

---

# Guiding a Diffusion Model with a Bad Version of Itself

---

**Tero Karras**  
NVIDIA

**Miika Aittala**  
NVIDIA

**Tuomas Kynkäänniemi**  
Aalto University

**Jaakko Lehtinen**  
NVIDIA, Aalto University

**Timo Aila**  
NVIDIA

**Samuli Laine**  
NVIDIA

## Abstract

The primary axes of interest in image-generating diffusion models are image quality, the amount of variation in the results, and how well the results align with a given condition, e.g., a class label or a text prompt. The popular classifier-free guidance approach uses an unconditional model to guide a conditional model, leading to simultaneously better prompt alignment and higher-quality images at the cost of reduced variation. These effects seem inherently entangled, and thus hard to control. We make the surprising observation that it is possible to obtain disentangled control over image quality without compromising the amount of variation by guiding generation using a smaller, less-trained version of the model itself rather than an unconditional model. This leads to significant improvements in ImageNet generation, setting record FIDs of 1.01 for  $64 \times 64$  and 1.25 for  $512 \times 512$ , using publicly available networks. Furthermore, the method is also applicable to unconditional diffusion models, drastically improving their quality.

## 1 Introduction

Denoising diffusion models [15, 46, 47, 48, 49] generate synthetic images by reversing a stochastic corruption process. Essentially, an image is revealed from pure noise by denoising it little by little in successive steps. A neural network that implements the denoiser (equivalently [53], the score function [20]) is a central design element, and various architectures have been proposed (e.g., [2, 7, 9, 19, 21, 24, 38]). Equally important are the details of the multi-step denoising process that corresponds mathematically to solving an ordinary [35, 47] or a stochastic [49] differential equation, for which many different parameterizations, solvers, and step schedules have been evaluated [22, 23, 30, 43, 57]. To control the output image, the denoiser is typically conditioned on a class label, an embedding of a text prompt, or some other form of conditioning input [36, 41, 45, 56].

The training objective of a diffusion model aims to cover the entire (conditional) data distribution. This causes problems in low-probability regions: The model gets heavily penalized for not representing them, but it does not have enough data to learn to generate good images corresponding to them. Classifier-free guidance (CFG) [16] has become the standard method for “lowering the sampling temperature”, i.e., focusing the generation on well-learned high-probability regions. By training a denoiser network to operate in both conditional and unconditional setting, the sampling process can be steered away from the unconditional result — in effect, the unconditional generation task specifies a result to *avoid*. This results in better prompt alignment and improved image quality, where the former effect is due to CFG implicitly raising the conditional part of the probability density to a power greater than one [10].

However, CFG has drawbacks that limit its usage as a general low-temperature sampling method. First, it is applicable only for conditional generation, as the guidance signal is based on the difference between conditional and unconditional denoising results. Second, because the unconditional and

conditional denoisers are trained to solve a different task, the sampling trajectory can overshoot the desired conditional distribution, which leads to skewed and often overly simplified image compositions [27]. Finally, the prompt alignment and quality improvement effects cannot be controlled separately, and it remains unclear how exactly they relate to each other.

In this paper, we provide new insights into why CFG improves image quality and show how this effect can be separated out into a novel method that we call *autoguidance*. Our method does not suffer from the task discrepancy problem because we use an inferior version of the main model itself as the guiding model, with unchanged conditioning. This guiding model can be obtained by simply limiting, e.g., model capacity and/or training time. We validate the effectiveness of autoguidance in various synthetic test cases as well as in practical image synthesis in class-conditional and text-conditional settings. In addition, our method enables guidance for unconditional synthesis. In quantitative tests, the generated image distributions are improved considerably when measured using FID [13] and  $FD_{\text{DINOv2}}$  [51] metrics, setting new records in ImageNet-512 and ImageNet-64 generation.

Our implementation and pre-trained models are available at <https://github.com/NVlabs/edm2>

## 2 Background

**Denosing diffusion.** Denosing diffusion generates samples from a distribution  $p_{\text{data}}(\mathbf{x})$  by iteratively denoising a sample of pure white noise, such that a noise-free random data sample is gradually revealed [15]. The idea is to consider heat diffusion of  $p_{\text{data}}(\mathbf{x})$  into a sequence of increasingly smoothed densities  $p(\mathbf{x}; \sigma) = p_{\text{data}}(\mathbf{x}) * \mathcal{N}(\mathbf{x}; \mathbf{0}, \sigma^2 \mathbf{I})$ . For a large enough  $\sigma_{\text{max}}$ , we have  $p(\mathbf{x}; \sigma_{\text{max}}) \approx \mathcal{N}(\mathbf{x}; \mathbf{0}, \sigma_{\text{max}}^2 \mathbf{I})$ , from which we can trivially sample by drawing normally distributed white noise. The resulting sample is then evolved backward towards low noise levels by a probability flow ODE [23, 47, 49]

$$d\mathbf{x}_\sigma = -\sigma \nabla_{\mathbf{x}_\sigma} \log p(\mathbf{x}_\sigma; \sigma) d\sigma \quad (1)$$

that maintains the property  $\mathbf{x}_\sigma \sim p(\mathbf{x}_\sigma; \sigma)$  for every  $\sigma \in [0, \sigma_{\text{max}}]$ . Upon reaching  $\sigma = 0$ , we obtain  $\mathbf{x}_0 \sim p(\mathbf{x}_0; 0) = p_{\text{data}}(\mathbf{x}_0)$  as desired.

In practice, the ODE is solved numerically by stepping along the trajectory defined by Equation 1. This requires evaluating the so-called score function [20]  $\nabla_{\mathbf{x}} \log p(\mathbf{x}; \sigma)$  for a given sample  $\mathbf{x}$  and noise level  $\sigma$  at each step. Rather surprisingly, we can approximate this vector using a neural network  $D_\theta(\mathbf{x}; \sigma)$  parameterized by weights  $\theta$  trained for the denoising task

$$\theta = \arg \min_{\theta} \mathbb{E}_{\mathbf{y} \sim p_{\text{data}}, \sigma \sim p_{\text{train}}, \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})} \|D_\theta(\mathbf{y} + \mathbf{n}; \sigma) - \mathbf{y}\|_2^2, \quad (2)$$

where  $p_{\text{train}}$  controls the noise level distribution during training. Given  $D_\theta$ , we can estimate  $\nabla_{\mathbf{x}} \log p(\mathbf{x}; \sigma) \approx (D_\theta(\mathbf{x}; \sigma) - \mathbf{x})/\sigma^2$ , up to approximation errors due to, e.g., finite capacity or training time [23, 53]. As such, we are free to interpret the network as predicting either a denoised sample or a score vector, whichever is more convenient for the analysis at hand. Many reparameterizations and practical ODE solvers are possible, as enumerated by Karras et al. [23]. We follow their recommendations, including the schedule  $\sigma(t) = t$  that lets us parameterize the ODE directly via noise level  $\sigma$  instead of a separate time variable  $t$ .

In most applications, each data sample  $\mathbf{x}$  is associated with a label  $\mathbf{c}$ , representing, e.g., a class index or a text prompt. At generation time, we control the outcome by choosing a label  $\mathbf{c}$  and seeking a sample from the conditional distribution  $p(\mathbf{x}|\mathbf{c}; \sigma)$  with  $\sigma = 0$ . In practice, this is achieved by training a denoiser network  $D_\theta(\mathbf{x}; \sigma, \mathbf{c})$  that accepts  $\mathbf{c}$  as an additional conditioning input.

**Classifier-free guidance.** For complex visual datasets, the generated images often fail to reproduce the clarity of the training images due to approximation errors made by finite-capacity networks. A broadly used trick called *classifier-free guidance* (CFG) [16] pushes the samples towards higher likelihood of the class label, sacrificing variety for “more canonical” images that the network appears to be better capable of handling.

