

# Perception Prioritized Training of Diffusion Models

Jooyoung Choi<sup>1</sup>    Jungbeom Lee<sup>1</sup>    Chaehun Shin<sup>1</sup>    Sungwon Kim<sup>1</sup>  
 Hyunwoo Kim<sup>3</sup>    Sungroh Yoon<sup>1,2,\*</sup>

<sup>1</sup> Data Science and AI Laboratory, Seoul National University

<sup>2</sup> AIIS, ASRI, INMC, ISRC, and Interdisciplinary Program in AI, Seoul National University

<sup>3</sup> LG AI Research

## Abstract

*Diffusion models learn to restore noisy data, which is corrupted with different levels of noise, by optimizing the weighted sum of the corresponding loss terms, i.e., denoising score matching loss. In this paper, we show that restoring data corrupted with certain noise levels offers a proper pretext task for the model to learn rich visual concepts. We propose to prioritize such noise levels over other levels during training, by redesigning the weighting scheme of the objective function. We show that our simple redesign of the weighting scheme significantly improves the performance of diffusion models regardless of the datasets, architectures, and sampling strategies.*

## 1. Introduction

Diffusion models [14, 36], a recent family of generative models, have achieved remarkable image generation performance. Diffusion models have been rapidly studied, as they offer several desirable properties for image synthesis, including stable training, easy model scaling, and good distribution coverage [27]. Starting from Ho et al. [14], recent works [8, 27, 40] have shown that the diffusion models can render high-fidelity images comparable to those generated by generative adversarial networks (GANs) [12], especially in class-conditional settings, by relying on additional efforts such as classifier guidance [8] and cascaded models [32]. However, the unconditional generation of single models still has considerable room for improvement, and performance has not been explored for various high-resolution datasets (e.g., FFHQ [20], MetFaces [18]) where other families of generative models [3, 11, 20, 22, 41] mainly compete.

Starting from tractable noise distribution, a diffusion model generates images by progressively removing noise. To achieve this, a model learns the reverse of the predefined *diffusion process*, which sequentially corrupts the con-

tents of an image with various levels of noise. A model is trained by optimizing the sum of denoising score matching losses [43] for various noise levels [39], which aims to learn the recovery of clean images from corrupted images. Instead of using a simple sum of losses, Ho et al. [14] observed that their empirically obtained weighted sum of losses was more beneficial to sample quality. Their weighted objective is the current de facto standard objective for training diffusion models [8, 25, 27, 32, 40]. However, surprisingly, it remains unknown why this performs well or whether it is optimal for sample quality. To the best of our knowledge, the design of a better weighting scheme to achieve better sample quality has not yet been explored.

Given the success of diffusion models with the standard weighted objective, we aim to amplify this benefit by exploring a more appropriate weighting scheme for the objective function. However, designing a weighting scheme is difficult owing to two factors. First, there are thousands of noise levels; therefore, an exhaustive grid search is impossible. Second, it is not clear what information the model learns at each noise level during training, therefore hard to determine the priority of each level.

In this paper, we first investigate what a diffusion model learns at each noise level. Our key intuition is that the diffusion model learns rich visual concepts by solving *pretext tasks* for each level, which is to recover the image from corrupted images. At the noise level where the images are slightly corrupted, images are already available for perceptually rich content and thus, recovering images does not require prior knowledge of image contexts. For example, the model can recover noisy pixels from neighboring clean pixels. Therefore, the model learns imperceptible details, rather than high-level contexts. In contrast, when images are highly corrupted so that the contents are unrecognizable, the model learns perceptually recognizable contents to solve the given pretext task. Our observations motivate us to propose P2 (perception prioritized) weighting, which aims to prioritize solving the pretext task of more important noise levels. We assign higher weights to the loss at levels where

\*Correspondence to: Sungroh Yoon (sryoon@snu.ac.kr)

the model learns perceptually rich contents while minimal weights to which the model learns imperceptible details.

