for sampling could differ from that used during training.

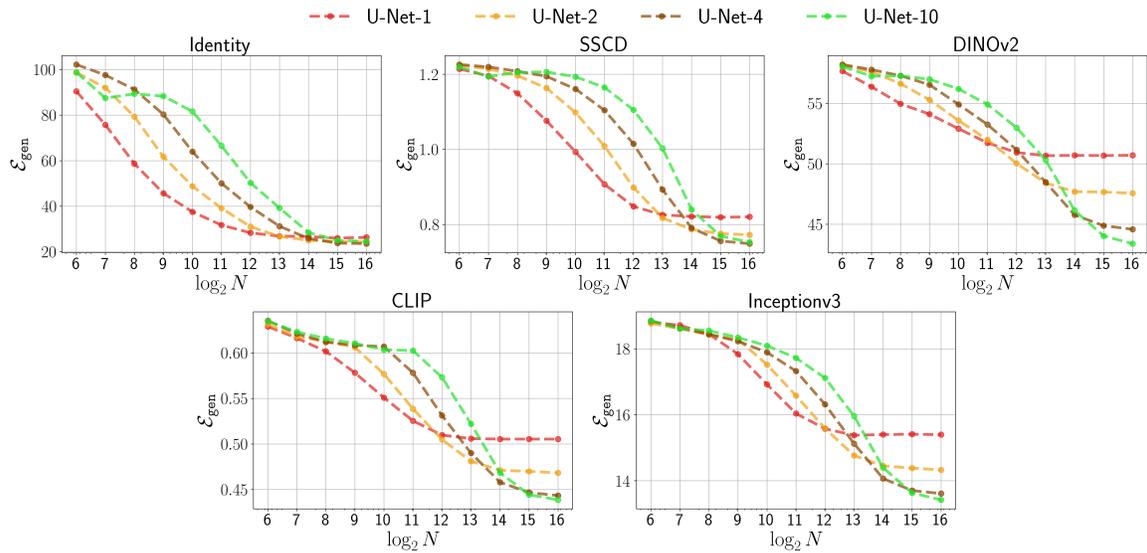


Figure 7: **Comparison between different image descriptors.** \mathcal{E}_{gen} plotted against $\log_2(N)$ for a range of U-Net architectures (U-Net-1, U-Net-2, U-Net-4, U-Net-10) using different image descriptors, including identity function, SSCD, DINOv2, CLIP, Inceptionv3.

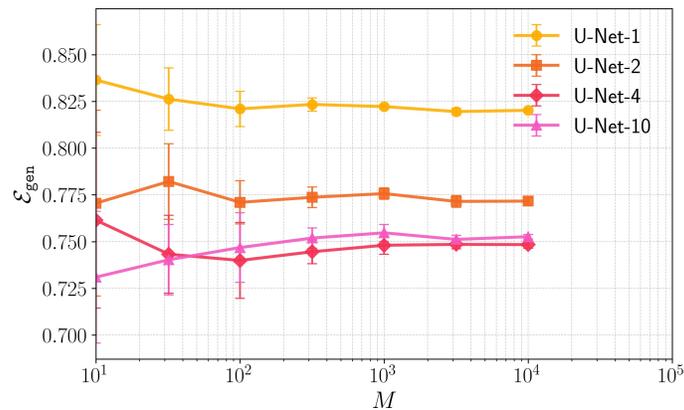


Figure 8: **Comparison across evaluation sample sizes.** The mean and variance of \mathcal{E}_{gen} are plotted against the number of evaluation samples M for various U-Net architectures (U-Net-1, U-Net-2, U-Net-4, U-Net-10), with a fixed number of training samples $N = 2^{16}$.

$\{10, 32, 100, 316, 1000, 3163, 10000\}$, and for each setting, generate 5 independent sets of $\{\mathbf{x}_{\text{gen},T}^{(i)}\}_{i=1}^M$ initial noise estimate \mathcal{E}_{gen} , computing both the mean and variance. In this experiment, we use DINOv2 as the image descriptor.

As shown in Figure 8, the variance of \mathcal{E}_{gen} approaches zero as M increases to 10,000, indicating that when $M \geq 10000$, the empirical estimate of \mathcal{E}_{gen} converges to its value over the underlying distribution. This result holds consistently across different model architectures.

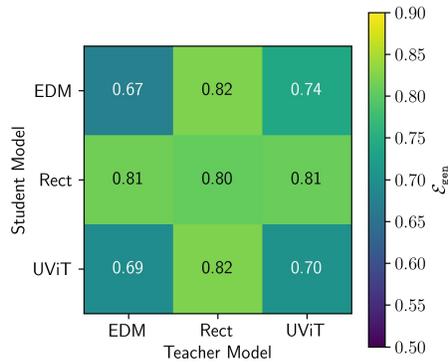


Figure 9: **Comparison of different teacher models.** The figure shows the \mathcal{E}_{gen} values for various student models (EDM, Rect, UViT) trained using different teacher models (EDM, Rect, UViT), with a fixed training data size of $N = 2^{16}$.

F.4. Teacher Model Architecture

We end this section by examining how different teacher models affect the evaluation protocol. Specifically, we consider three types of diffusion models: EDM, Rectified Flow (Rect) [50], and UViT. Using the CIFAR-10 dataset, we train three teacher models, one for each of these diffusion types. For each teacher model, we then evaluate all three diffusion models as student models. We report their corresponding \mathcal{E}_{gen} values. Both teacher and student models use the same sampling method, the second-order Heun solver with 18 steps.

As shown in Figure 9, the \mathcal{E}_{gen} is approximately 0.7 when both the student and teacher models are selected from EDM or UViT. However, \mathcal{E}_{gen} increases to around 0.8 when either the student or teacher model is Rect. According to its original paper, Rect has the poorest generative quality among the three, as measured by FID. This suggests that the teacher model should possess strong generative performance to serve as an underlying distribution that is close to the real-world data distribution. Therefore, in this paper, we adopt EDM as the teacher model, as it achieves the lowest FID among the three models.