In a general setting, guidance in a diffusion model involves two denoiser networks  $D_0(\mathbf{x}; \sigma, \mathbf{c})$  and  $D_1(\mathbf{x}; \sigma, \mathbf{c})$ . The guiding effect is achieved by *extrapolating* between the two denoising results by a factor  $w$ :

$$D_w(\mathbf{x}; \sigma, \mathbf{c}) = wD_1(\mathbf{x}; \sigma, \mathbf{c}) + (1 - w)D_0(\mathbf{x}; \sigma, \mathbf{c}). \quad (3)$$

Trivially, setting  $w = 0$  or  $w = 1$  recovers the output of  $D_0$  and  $D_1$ , respectively, while choosing  $w > 1$  over-emphasizes the output of  $D_1$ . Recalling the equivalence of denoisers and scores [53], we can write

$$D_w(\mathbf{x}; \sigma, \mathbf{c}) \approx \mathbf{x} + \sigma^2 \nabla_{\mathbf{x}} \log \underbrace{\left( p_0(\mathbf{x}|\mathbf{c}; \sigma) \left[ \frac{p_1(\mathbf{x}|\mathbf{c}; \sigma)}{p_0(\mathbf{x}|\mathbf{c}; \sigma)} \right]^w \right)}_{\propto: p_w(\mathbf{x}|\mathbf{c}; \sigma)}. \quad (4)$$

Thus, guidance grants us access to the score of the density  $p_w(\mathbf{x}|\mathbf{c}; \sigma)$  implied in the parentheses. This score can be further written as [10, 16]

$$\nabla_{\mathbf{x}} \log p_w(\mathbf{x}|\mathbf{c}; \sigma) = \nabla_{\mathbf{x}} \log p_1(\mathbf{x}|\mathbf{c}; \sigma) + (w - 1) \nabla_{\mathbf{x}} \log \frac{p_1(\mathbf{x}|\mathbf{c}; \sigma)}{p_0(\mathbf{x}|\mathbf{c}; \sigma)}. \quad (5)$$

Substituting this expression into the ODE of Equation 1, this yields the standard evolution for generating images from  $p_1$ , plus a perturbation that increases (for  $w > 1$ ) the ratio of  $p_1$  and  $p_0$  as evaluated at the sample. The latter can be interpreted as increasing the likelihood that a hypothetical classifier would attribute for the sample having come from density  $p_1$  rather than  $p_0$ .

In CFG, we train an auxiliary *unconditional* denoiser  $D_\theta(\mathbf{x}; \sigma)$  to denoise the distribution  $p(\mathbf{x}; \sigma)$  marginalized over  $\mathbf{c}$ , and use this as  $D_0$ . In practice, this is typically [16] done using the same network  $D_\theta$  with an empty conditioning label, setting  $D_0 := D_\theta(\mathbf{x}; \sigma, \emptyset)$  and  $D_1 := D_\theta(\mathbf{x}; \sigma, \mathbf{c})$ . By Bayes’ rule, the extrapolated score vector becomes  $\nabla_{\mathbf{x}} \log p(\mathbf{x}|\mathbf{c}; \sigma) + (w - 1) \nabla_{\mathbf{x}} \log p(\mathbf{c}|\mathbf{x}; \sigma)$ . During sampling, this guides the image to more strongly align with the specified class  $\mathbf{c}$ .

It would be tempting to conclude that solving the diffusion ODE with the score function of Equation 5 produces samples from the data distribution specified by  $p_w(\mathbf{x}|\mathbf{c}; 0)$ . Unfortunately this is *not* the case, because  $p_w(\mathbf{x}|\mathbf{c}; \sigma)$  does not represent a valid heat diffusion of  $p_w(\mathbf{x}|\mathbf{c}; 0)$  [58]. Therefore, solving the ODE does not, in fact, follow the density. Instead, the samples are blindly pushed towards higher values of the implied density at each noise level during sampling. This can lead to distorted sampling trajectories, greatly exaggerated truncation, and mode dropping in the results [27], as well as over-saturation of colors [45]. Nonetheless, the improvement in image quality is often remarkable, and high guidance values are commonly used despite the drawbacks (e.g., [14, 39, 41, 45]).

### 3 Why does CFG improve image quality?

We begin by identifying the mechanism by which CFG improves image quality instead of only affecting prompt alignment. To illustrate why unguided diffusion models often produce unsatisfactory images, and how CFG remedies the problem, we study a 2D toy example where a small-scale denoiser network is trained to perform conditional diffusion in a synthetic dataset (Figure 1). The dataset is designed to exhibit low local dimensionality (i.e., highly anisotropic and narrow support) and hierarchical emergence of local detail upon noise removal. These are both properties that can be expected from the actual manifold of realistic images [5, 40]. For details of the setup, see Appendix C.

**Score matching leads to outliers.** Compared to sampling directly from the underlying distribution (Figure 1a), the unguided diffusion in Figure 1b produces a large number of extremely unlikely samples outside the bulk of the distribution. In the image generation setting, these would correspond to unrealistic and broken images.

We argue that the outliers stem from the limited capability of the score network combined with the score matching objective. It is well known that maximum likelihood (ML) estimation leads to a “conservative” fit of the data distribution [3] in the sense that the model attempts to cover all training samples. This is because the underlying Kullback–Leibler divergence incurs extreme penalties if the model severely underestimates the likelihood of any training sample. While score matching is generally not equal to ML estimation, they are closely related [15, 32, 49] and appear to exhibit broadly similar behavior. For example, it is known that for a multivariate Gaussian model, the optimal score matching fit coincides with the ML estimate [20]. Figures 2a and 2b show the learned score field and implied density in our toy example for two models of different capacity at an intermediate noise level. The stronger model envelops the data more tightly, while the weaker model’s density is more spread out.

From the perspective of image generation, a tendency to cover the entire training data becomes a problem: The model ends up producing strange and unlikely images from the data distribution’s




Figure 1: A fractal-like 2D distribution with two classes indicated with gray and orange regions. Approximately 99% of the probability mass is inside the shown contours. **(a)** Ground truth samples drawn directly from the orange class distribution. **(b)** Conditional sampling using a small denoising diffusion model generates outliers. **(c)** Classifier-free guidance ( $w = 4$ ) eliminates outliers but reduces diversity by over-emphasizing the class. **(d)** Naive truncation via lengthening the score vectors. **(e)** Our method concentrates samples on high-probability regions without reducing diversity.

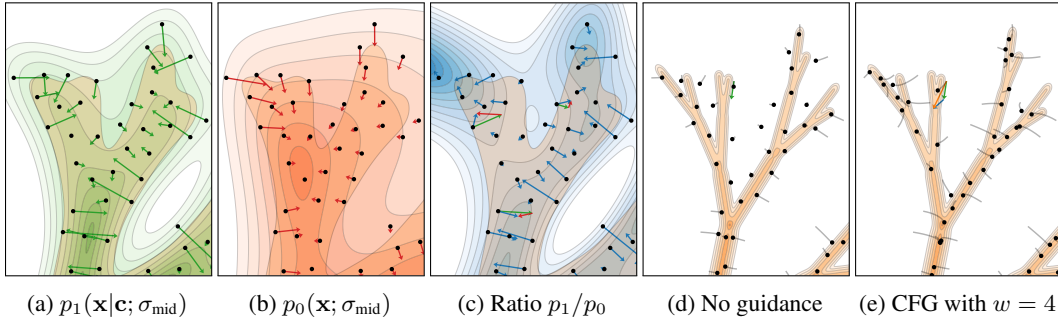


Figure 2: Closeup of the region highlighted in Figure 1c. **(a)** The implied learned density  $p_1(\mathbf{x}|\mathbf{c}; \sigma_{\text{mid}})$  (green) at an intermediate noise level  $\sigma_{\text{mid}}$  and its score vectors (log-gradients), plotted at representative sample points. The learned density approximates the underlying ground truth  $p(\mathbf{x}|\mathbf{c}; \sigma_{\text{mid}})$  (orange) but fails to replicate its sharper details. **(b)** The weaker unconditional model learns a further spread-out density  $p_0(\mathbf{x}; \sigma_{\text{mid}})$  (red) with a looser fit to the data. **(c)** Guidance moves the points according to the gradient of the (log) ratio of the two learned densities (blue). As the higher-quality model is more sharply concentrated at the data, this field tends inward towards the data distribution. The corresponding gradient is simply the difference of respective gradients in (a) and (b), illustrated at selected points. **(d)** Sampling trajectories taken by standard unguided diffusion following the learned score  $\nabla_{\mathbf{x}} \log p_1(\mathbf{x}|\mathbf{c}; \sigma)$ , from noise level  $\sigma_{\text{mid}}$  to 0. The contours (orange) represent the ground truth noise-free density. **(e)** Guidance introduces an additional force shown in (c), causing the points to concentrate at the core of the data density during sampling.

extremities that are not learnt accurately but included just to avoid the high loss penalties. Furthermore, during training, the network has only seen real noisy images as inputs, and during sampling it may not be prepared to deal with the unlikely samples it is handed down from the higher noise levels.

**CFG eliminates outliers.** The effect of applying classifier-free guidance during generation is demonstrated in Figure 1c. As expected, the samples avoid the class boundary (i.e., there are no samples in the vicinity of the gray area), and entire branches of the distribution are dropped. We also observe a second phenomenon, where the samples have been pulled in towards the core of the manifold, and away from the low-probability intermediate regions. Seeing that this eliminates the unlikely outlier samples, we attribute the image quality improvement to it. However, mere boosting of the class likelihood does not explain this increased concentration.

We argue that this phenomenon stems from a quality difference between the conditional and unconditional denoiser networks. The denoiser  $D_0$  faces a more difficult task of the two: It has to generate from all classes at once, whereas  $D_1$  can focus on a single class for any specific sample. Given the more difficult task, and typically only a small slice of the training budget, the network  $D_0$  attains

a worse fit to the data.<sup>1</sup> This difference in accuracy is apparent in respective plots of the learned densities in Figures 2a and 2b.