To validate the effectiveness of the proposed P2 weighting, we first compare diffusion models trained with previous standard weighting scheme and P2 weighting on various datasets. Models trained with our objective are consistently superior to the previous standard objective by large margins. Moreover, we show that diffusion models trained with our objective achieve state-of-the-art performance on CelebA-HQ [17] and Oxford-flowers [28] datasets, and comparable performance on FFHQ [20] among various types of generative models, including generative adversarial networks (GANs) [12]. We further analyze whether P2 weighting is effective to various model configurations and sampling steps. Our main contributions are as follows:

- We introduce a simple and effective weighting scheme of training objectives to encourage the model to learn rich visual concepts.
- We investigate how the diffusion models learn visual concepts from each noise level.
- We show consistent improvement of diffusion models across various datasets, model configurations, and sampling steps.

## 2. Background

### 2.1. Definitions

Diffusion models [14, 36] transform complex data distribution  $p_{data}(x)$  into simple noise distribution  $\mathcal{N}(0, \mathbf{I})$  and learn to recover data from noise. The *diffusion process* of diffusion models gradually corrupts data  $x_0$  with predefined noise scales  $0 < \beta_1, \beta_2, \dots, \beta_T < 1$ , indexed by time step  $t$ . Corrupted data  $x_1, \dots, x_T$  are sampled from data  $x_0 \sim p_{data}(x)$ , with a diffusion process, which is defined as Gaussian transition:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t\mathbf{I}). \quad (1)$$

Noisy data  $x_t$  can be sampled from  $x_0$  directly:

$$x_t = \sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon, \quad (2)$$

where  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$  and  $\alpha_t := \prod_{s=1}^t (1 - \beta_s)$ . We note that data  $x_0$ , noisy data  $x_1, \dots, x_T$ , and noise  $\epsilon$  are of the same dimensionality. To ensure  $p(x_T) \sim \mathcal{N}(0, \mathbf{I})$  and the reversibility of the diffusion process [36], one should set  $\beta_t$  to be small and  $\alpha_T$  to be near zero. To this end, Ho et al. [14] and Dhariwal et al. [8] use a linear noise schedule where  $\beta_t$  increases linearly from  $\beta_1$  to  $\beta_T$ . Nichol et al. [27] use a cosine schedule where  $\alpha_t$  resembles the cosine function.

Diffusion models generate data  $x_0$  with the learned *denoising process*  $p_\theta(x_{t-1}|x_t)$  which reverses the diffusion

process of Eq. (1). Starting from noise  $x_T \sim \mathcal{N}(0, \mathbf{I})$ , we iteratively subtract the noise predicted by noise predictor  $\epsilon_\theta$ :

$$x_{t-1} = \frac{1}{\sqrt{1 - \beta_t}}(x_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}}\epsilon_\theta(x_t, t)) + \sigma_t z, \quad (3)$$

where  $\sigma_t^2$  is a variance of the denoising process and  $z \sim \mathcal{N}(0, \mathbf{I})$ . Ho et al. [14] used  $\beta_t$  as  $\sigma_t^2$ .

Recent work Kingma et al. [23] simplified the noise schedules of diffusion models in terms of *signal-to-noise ratio* (SNR). SNR of corrupted data  $x_t$  is a ratio of squares of mean and variance from Eq. (2), which can be written as:

$$\text{SNR}(t) = \alpha_t / (1 - \alpha_t), \quad (4)$$

and thus the variance of noisy data  $x_t$  can be written in terms of SNR:  $\alpha_t = 1 - 1 / (1 + \text{SNR}(t))$ . We would like to note that  $\text{SNR}(t)$  is a monotonically decreasing function.

### 2.2. Training Objectives

The diffusion model is a type of variational auto-encoder (VAE); where the encoder is defined as a fixed *diffusion process* rather than a learnable neural network, and the decoder is defined as a learnable *denoising process* that generates data. Similar to VAE, we can train diffusion models by optimizing a variational lower bound (VLB), which is a sum of denoising score matching losses [43]:  $L_{vlb} = \sum_t L_t$ , where weights for each loss term are uniform. For each step  $t$ , denoising score matching loss  $L_t$  is a distance between two Gaussian distributions, which can be rewritten in terms of noise predictor  $\epsilon_\theta$  as:

$$\begin{aligned} L_t &= D_{KL}(q(x_{t-1}|x_t, x_0) || p_\theta(x_{t-1}|x_t)) \\ &= \mathbb{E}_{x_0, \epsilon} \left[ \frac{\beta_t}{(1 - \beta_t)(1 - \alpha_t)} ||\epsilon - \epsilon_\theta(x_t, t)||^2 \right]. \end{aligned} \quad (5)$$

Intuitively, we train a neural network  $\epsilon_\theta$  to predict the noise  $\epsilon$  added in noisy image  $x_t$  for given time step  $t$ .