From our interpretation in Section 2, it follows that CFG is not only boosting the likelihood of the sample having come from the class  $c$ , but *also* that of having come from the higher-quality implied distribution. Recall that guidance boils down to an additional force (Equation 5) that pulls the samples towards higher values of  $\log [p_1(\mathbf{x}|c; \sigma)/p_0(\mathbf{x}|c; \sigma)]$ . Plotting this ratio for our toy example in Figure 2c, along with corresponding gradients that guidance contributes to the ODE vector field, we see that the ratio generally decreases with distance from the manifold due to the denominator  $p_0$  representing a more spread-out distribution, and hence falling off slower than the numerator  $p_1$ . Consequently, the gradients point inward towards the data manifold. Each contour of the density ratio corresponds to a specific likelihood that a hypothetical classifier would assign on a sample being drawn from  $p_1$  instead of  $p_0$ . Because the contours roughly follow the local orientation and branching of the data manifold, pushing samples deeper into the “good side” concentrates them at the manifold.<sup>2</sup>

**Discussion.** We can expect the two models to suffer from inability to fit at similar places, but to a different degree. The predictions of the denoisers will disagree more decisively in these regions. As such, CFG can be seen as a form of adaptive truncation that identifies when a sample is likely to be under-fit and pushes it towards the general direction of better samples. Figures 2d and 2e show the effect over the course of generation: The truncation “overshoots” the correction and produces a narrower distribution than the ground truth, but in practice this does not appear to have an adverse effect on the images.

In contrast, a naive attempt at achieving this kind of truncation — inspired by, e.g., the truncation trick in GANs [4, 33] or lowering temperature in generative language models — would counteract the smoothing by uniformly lengthening the score vectors by a factor  $w > 1$ . This is illustrated in Figure 1d, where the samples are indeed concentrated in high-probability regions, but in an isotropic fashion that leaves the outer branches empty. In practice, images generated this way tend to show reduced variation, oversimplified details, and monotone texture.

## 4 Our method

We propose to isolate the image quality improvement effect by directly guiding a high-quality model  $D_1$  with a poor model  $D_0$  trained on the *same task, conditioning, and data distribution*, but suffering from certain additional degradations, such as low capacity and/or under-training. We call this procedure *autoguidance*, as the model is guided with an inferior version of itself.

In the context of our 2D toy example, this turns out to work surprisingly well. Figure 1e demonstrates the effect of using a smaller  $D_0$  with fewer training iterations. As desired, the samples are pulled close to the distribution without systematically dropping any part of it.

To analyze why this technique works, recall that under limited model capacity, score matching tends to over-emphasize low-probability (i.e., implausible and under-trained) regions of the data distribution. Exactly where and how the problems appear depend on various factors such as network architecture, dataset, training details, etc., and we cannot expect to identify and characterize the specific issues a priori. However, we can expect a weaker version of the *same* model to make broadly similar errors in the same regions, only stronger. Autoguidance seeks to identify and reduce the errors made by the stronger model by measuring its difference to the weaker model’s prediction, and boosting it. When the two models agree, the perturbation is insignificant, but when they disagree, the difference indicates the general direction towards better samples.

As such, we can expect autoguidance to work if the two models suffer from degradations that are compatible with each other. Since any  $D_1$  can be expected to suffer from, e.g., lack of capacity and

<sup>1</sup>The visual quality difference is obvious if we simply inspect the unconditional images generated by current large-scale models. Furthermore, the unconditional case tends to work so poorly that the corresponding quantitative numbers are hardly ever reported. The EDM2-S model [24] trained with ImageNet-512, for example, yields a FID of 2.56 in the class-conditional setting and 11.67 in the unconditional setting.

<sup>2</sup>Discriminator guidance [25] trains an explicit classifier to discriminate between the generated samples and noisy training samples at each noise level and uses its log-gradient to guide the sampling. Our analysis is not applicable in this situation; in CFG the task is implicit and the distinction is between  $p_1$  and  $p_0$ .

lack of training — at least to some degree — it makes sense to choose  $D_0$  so that it further exacerbates these aspects.

In practice, models that are trained separately or for a different number of iterations differ not only in accuracy of fit, but also in terms of random initialization, shuffling of the training data, etc. For guidance to be successful, the quality gap should be large enough to make the systematic spreading-out of the density outweigh these random effects.

**Study on synthetic degradations.** To validate our hypothesis that the two models must suffer from the same kind of degradations, we perform a controlled experiment using synthetic corruptions applied to a well-trained real-world image diffusion model. We create the main and guiding networks,  $D_1$  and  $D_0$ , by applying different degrees of a synthetic corruption to the base model. This construction allows us to use the untouched base model as grounding when measuring the FID effect of the various combinations of corruptions applied to  $D_1$  and  $D_0$ . We find that as long as the degradations are compatible, autoguidance largely undoes the damage caused by the corruptions:

- **Base model:** As the base model, we use EDM2-S trained on ImageNet-512 without dropout (FID = 2.56).
- **Dropout:** We construct  $D_1$  by applying 5% dropout to the base model in a post-hoc fashion (FID = 4.98), and  $D_0$  by applying 10% dropout to the base model (FID = 15.00). Applying autoguidance, we reach the best result (FID = 2.55) with  $w = 2.25$ , matching the base model’s FID.
- **Input noise:** We construct  $D_1$  by modifying the base model to add noise to the input images so that their noise level is increased by 10% (FID = 3.96). The  $\sigma$  conditioning input of the denoiser is adjusted accordingly. The guiding model  $D_0$  is constructed similarly, but with a noise level increase of 20% (FID = 9.73). Applying autoguidance, we reach the best result (FID = 2.56) with  $w = 2.00$ , again matching the base model’s FID.
- **Mismatched degradations:** If we corrupt  $D_1$  by dropout and  $D_0$  by input noise, or vice versa, guidance does not improve the results at all; in these cases, the best FID is obtained by setting  $w = 1$ , i.e., by disabling guidance and using the less corrupted  $D_1$  exclusively.

While this experiment corroborates our main hypothesis, we do not suggest that guiding with these synthetic degradations would be useful in practice. A realistic diffusion model will not suffer from these particular degradations, so creating a guiding model by introducing them would not yield consistent truncation towards the data manifold.

## 5 Results

Our primary evaluation is carried out using ImageNet (ILSVRC2012) [8] at two resolutions:  $512 \times 512$  and  $64 \times 64$ . For ImageNet-512 we use latent diffusion [42], while ImageNet-64 works directly on RGB pixels. We take the current state-of-the-art diffusion model EDM2 [24] as our baseline.<sup>3</sup> We use the EDM2-S and EDM2-XXL models with default sampling parameters: 32 deterministic steps with a 2<sup>nd</sup> order Heun sampler [23]. For most setups, a pre-trained model is publicly available, and in the remaining cases we train the models ourselves using the official implementation (Appendix B).

We use two degradations for the guiding model: shorter training time and reduced capacity compared to the main model. We obtain the best results by having both of these enabled. With EDM2-S, for example, we use an XS-sized guiding model that receives  $1/16^{\text{th}}$  of the training iterations of the main model. We ablate the relative importance of the degradations as well as the sensitivity to these specific choices in Section 5.1. As the EDM2 networks are known to be sensitive to the guidance weight and EMA length [24], we search the optimal values for each case using a grid search.

Table 1 shows that our method improves FID [13] and  $\text{FD}_{\text{DINOv2}}$  [51] considerably. Using the small model (EDM2-S) in ImageNet-512, our autoguidance improves FID from 2.56 to 1.34. This beats the 1.68 achieved by the concurrently proposed CFG + Guidance Interval [27], and is the best result reported for this dataset regardless of the model size. Using the largest model (EDM2-XXL) further improves the record to 1.25. The  $\text{FD}_{\text{DINOv2}}$  records are similarly improved, with the large model

<sup>3</sup><https://github.com/NVlabs/edm2>

Method		FID	$w$	EMA <sub>m</sub>	EMA <sub>g</sub>	FD <sub>DINOv2</sub>	$w$	EMA <sub>m</sub>	EMA <sub>g</sub>
512×512	EDM2-S [24]	2.56	–	0.130	–	68.64	–	0.190	–
	+ Classifier-free guidance [24]	2.23	1.40	0.025	0.025	52.32	1.90	0.085	0.085
	+ Guidance interval [27]	1.68	2.10	0.025	0.025	46.25	3.20	0.085	0.085
	+ Autoguidance (XS, T/16) Ours	<b>1.34</b>	2.10	0.070	0.125	<b>36.67</b>	2.45	0.120	0.165
	– Same EMA for both	1.53	1.95	0.050	0.050	40.81	2.25	0.115	0.115
	– Reduce training only	1.51	2.20	0.090	0.130	42.27	2.55	0.130	0.170
	– Reduce capacity only	2.13	1.80	0.120	0.160	59.89	1.90	0.140	0.085
	EDM2-XXL [24]	1.91	–	0.070	–	42.84	–	0.150	–
	+ Classifier-free guidance [24]	1.81	1.20	0.015	0.015	33.09	1.70	0.015	0.015
	+ Guidance interval [27]	1.40	2.00	0.015	0.015	29.16	2.90	0.015	0.015
+ Autoguidance (M, T/3.5) Ours	<b>1.25</b>	2.05	0.075	0.155	<b>24.18</b>	2.30	0.130	0.205	
EDM2-S, unconditional	11.67	–	0.145	–	209.53	–	0.170	–	
+ Autoguidance (XS, T/16) Ours	<b>3.86</b>	2.85	0.070	0.110	<b>90.39</b>	2.90	0.090	0.125	
64×64	RIN [21]	1.23	–	0.033	–	–	–	–	–
	EDM2-S [24]	1.58	–	0.075	–	58.52	–	0.160	–
	+ Classifier-free guidance	1.48	1.15	0.030	0.030	41.84	1.85	0.040	0.040
	+ Autoguidance (XS, T/8) Ours	<b>1.01</b>	1.70	0.045	0.110	<b>31.85</b>	2.20	0.105	0.175

Table 1: Results on ImageNet-512 and ImageNet-64. The parameters of autoguidance refer to the capacity and amount training received by the guiding model. The latter is given relative to the number of training images shown to the main model ( $T$ ). The columns EMA<sub>m</sub> and EMA<sub>g</sub> indicate the length parameter of the post-hoc EMA technique [24] for the main and guiding model, respectively.




Figure 3: Sensitivity w.r.t. autoguidance parameters, using EDM2-S on ImageNet-512. The shaded regions indicate the min/max FID over 3 evaluations. (a) Sweep over guidance weight  $w$  while keeping all other parameters unchanged. The curves correspond to how much the guiding model was trained relative to the number of images shown to the main model. (b) Sweep over guidance weight for different guiding model capacities. (c) Sweep over the two EMA length parameters for our best configuration, denoted with \* in (a) and (b).

lowering the record from 29.16 to 24.18. In ImageNet-64, the improvement is even larger; in this dataset, we set the new record FID and FD<sub>DINOv2</sub> of 1.01 and 31.85, respectively.

A particular strength of autoguidance is that it can be applied to unconditional models as well. While conditional ImageNet generation may be getting close to saturation, the unconditional results remain surprisingly poor. EDM2-S achieves a FID of 11.67 in the unconditional setting, indicating that practically none of the generated images are of presentable quality. Enabling autoguidance lowers the FID substantially to 3.86, and the improvement in FD<sub>DINOv2</sub> is similarly significant.

## 5.1 Ablations

Table 1 further shows that it is beneficial to allow independent EMA lengths for the main and guiding models. When both are forced to use the same EMA, FID worsens from 1.34 to 1.53 in ImageNet-512 (EDM2-S). We also measure the effect of each degradation (reduced training time, capacity) in isolation. If we set the guiding model to the same capacity as the main model and only train it for a



Figure 4: Example results for the *Tree frog*, *Palace*, *Mushroom*, *Castle* classes of ImageNet-512 using EDM2-S. Guidance weight increases to the right; rows are classifier-free guidance and our method.

shorter time, FID worsens to 1.51. If we instead train the reduced-capacity guiding model for as long as the main model, FID suffers a lot more, to 2.13. We can thus conclude that both degradations are beneficial and orthogonal, but a majority of the improvement comes from reduced training of the guiding model. Notably, all these ablations still outperform standard CFG in terms of FID.

Figure 3 probes the sensitivity to various hyperparameters. Our best result is obtained by training the guiding model  $1/16^{\text{th}}$  as much as the main model, in terms of images shown during training. Further halving the training budget is almost equally good, while doubling the amount of training starts to slowly compromise the results. The results are quite insensitive to the choice of the guidance weight. In terms of the capacity of the guiding model, one step smaller (XS for EDM2-S) gave the best result. Two steps smaller (XXS) was also better than no capacity reduction (S), but started to show excessive sensitivity to the guidance weight. The results are also sensitive to the EMA length, similarly to the original EDM2. Post-hoc EMA [24] allows us to search the optimal parameters at a feasible cost.

We also explored several other degradations for the guiding model but did not find them to be beneficial. First, we tried reducing the amount of training data used for the guiding model, but this did not seem to improve the results over the baseline. Second, applying guidance interval [27] on top of our method reduced its benefits to some extent, suggesting that autoguidance is helpful at all noise levels. Third, deriving the guiding model from the main model using synthetic degradations did not work at all, providing further evidence that the guiding model needs to exhibit the same kinds of degradations that the main model suffers from. Fourth, we found that if the main model had been quantized, e.g., to improve inference speed, quantizing it to an even lower precision did not yield a useful guiding model.

One limitation of autoguidance is the need to train a separate guiding model. That said, the additional training cost can be quite modest when using a smaller model and shorter training time for the guiding model. For example, the EDM2-M model trains approximately  $2.7\times$  as fast as EDM2-XXL per iteration, and we train it for  $1/3.5$  of iterations, so the additional cost is around +11%. For the EDM2-S/XS pair used in most of our experiments, the added training cost is only +3.6%.






Figure 5: Results for DeepFloyd IF [50] using the prompt “A blue jay standing on a large basket of rainbow macarons”. The rows correspond to guidance weights  $w \in \{1, 2, 3, 4\}$ . The leftmost column shows results for CFG and the rightmost for autoguidance (XL-sized model guided by M-sized one). The middle columns correspond to blending between the two. See Appendix A for more examples.

## 5.2 Qualitative results

Figure 4 shows examples of generated images for ImageNet-512. Both CFG and our method tend to improve the perceptual quality of images, guiding the results towards clearer realizations as the guidance weight increases. However, CFG seems to have a tendency to head towards a more limited number of canonical images [27] per class, while our method produces a wider gamut of image compositions. An example is the atypical image of a *Palace* at  $w = 1$ , which CFG converts to a somewhat idealized depiction as  $w$  increases. Sometimes the unguided sample contains incompatible elements of multiple possible images, such as the *Castle* image, which includes a rough sketch of two or three castles of unrelated styles. In this instance, CFG apparently struggles to decide what to do, whereas our method first builds the large red element into a castle, and with increased guidance focuses on the red foreground object. A higher number of possible output images is consistent with a lower FID, implying better coverage of the training data.

In order to study our method in the context of large-scale image generators, we apply it to DeepFloyd IF [50]. We choose this baseline because multiple differently-sized models are publicly available. Ideally we could have also used an earlier snapshot as the guiding model, but those were not available. DeepFloyd IF generates images as a cascade of three diffusion models: a base model and two super-resolution stages. We apply our method to the base model only, while the subsequent stages always use CFG. Figure 5 demonstrates the effect of CFG, our method, and their various combinations. To combine autoguidance with CFG, we extend Equation 3 to cover multiple guiding models as proposed by Liu et al. [31] and distribute the total guidance weight among them using linear interpolation (see Appendix B.2 for details). While CFG improves the image quality significantly, it also simplifies the style and layout of the image towards a canonical depiction. Our method similarly improves the

image quality, but it better preserves the image’s style and visual complexity. We hope that using both guiding methods simultaneously will serve as a new, useful artistic control.

## 6 Discussion and Future work

We have shown that classifier-free guidance entangles several phenomena together, and that a different perspective together with simple practical changes opens up an entire new design space. In addition to removing the superfluous connection to conditioning, this enables significantly better results.

Potential directions for future work include formally proving the conditions that allow autoguidance to be beneficial, and deriving good rules of thumb for selecting the best guiding model. Our suggestion — an early snapshot of a smaller model — is easy to satisfy in principle, but these are not available for current large-scale image generators in practice. Such generators are also often trained in successive stages where the training data may change at some point, causing potential distribution shifts between snapshots that would violate our assumptions. Various modifications to guidance [1, 17, 18] can be seen as inducing degradations through perturbation of attention maps or denoiser inputs. Whether these approaches could provide additional benefits in our setup remains an open question. Autoguidance also bears conceptual similarity to contrastive decoding [29] used in large language models to reduce the repetitiveness of generations, and there may be opportunities for sharing observations between the two domains.

Recently, several studies [6, 11, 27, 44, 54] have reduced the downsides of CFG by making the guidance weight noise level-dependent. A key benefit from these schedules appears to be the suppression of CFG at high noise levels, where its image quality benefit is overshadowed by the undesirable reduction in variation that is caused by large differences in the content of the differently conditioned distributions. In contrast, autoguidance is not expected to suffer from this problem at high noise levels, as both models target the same distribution. So far we have compared autoguidance only with the interval method [27], which we did not find beneficial in combination. A further study on the various possible combinations, in terms of quantitative performance as well as artistic control, is a natural next step. It could also be interesting to further isolate the origin of the improvement using alternative metrics, such as precision and recall [28], Human Preference Score [55], or PickScore [26].