Ho et al. [14] empirically observed that the following simplified objective is more beneficial to sample quality:

$$L_{simple} = \sum_t \mathbb{E}_{x_0, \epsilon} [||\epsilon - \epsilon_\theta(x_t, t)||^2]. \quad (6)$$

In terms of VLB, their objective is  $L_{simple} = \sum_t \lambda_t L_t$  with weighting scheme  $\lambda_t = (1 - \beta_t)(1 - \alpha_t) / \beta_t$ . In a continuous-time setting, this scheme can be expressed in terms of SNR:

$$\lambda_t = -1 / \log\text{-SNR}'(t) = -\text{SNR}(t) / \text{SNR}'(t), \quad (7)$$

where  $\text{SNR}'(t) = \frac{d\text{SNR}(t)}{dt}$ . See appendix for derivations.

While Ho et al. [14] use fixed values for the variance  $\sigma_t$ , Nichol et al. [27] propose to learn it with hybrid objective  $L_{hybird} = L_{simple} + cL_{vlb}$ , where  $c = 1e^{-3}$ . They




Figure 1. **Information removal of a diffusion process.** (Left) Perceptual distance of corrupted images as a function of signal-to-noise ratio (SNR). Distances are measured between two noisy images either corrupted from the same image (blue) or different images (orange). We averaged distances measured with 200 random triplets from CelebA-HQ. Perceptually recognizable contents are removed when SNR magnitude is between  $10^{-2}$  and  $10^0$ . (Right) Illustration of the diffusion process.

observed that learning  $\sigma_t$  enables reducing sampling steps while maintaining the generation performance. We inherit their hybrid objective for efficient sampling and modify  $L_{simple}$  to improve performance.

### 2.3. Evaluation Metrics

We use FID [13] and KID [2] for quantitative evaluations. FID is well-known to be analogous to human perception [13] and well-used as a default metric [8, 11, 18, 20, 30] for measuring generation performances. KID is a well-used metric to measure performance on small datasets [18, 19, 30]. However, since both metrics are sensitive to the preprocessing [30], we use a correctly implemented library [30]. We compute FID and KID between the generated samples and the entire training set. We measured final scores with 50k samples and conducted ablation studies with 10k samples for efficiency, following [8]. We denote them as FID-50k and FID-10k respectively.

## 3. Method

We first investigate what the model learns at each diffusion step in Sec. 3.1. Then, we propose our weighting scheme in Sec. 3.2. We provide discussions on how our weighting scheme is effective in Sec. 3.3.

### 3.1. Learning to Recover Signals from Noisy Data

Diffusion models learn visual concepts by solving *pre-text task* at each noise level, which is to recover signals from corrupted signals. More specifically, the model predicts the noise component  $\epsilon$  of a noisy image  $x_t$ , where the time step  $t$  is an index of the noise level. While the output of diffusion models is noise, other generative models (VAE, GAN) directly output images. Because noise does not contain any content or signals, it is difficult to understand how the noise predictions contribute to learning rich visual concepts. Such nature of diffusion models arises the following

question: *what information does the model learn at each step during training?*

**Investigating diffusion process.** We first investigate the predefined diffusion process to explore what the model can learn from each noise level. Let say we have two different clean images  $x_0, x'_0$  and three noisy images  $x_{tA}, x_{tB} \sim q(x_t|x_0), x'_t \sim q(x_t|x'_0)$ , where  $q$  is the diffusion process. In Fig. 1 (left), we measure perceptual distances (LPIPS [46]) in two cases: the distance between  $x_{tA}$  and  $x_{tB}$  (blue line), which share the same  $x_0$ , and the distance between  $x_{tA}$  and  $x'_t$  (orange line), which were synthesized from different images  $x_0$  and  $x'_0$ . We present the distances of the two cases as functions of the *signal-to-noise ratio* (SNR) introduced in Eq. (4), which characterizes the noise level at each step. To briefly review, SNR decreases through the diffusion process, as shown in Fig. 1 (right), and increases through the denoising process.