## Acknowledgments

We thank David Luebke, Janne Hellsten, Ming-Yu Liu, and Alex Keller for discussions and comments, and Tero Kuosmanen and Samuel Klenberg for maintaining our compute infrastructure.

## References

- [1] D. Ahn, H. Cho, J. Min, W. Jang, J. Kim, S. Kim, H. H. Park, K. H. Jin, and S. Kim. Self-rectifying diffusion sampling with perturbed-attention guidance. In *Proc. ECCV*, 2024.
- [2] F. Bao, S. Nie, K. Xue, Y. Cao, C. Li, H. Su, and J. Zhu. All are worth words: A ViT backbone for diffusion models. In *Proc. CVPR*, 2023.
- [3] C. M. Bishop. *Neural networks for pattern recognition*. Oxford University Press, USA, 1995.
- [4] A. Brock, J. Donahue, and K. Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *Proc. ICLR*, 2019.
- [5] B. C. A. Brown, A. L. Caterini, B. L. Ross, J. C. Cresswell, and G. Loaiza-Ganem. Verifying the union of manifolds hypothesis for image data. In *Proc. ICLR*, 2023.
- [6] H. Chang, H. Zhang, J. Barber, A. Maschinot, J. Lezama, L. Jiang, M.-H. Yang, K. Murphy, W. T. Freeman, M. Rubinstein, Y. Li, and D. Krishnan. Muse: Text-to-image generation via masked generative transformers. In *Proc. ICML*, 2023.
- [7] K. Crowson, S. A. Baumann, A. Birch, T. M. Abraham, D. Z. Kaplan, and E. Shippole. Scalable high-resolution pixel-space image synthesis with hourglass diffusion transformers. In *Proc. ICML*, 2024.
- [8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proc. CVPR*, 2009.
- [9] P. Dhariwal and A. Nichol. Diffusion models beat GANs on image synthesis. In *Proc. NeurIPS*, 2021.
- [10] S. Dieleman. Guidance: A cheat code for diffusion models. Blog post. <https://sander.ai/2022/05/26/guidance.html>, 2022.
- [11] S. Gao, P. Zhou, M.-M. Cheng, and S. Yan. MDTv2: Masked diffusion transformer is a strong image synthesizer. In *Proc. ICCV*, 2023.
- [12] D. Hendrycks and K. Gimpel. Gaussian error linear units (GELUs). *CoRR*, abs/1606.08415, 2016.

- [13] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Proc. NIPS*, 2017.
- [14] J. Ho, W. Chan, C. Saharia, J. Whang, R. Gao, A. Gritsenko, D. P. Kingma, B. Poole, M. Norouzi, D. J. Fleet, and T. Salimans. Imagen Video: High definition video generation with diffusion models. *CoRR*, abs/2210.02303, 2022.
- [15] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *Proc. NeurIPS*, 2020.
- [16] J. Ho and T. Salimans. Classifier-free diffusion guidance. In *Proc. NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021.
- [17] S. Hong. Smoothed energy guidance: Guiding diffusion models with reduced energy curvature of attention. In *Proc. NeurIPS*, 2024.
- [18] S. Hong, G. Lee, W. Jang, and S. Kim. Improving sample quality of diffusion models using self-attention guidance. In *Proc. ICCV*, 2023.
- [19] E. Hoogeboom, J. Heek, and T. Salimans. Simple diffusion: End-to-end diffusion for high resolution images. In *Proc. ICML*, 2023.
- [20] A. Hyvärinen. Estimation of non-normalized statistical models by score matching. *JMLR*, 6(24):695–709, 2005.
- [21] A. Jabri, D. J. Fleet, and T. Chen. Scalable adaptive computation for iterative generation. In *Proc. ICML*, 2023.
- [22] A. Jolicoeur-Martineau, K. Li, R. Piché-Taillefer, T. Kachman, and I. Mitliagkas. Gotta go fast when generating data with score-based models. *CoRR*, abs/2105.14080, 2021.
- [23] T. Karras, M. Aittala, T. Aila, and S. Laine. Elucidating the design space of diffusion-based generative models. In *Proc. NeurIPS*, 2022.
- [24] T. Karras, M. Aittala, J. Lehtinen, J. Hellsten, T. Aila, and S. Laine. Analyzing and improving the training dynamics of diffusion models. In *Proc. CVPR*, 2024.
- [25] D. Kim, Y. Kim, S. J. Kwon, W. Kang, and I.-C. Moon. Refining generative process with discriminator guidance in score-based diffusion models. In *Proc. ICML*, 2023.
- [26] Y. Kirstain, A. Polyak, U. Singer, S. Matiana, J. Penna, and O. Levy. Pick-a-Pic: An open dataset of user preferences for text-to-image generation. In *Proc. NeurIPS*, 2023.
- [27] T. Kynkäänniemi, M. Aittala, T. Karras, S. Laine, T. Aila, and J. Lehtinen. Applying guidance in a limited interval improves sample and distribution quality in diffusion models. In *Proc. NeurIPS*, 2024.
- [28] T. Kynkäänniemi, T. Karras, S. Laine, J. Lehtinen, and T. Aila. Improved precision and recall metric for assessing generative models. In *Proc. NeurIPS*, 2019.
- [29] X. L. Li, A. Holtzman, D. Fried, P. Liang, J. Eisner, T. Hashimoto, L. Zettlemoyer, and M. Lewis. Contrastive decoding: Open-ended text generation as optimization. In *Proc. ACL*, 2023.
- [30] L. Liu, Y. Ren, Z. Lin, and Z. Zhao. Pseudo numerical methods for diffusion models on manifolds. In *Proc. ICLR*, 2022.
- [31] N. Liu, S. Li, Y. Du, A. Torralba, and J. B. Tenenbaum. Compositional visual generation with composable diffusion models. In *Proc. ECCV*, 2022.
- [32] S. Lyu. Interpretation and generalization of score matching. In *Proc. UAI*, 2009.
- [33] M. Marchesi. Megapixel size image creation using generative adversarial networks. *CoRR*, abs/1706.00082, 2017.
- [34] P. Mishkin, L. Ahmad, M. Brundage, G. Krueger, and G. Sastry. DALL·E 2 preview – risks and limitations. *OpenAI*, 2022.
- [35] A. Nichol and P. Dhariwal. Improved denoising diffusion probabilistic models. In *Proc. ICML*, 2021.
- [36] A. Nichol, P. Dhariwal, A. Ramesh, P. Shyam, P. Mishkin, B. McGrew, I. Sutskever, and M. Chen. GLIDE: Towards photorealistic image generation and editing with text-guided diffusion models. In *Proc. ICML*, 2022.
- [37] M. Oquab, T. Darcet, T. Moutakanni, H. V. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby, M. Assran, N. Ballas, W. Galuba, R. Howes, P.-Y. Huang, S.-W. Li, I. Misra, M. Rabbat, V. Sharma, G. Synnaeve, H. Xu, H. Jegou, J. Mairal, P. Labatut, A. Joulin, and P. Bojanowski. DINOv2: Learning robust visual features without supervision. *TMLR*, 2024.
- [38] W. Peebles and S. Xie. Scalable diffusion models with transformers. In *Proc. ICCV*, 2023.
- [39] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall. DreamFusion: Text-to-3D using 2D diffusion. In *Proc. ICLR*, 2023.
- [40] P. Pope, C. Zhu, A. Abdelkader, M. Goldblum, and T. Goldstein. The intrinsic dimension of images and its impact on learning. In *Proc. ICLR*, 2021.
- [41] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with CLIP latents. *CoRR*, abs/2204.06125, 2022.
- [42] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proc. CVPR*, 2022.
- [43] A. Sabour, S. Fidler, and K. Kreis. Align your steps: Optimizing sampling schedules in diffusion models. In *Proc. ICML*, 2024.

- [44] S. Sadat, J. Buhmann, D. Bradley, O. Hilliges, and R. M. Weber. CADs: Unleashing the diversity of diffusion models through condition-annealed sampling. In *Proc. ICLR*, 2024.
- [45] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. S. Ghasemipour, B. K. Ayan, S. S. Mahdavi, R. G. Lopes, T. Salimans, J. Ho, D. J. Fleet, and M. Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *Proc. NeurIPS*, 2022.
- [46] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proc. ICML*, 2015.
- [47] J. Song, C. Meng, and S. Ermon. Denoising diffusion implicit models. In *Proc. ICLR*, 2021.
- [48] Y. Song and S. Ermon. Generative modeling by estimating gradients of the data distribution. In *Proc. NeurIPS*, 2019.
- [49] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations. In *Proc. ICLR*, 2021.
- [50] Stability AI. DeepFloyd IF. GitHub repository. <https://github.com/deep-floyd/IF>, 2023.
- [51] G. Stein, J. C. Cresswell, R. Hosseinzadeh, Y. Sui, B. L. Ross, V. Villicroze, Z. Liu, A. L. Caterini, J. E. T. Taylor, and G. Loaiza-Ganem. Exposing flaws of generative model evaluation metrics and their unfair treatment of diffusion models. In *Proc. NeurIPS*, 2023.
- [52] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the Inception architecture for computer vision. In *Proc. CVPR*, 2016.
- [53] P. Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011.
- [54] X. Wang, N. Dufour, N. Andreou, M.-P. Cani, V. F. Abrevaya, D. Picard, and V. Kalogeiton. Analysis of classifier-free guidance weight schedulers. *TMLR*, 2024.
- [55] X. Wu, Y. Hao, K. Sun, Y. Chen, F. Zhu, R. Zhao, and H. Li. Human Preference Score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *CoRR*, abs/2306.09341, 2023.
- [56] L. Zhang, A. Rao, and M. Agrawala. Adding conditional control to text-to-image diffusion models. In *Proc. ICCV*, 2023.
- [57] Q. Zhang and Y. Chen. Fast sampling of diffusion models with exponential integrator. In *Proc. ICLR*, 2023.
- [58] C. Zheng and Y. Lan. Characteristic guidance: Non-linear correction for diffusion model at large guidance scale. In *Proc. ICML*, 2024.

# Appendices

## A Additional results

Figure 6 shows additional results using DeepFloyd IF, similar to Figure 5.



Figure 6: Additional results for DeepFloyd IF [50]. The rows correspond to guidance weights  $w \in \{1, 2, 3, 4\}$ . CFG and our method (XL-sized model guided by M-sized one) on the leftmost and rightmost column, respectively. The middle columns correspond to blending between the two.




Figure 7: Sweep over guidance weight  $w$  using EDM2-S on ImageNet-512. The optimal EMA length was searched separately for the three methods and two metrics (FID and  $FD_{DINOv2}$ ).




Figure 8: Resulting variation for CFG and our method in *Tree frog*, *Minibus*, *Mushroom* classes of ImageNet-512 using EDM2-S ( $FD_{DINOv2}$ -optimized). In this image, we have exaggerated the amount of guidance ( $w = 4$ ) to make its effect on variation more clearly visible. This causes excessive saturation and other image artifacts, but clearly shows that CFG steers towards canonical templates, while our method preserves much greater variation.

Figure 7 plots FID and  $FD_{DINOv2}$  as functions of guidance weight, showing that autoguidance is less sensitive to the exact choice of  $w$  than CFG. Figure 8 shows example grids that demonstrate that autoguidance retains much higher variation than CFG. Figure 9 provides additional visualizations from our toy example, showing how the implied densities evolve during inference with CFG and autoguidance.

## B Implementation details

We performed our main experiments on top of the publicly available EDM2 [24] codebase<sup>4</sup> using NVIDIA A100 GPUs, Python 3.11.7, PyTorch 2.2.0, CUDA 11.8, and CuDNN 8.9.7. Since our method only involves using a different guiding model during sampling, we were able to perform all

<sup>4</sup><https://github.com/NVlabs/edm2>

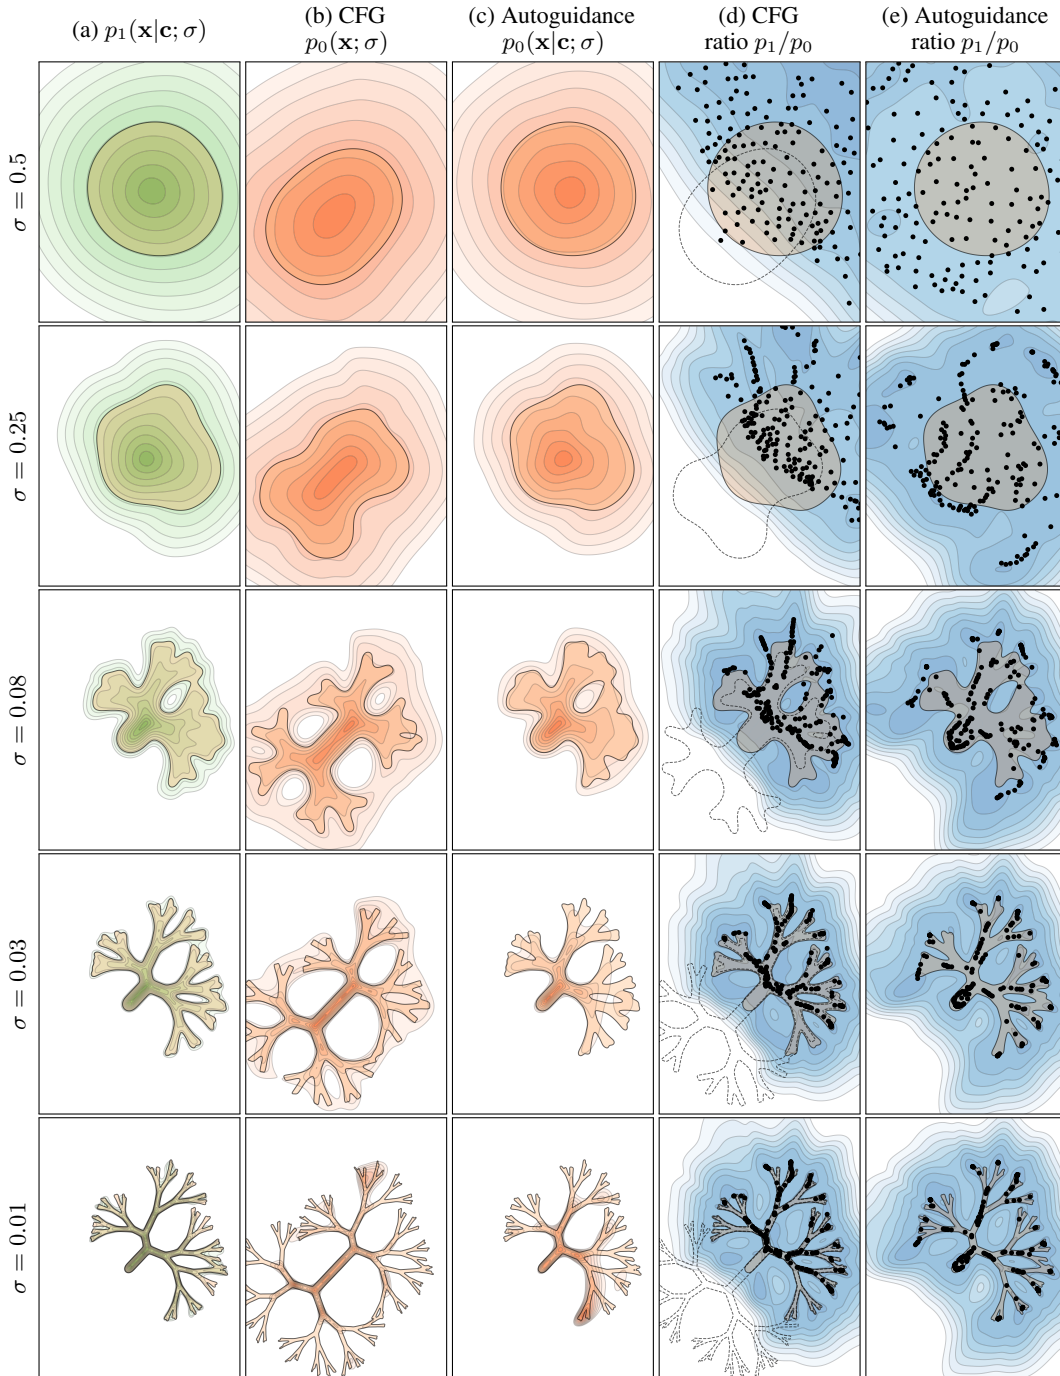


Figure 9: Progression of implied learned densities during sampling over various  $\sigma$  in a setup similar to Figure 2. Contours of the corresponding ground truth distributions are also shown. **(a)** Main model density  $p_1(\mathbf{x}|\mathbf{c}; \sigma)$ . **(b)** Unconditional guiding model density  $p_0(\mathbf{x}; \sigma)$  in CFG. **(c)** Conditional guiding model density  $p_0(\mathbf{x}|\mathbf{c}; \sigma)$  in autoguidance. **(d)** With CFG, guidance towards higher ratio  $p_1/p_0$  pushes samples towards top right, especially at high  $\sigma$  (top rows). **(e)** With autoguidance, this anomalous effect does not occur and samples cover the entire class  $\mathbf{c}$ .

measurements using the existing command-line scripts. We only had to modify one line of code in the sampling loop to pass in the class label to the guiding network in addition to the main network. Algorithm 1 demonstrates the steps needed to reproduce one of our results from Table 1; the rest can be reproduced by repeating the same steps with different models and hyperparameters as indicated in the table. For the four cases where a pre-trained model was unavailable, we trained the models ourselves as detailed in Algorithm 2.