The early steps of the diffusion process have large SNRs, which indicates invisibly small noise; thus, noisy images  $x_t$  retain a large amount of contents from the clean image  $x_0$ . Therefore, in the early steps,  $x_{tA}$  and  $x_{tB}$  are perceptually similar, while  $x_{tA}$  and  $x'_t$  are perceptually different, as shown by the large SNR side in Fig. 1 (left). A model can recover signals without understanding holistic contexts, as perceptually rich signals are already prepared in the image. Thus the model will learn only imperceptible details by solving recovery tasks when SNR is large.

In contrast, the late steps have small SNRs, indicating a sufficiently large noise to remove the contents of  $x_0$ . Therefore, distances of both cases start to converge to a constant value, as the noisy images become difficult to recognize the high-level contents. It is shown in the small SNR side in Fig. 1 (left). Here, a model needs prior knowledge to recover signals because the noisy images lack recognizable content. We argue that the model will learn perceptually rich contents by solving recovery tasks when SNR is small.




Figure 2. **Stochastic reconstruction.** (Left) Illustration of reconstruction, where sample are obtained from full sampling chain. (Right) Reconstructions  $\hat{x}_0$  with input images  $x_0$  on the rightmost column and SNR of  $x_t$  on the bottom. Samples in the 1st, 2nd columns share only the coarse attributes (e.g., global color structure) with the input. The 3rd, 4th columns share perceptually discriminative contents with the input. 5th column are almost identical to the input, suggesting that the model learns imperceptible details when SNRs are large.

**Investigating a trained model.** We would like to verify the aforementioned discussions with a trained model. Given an input image  $x_0$ , we first perturb it to  $x_t$  using a diffusion process  $q(x_t|x_0)$  and reconstruct it with the learned denoising process  $p_\theta(\hat{x}_0|x_t)$ , as illustrated in Fig. 2 (left). When  $t$  is small, the reconstruction  $\hat{x}_0$  will be highly similar to the input  $x_0$  as the diffusion process removes a small amount of signals, while  $\hat{x}_0$  will share less content with  $x_0$  when  $t$  is large. In Fig. 2 (right), we compare  $x_0$  and  $\hat{x}_0$  among various  $t$  to show how each step contributes to the sample. Samples in the first two columns share only coarse features (e.g., global color scheme) with the input on the rightmost column, whereas samples in the third and fourth columns share perceptually discriminative contents. This suggests that the model learns coarse features when the SNR of step  $t$  is smaller than  $10^{-2}$  and the model learns the content when SNR is between  $10^{-2}$  and  $10^0$ . When the SNR is larger than  $10^0$  (fifth column), reconstructions are perceptually identical to the inputs, suggesting that the model learns imperceptible details that do not contribute to perceptually recognizable contents.

Based on the above observations, we hypothesize that diffusion models learn coarse features (e.g., global color structure) at steps of small SNRs ( $0-10^{-2}$ ), perceptually rich contents at medium SNRs ( $10^{-2}-10^0$ ), and remove remaining noise at large SNRs ( $10^0-10^4$ ). According to our hypothesis, we group noise levels into three stages, which we term *coarse*, *content*, and *clean-up* stages.

### 3.2. Perception Prioritized Weighting

In the previous section, we explored what the diffusion model learns from each step in terms of SNR. We discussed that the model learns coarse features (e.g., global color structure), perceptually rich contents, and to clean up the remaining noise at three groups of noise levels. We pointed out that the model learns imperceptible details at the clean-

up stage. In this section, we introduce Perception Prioritized (P2) weighting, a new weighting scheme for the training objective, which aims to prioritize learning from more important noise levels.