### B.1 Hyperparameter search

We optimized the autoguidance parameters for each configuration in Table 1 and Figure 3 using automated grid search. We performed the search separately for FID and  $FD_{DINOv2}$  across the space of five parameters:

- **Model capacity:** All available model capacities (EDM2-XXS, EDM2-XS, etc.) up to and including the capacity of the main model.
- **Training time:** Number of training images in powers of two. We found that the sweet spot was always within the range  $\{2^{26}, 2^{27}, 2^{28}, 2^{29}\}$  — there was no need to go for lower or higher values.
- **Guidance weight:** All values within  $[1.00, 3.50]$  at regular intervals of 0.05.
- **EMA lengths:** All values within  $[0.010, 0.250]$  at regular intervals of 0.005. We treated the EMA length of the main model and the guiding model as two separate parameters.

In order to reduce the computational workload, we pruned the search space by considering only a local neighborhood of parameters around the best FID or  $FD_{DINOv2}$  found thus far. Whenever the result improved, we placed a new local grid around the corresponding parameters, resulting in gradual convergence towards the global optimum. Once we reached the optimum, we further re-evaluated the nearby parameter choices two more times to account for the effect of random noise. As such, each result reported in Table 1 and Figure 3 represents the best of three evaluations. The typical range of random variation is indicated by the shaded regions in Figure 3.

In total, the grid search resulted in roughly 30,000 metric evaluations across all of our configurations. In each metric evaluation, the main cost comes from generating 50,000 random images, which takes around 30–60 minutes using eight A100 GPUs and consumes approximately 2–5 kWh of energy, depending on model size. The overall energy consumption of our entire project was thus in the ballpark of 60–150 MWh.

### B.2 DeepFloyd IF experiments

In order to apply CFG and autoguidance simultaneously, we extend Equation 3 to cover multiple guiding models as proposed by Liu et al. [31]. In this case, we have three models: the main model  $D_m$ , an unconditional model  $D_u$ , and a reduced-capacity conditional model  $D_c$ . The guided denoising result is then defined as

$$D_w(\mathbf{x}; \sigma, \mathbf{c}) := D_m(\mathbf{x}; \sigma, \mathbf{c}) + \sum_{i \in \{u, c\}} (w_i - 1)(D_m(\mathbf{x}; \sigma, \mathbf{c}) - D_i(\mathbf{x}; \sigma, \mathbf{c})), \quad (6)$$

where  $w_u$  and  $w_c$  correspond to the guidance weights for CFG and autoguidance, respectively. To interpolate between the two methods, we further define  $w_u := (1 - \alpha)(w - 1) + 1$  and  $w_c := \alpha(w - 1) + 1$ , where  $w$  indicates the desired total amount of guidance and  $\alpha \in [0, 1]$  is a linear interpolation factor.

DeepFloyd IF [50] uses a three-stage cascade: a base model followed by two super-resolution stages. We apply our method only to the base model, while the super-resolution stages always use CFG with their default weights (4 and 9). We use their Stochastic DDPM sampler with default settings: 100, 75, 50 steps for the base model and subsequent super-resolution stages, respectively. Dynamic thresholding [45] is used for all stages.

## C Details of the 2D toy example

In this section, we describe the construction of the 2D toy dataset used in the analysis of Section 3, as well as the associated model architecture, training setup, and sampling parameters. The related code is available at <https://github.com/NVlabs/edm2>



---

**Algorithm 1** Reproducing our FID result for the “Autoguidance (XS,  $T/16$ )” row in Table 1.

---

```
1 # Download EDM2 source code.
2 git clone https://github.com/NVlabs/edm2.git
3 cd edm2
4 git checkout 32ecad3
5
6 # Patch the sampler so that class labels are passed in to the guiding network.
7 sed -i 's/gnet(x, t)/gnet(x, t, labels)/g' generate_images.py
8
9 # Download the necessary EDM2 models.
10 rclone copy --progress --http-url https://nvlabs-fi-cdn.nvidia.com/edm2 \
11     :http:raw-snapshots/edm2-img512-s/ raw-snapshots/edm2-img512-s/
12 rclone copy --progress --http-url https://nvlabs-fi-cdn.nvidia.com/edm2 \
13     :http:raw-snapshots/edm2-img512-xs/ raw-snapshots/edm2-img512-xs/
14
15 # Reconstruct the corresponding post-hoc EMA models.
16 python reconstruct_phema.py --indir=raw-snapshots/edm2-img512-s --outdir=autoguidance/phema \
17     --outprefix=img512-s --outking=2147483 --outstd=0.070
18 python reconstruct_phema.py --indir=raw-snapshots/edm2-img512-xs --outdir=autoguidance/phema \
19     --outprefix=img512-xs --outking=134217 --outstd=0.125
20
21 # Generate 50,000 images using 8 GPUs.
22 torchrun --standalone --nproc_per_node=8 generate_images.py --seeds=0-49999 --subdirs \
23     --outdir=autoguidance/images/img512-s --net=autoguidance/phema/img512-s-2147483-0.070.pkl \
24     --gnet=autoguidance/phema/img512-xs-0134217-0.125.pkl --guidance=2.10
25
26 # Calculate FID.
27 python calculate_metrics.py calc --images=autoguidance/images/img512-s \
28     --ref=https://nvlabs-fi-cdn.nvidia.com/edm2/dataset-refs/img512.pkl
```

---

**Dataset.** For each of the two classes  $\mathbf{c}$ , we model the fractal-like data distribution as a mixture of Gaussians  $\mathcal{M}_{\mathbf{c}} = (\{\phi_i\}, \{\mu_i\}, \{\Sigma_i\})$ , where  $\phi_i$ ,  $\mu_i$ , and  $\Sigma_i$  represent the weight, mean, and  $2 \times 2$  covariance matrix of each component  $i$ , respectively. This lets us calculate the ground truth scores and probability densities analytically and, consequently, to visualize them without making any additional assumptions. The probability density for a given class is given by

$$p_{\text{data}}(\mathbf{x}|\mathbf{c}) = \sum_{i \in \mathcal{M}_{\mathbf{c}}} \phi_i \mathcal{N}(\mathbf{x}; \mu_i, \Sigma_i), \quad \text{where} \quad (7)$$

$$\mathcal{N}(\mathbf{x}; \mu, \Sigma) = \frac{1}{\sqrt{(2\pi)^2 \det(\Sigma)}} \exp\left(-\frac{1}{2}(\mathbf{x} - \mu)^\top \Sigma^{-1}(\mathbf{x} - \mu)\right). \quad (8)$$

Applying heat diffusion to  $p_{\text{data}}(\mathbf{x}|\mathbf{c})$ , we obtain a sequence of increasingly smoothed densities  $p(\mathbf{x}|\mathbf{c}; \sigma)$  parameterized by noise level  $\sigma$ :

$$p(\mathbf{x}|\mathbf{c}; \sigma) = \sum_{i \in \mathcal{M}_{\mathbf{c}}} \phi_i \mathcal{N}(\mathbf{x}; \mu_i, \Sigma_{i,\sigma}^*), \quad \text{where} \quad \Sigma_{i,\sigma}^* = \Sigma_i + \sigma^2 \mathbf{I}. \quad (9)$$

The score function of  $p(\mathbf{x}|\mathbf{c}; \sigma)$  is then given by

$$\nabla_{\mathbf{x}} \log p(\mathbf{x}|\mathbf{c}; \sigma) = \frac{\sum_{i \in \mathcal{M}_{\mathbf{c}}} \phi_i \mathcal{N}(\mathbf{x}; \mu_i, \Sigma_{i,\sigma}^*) (\Sigma_{i,\sigma}^*)^{-1} (\mu_i - \mathbf{x})}{\sum_{i \in \mathcal{M}_{\mathbf{c}}} \phi_i \mathcal{N}(\mathbf{x}; \mu_i, \Sigma_{i,\sigma}^*)}. \quad (10)$$

We construct  $\mathcal{M}_{\mathbf{c}}$  to represent a thin tree-like structure by starting with one main “branch” and recursively subdividing it into smaller ones. Each branch is represented by 8 anisotropic Gaussian components and the subdivision is performed 6 times, decaying  $\phi$  after each subdivision and slightly randomizing the lengths and orientations of the two resulting sub-branches. This yields  $127 \times 8 = 1016$  components per class and  $1016 \times 2 = 2032$  components in total. We define the coordinate system so that the mean and standard deviation of  $p_{\text{data}}$ , marginalized over  $\mathbf{c}$ , are equal to 0 and  $\sigma_{\text{data}} = 0.5$  along each axis, respectively, matching the recommendations by Karras et al. [23].

---

**Algorithm 2** Training the additional EDM2 models needed in Section 5.
 

---

```

1 # Unconditional EDM2-S with ImageNet-512, used as a main model in Table 1.
2 torchrun --nnodes=4 --nproc_per_node=8 train_edm2.py \
3   --outdir=autoguidance/train/img512-s-uncond --data=datasets/img512-sd.zip --cond=0 \
4   --preset=edm2-img512-s --duration=2048Mi
5
6 # Unconditional EDM2-XS with ImageNet-64, used as a guiding model in Table 1.
7 torchrun --nnodes=4 --nproc_per_node=8 train_edm2.py \
8   --outdir=autoguidance/train/img64-xs-uncond --data=datasets/img64.zip --cond=0 \
9   --preset=edm2-img64-s --channels=128 --lr=0.0120 --duration=2048Mi
10
11 # Conditional EDM2-XS with ImageNet-64, used as a guiding model in Table 1.
12 torchrun --nnodes=4 --nproc_per_node=8 train_edm2.py \
13   --outdir=autoguidance/train/img64-xs --data=datasets/img64.zip --cond=1 \
14   --preset=edm2-img64-s --channels=128 --lr=0.0120 --duration=512Mi
15
16 # Conditional EDM2-XXS with ImageNet-512, used as a guiding model in Figure 3b.
17 torchrun --nnodes=4 --nproc_per_node=8 train_edm2.py \
18   --outdir=autoguidance/train/img512-xxs --data=datasets/img512-sd.zip --cond=1 \
19   --preset=edm2-img512-xs --channels=64 --lr=0.0170 --duration=512Mi

```

---

**Models.** We implement the denoiser models  $D_0$  and  $D_1$  as simple multi-layer perceptrons, utilizing the magnitude-preserving design principles from EDM2 [24]. To be able to visualize the implied probability densities in Figure 2, we design the model interface so that for a given noisy sample, each model outputs a single scalar representing the logarithm of the corresponding unnormalized probability density, as opposed to directly outputting the denoised sample or the score vector. Concretely, let us denote the output of a given model by  $G_\theta(\mathbf{x}; \sigma, \mathbf{c})$ . The corresponding normalized probability density is then given by

$$p_\theta(\mathbf{x}|\mathbf{c}; \sigma) = \exp(G_\theta(\mathbf{x}; \sigma, \mathbf{c})) / \int \exp(G_\theta(\mathbf{x}; \sigma, \mathbf{c})) d\mathbf{x}. \quad (11)$$

By virtue of defining  $G_\theta$  this way, we can derive the score vector, and by extension, the denoised sample, from  $G_\theta$  through automatic differentiation:

$$\nabla_{\mathbf{x}} \log p_\theta(\mathbf{x}|\mathbf{c}; \sigma) = \nabla_{\mathbf{x}} G_\theta(\mathbf{x}; \sigma, \mathbf{c}) \quad (12)$$

$$D_\theta(\mathbf{x}; \sigma, \mathbf{c}) = \mathbf{x} + \sigma^2 \nabla_{\mathbf{x}} G_\theta(\mathbf{x}; \sigma, \mathbf{c}). \quad (13)$$

Besides Equation 12, we also tried out the alternative formulations where the model outputs the score vector or the denoised sample directly. The results produced by all these variants were qualitatively more or less identical; we chose to go with the formulation above purely for convenience.

To connect the above definition of  $G_\theta$  to the raw network layers, we apply preconditioning using the same general principles as in EDM [23]. Denoting the function represented by the raw network layers as  $F_\theta$ , we define  $G_\theta$  as

$$G_\theta(\mathbf{x}; \sigma, \mathbf{c}) = -\frac{1}{2} \|\mathbf{x}^*\|_2^2 - \frac{g_\theta}{\sigma n} \sum_{i=1}^n F_{\theta,i}(\mathbf{x}^*; \frac{1}{4} \log \sigma, \mathbf{c})^2, \quad \text{where } \mathbf{x}^* = \frac{\mathbf{x}}{\sqrt{\sigma^2 + \sigma_{\text{data}}^2}} \quad (14)$$

and the sum is taken over the  $n$  output features of  $F_\theta$ . We scale the output of  $F_\theta$  by a learned scaling factor  $g_\theta$  that we initialize to zero.

The goal of Equation 14 is to satisfy the following three requirements:

- The input of  $F_\theta$  should have zero mean and unit magnitude. This is achieved through the division by  $\sqrt{\sigma^2 + \sigma_{\text{data}}^2}$ .
- After initialization,  $G_\theta$  should represent the best possible first-order approximation of the correct solution. This is achieved through the  $-\frac{1}{2} \|\mathbf{x}^*\|_2^2$  term, as well as the fact that  $g_\theta = 0$  after initialization.
- After training,  $\sqrt{g_\theta} \cdot F_\theta$  should have approximately unit magnitude. This is achieved through the division by  $\sigma n$ .

In practice, we use an MLP with one input layer and four hidden layers, interspersed with SiLU [12] activation functions and implemented using the magnitude-preserving primitives from EDM2 [24]. The input is a 4-dimensional vector  $[\mathbf{x}_x^*; \mathbf{x}_y^*; \frac{1}{4} \log \sigma; 1]$  and the output of each hidden layer has  $n$  features, where  $n = 64$  for  $D_1$  and 32 for  $D_0$ .

**Training.** Given that we have the exact score function of the ground truth distribution readily available (Equation 9), we train the models using exact score matching [20] for simplicity and increased robustness. We thus define the loss function as

$$\mathcal{L}(\theta) = \mathbb{E}_{\sigma \sim p_{\text{train}}, \mathbf{x} \sim p(\mathbf{x}; \sigma)} \sigma^2 \|\nabla_{\mathbf{x}} \log p_{\theta}(\mathbf{x}; \sigma) - \nabla_{\mathbf{x}} \log p(\mathbf{x}; \sigma)\|_2^2, \quad (15)$$

where  $\sigma \sim p_{\text{train}}$  is realized as  $\log(\sigma) \sim \mathcal{N}(P_{\text{mean}}, P_{\text{std}})$  [23]. As an alternative to exact score matching, we also experimented with the more commonly used denoising score matching, but did not observe any noticeable differences in model behavior or training dynamics.

We train  $D_1$  for 4096 iterations using a batch size of 4096 samples, and  $D_0$  for 512 iterations. We set  $P_{\text{mean}} = -2.3$  and  $P_{\text{std}} = 1.5$ , and use  $\alpha_{\text{ref}} / \sqrt{\max(t/t_{\text{ref}}, 1)}$  learning rate decay schedule with  $\alpha_{\text{ref}} = 0.01$  and  $t_{\text{ref}} = 512$  iterations, along with a power function EMA profile [24] with  $\sigma_{\text{rel}} = 0.010$ . Overall, the setup is robust with respect to the hyperparameters; the phenomena illustrated in Figures 1 and 2 remain unchanged across a wide range of parameter choices.

**Sampling.** We use the standard EDM sampler [23] with  $N = 32$  Heun steps (NFE = 63),  $\sigma_{\text{min}} = 0.002$ ,  $\sigma_{\text{max}} = 5$ , and  $\rho = 7$ . We chose the values of  $N$  and  $\sigma_{\text{max}}$  to be much higher than what is actually needed for this dataset in order to avoid potential discretization errors from affecting our conclusions. In Figure 1, we set  $w = 4$  for CFG and  $w = 3$  for autoguidance, and multiply the score vectors (Equation 12) by 1.40 for naive truncation. In Figure 2, we set  $w = 4$  and  $\sigma_{\text{mid}} = 0.03$ .

## D Broader societal impact

Generative modeling, including images and videos, has significant misuse potential. It can trigger negative consequences within the society in several ways. The primary concerns include various types of disinformation, but also the potential to amplify stereotypes and unwanted biases [34]. Our improvements to the sample quality can make the results even more believable, even when used for disinformation. That said, we do not unlock any novel uses of the technology.

## E Licenses

- EDM2 models [24]: Creative Commons BY-NC-SA 4.0 license
- DeepFloyd IF models [50]: Modified MIT license
- Stable Diffusion VAE model [42]: CreativeML Open RAIL++-M license
- InceptionV3 model [52]: Apache 2.0 license
- DINOv2 model [37]: Apache 2.0 license
- ImageNet dataset [8]: Custom non-commercial license

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Both the quantitative and qualitative claims are substantiated by our results.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Primarily covered in Section 6.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [No]

Justification: Our equations build on previously known results, and highlight their practical aspects through mostly readily verifiable algebraic manipulations. We do not explicitly present all details of the derivations, and assume that the previous results are sufficiently rigorously proven in the respective literature.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: Appendix B describes how publicly available code can be used to reproduce our results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

## 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Appendix B describes how publicly available code can be used to reproduce our results. The models that are not already publicly available, as well as the implementation of the 2D toy example, will be made available prior to the conference.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: In Section 5 and Appendix B.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Shaded regions in Figure 3.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

#### 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: In Appendix B.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

#### 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have reviewed the Code of Ethics, and do not find concerns beyond those relating to the general topic of generative modeling of images.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

#### 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: In Appendix D.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [No]

Justification: While our work falls into a high-risk category, the nature of our work (basically reusing pre-trained models) is such that we are not really in a position to introduce new safeguards.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: In Appendix E.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.



- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

### 13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: —

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: —

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: —

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.