We opt to assign minimal weights to the unnecessary clean-up stage thereby assigning relatively higher weights to the rest. In particular, we aim to emphasize training on the content stage to encourage the model to learn perceptually rich contexts. To this end, we construct the following weighting scheme:

$$\lambda'_t = \frac{\lambda_t}{(k + \text{SNR}(t))^\gamma}, \quad (8)$$

where  $\lambda_t$  is the previous standard weighting scheme (Eq. (7)) and  $\gamma$  is a hyperparameter that controls the strength of down-weighting focus on learning imperceptible details.  $k$  is a hyperparameter that prevents exploding weights for extremely small SNRs and determines sharpness of the weighting scheme. While multiple designs are possible, we show that even the simplest choice (P2) outperforms the standard scheme  $\lambda_t$ . Our method is applicable to existing diffusion models by replacing  $\sum_t \lambda_t L_t$  with  $\sum_t \lambda'_t L_t$ .

In fact, our weighting scheme  $\lambda'_t$  is a generalization of the popularly used [8, 25, 27, 40] weighting scheme  $\lambda_t$  of Ho et al. [14] (Eq. (7)), where  $\lambda'_t$  arrives at  $\lambda_t$  when  $\gamma = 0$ . We refer to  $\lambda_t$  as the *baseline* herein.

### 3.3. Effectiveness of P2 Weighting

Prior works [14, 27] empirically suggest that the baseline objective  $\sum_t \lambda_t L_t$  offers a better inductive bias for sample quality than the VLB objective  $\sum_t L_t$ , which does not impose any inductive bias during training. Fig. 3 exhibits  $\lambda'_t$  and  $\lambda_t$  for both linear [14] and cosine [27] noise schedules, which are explained in Sec. 2.1, indicating that both weighting schemes focus training on the content stage the most and the cleaning stage the least. The success of the baseline






Figure 3. **Weighting schemes.** (Left) Signal-to-noise ratio (SNR) of linear and cosine noise schedules for reference. (Middle) Weights of our P2 weighting and the baseline with a cosine schedule. (Right) Weights of P2 weighting and the baseline with a linear schedule. Compared to the baseline, P2 weighting suppresses weights for large SNRs, where the model learns imperceptible details.

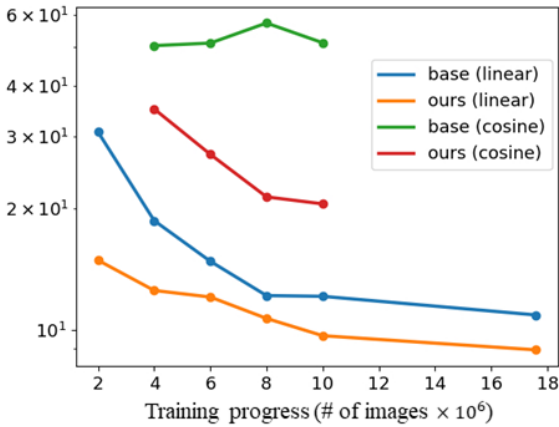


Figure 4. **Comparison of FID-10k through training on FFHQ.** P2 weighting consistently improves performance for both linear and cosine schedules. Training progress refers to the number of images seen by the model. Samples are generated with 250 steps.

weighting is in line with our previous hypothesis that models learn perceptually rich content by solving pretext tasks at the content stage.

However, despite the success of the baseline objective, we argue that the baseline objective still imposes an undesired focus on learning imperceptible details and prevents from learning perceptually rich content. Fig. 3 shows that our  $\lambda'_t$  further suppresses the weights for the cleaning stage, which relatively uplifts the weights for the coarse and the content stages. To visualize relative changes of weights, we exhibit normalized weighting schemes. Fig. 4 supports our method in that FID of the diffusion model trained with our weighting scheme ( $\gamma = 1$ ) beats the baseline for both linear and cosine schedules throughout the training.

Another notable result from Fig. 4 is that the cosine schedule is inferior to the linear schedule by a large margin, although our weighting scheme improves the FID by

a large gap. Eq. (5) indicates that the weighting scheme is closely related to the noise schedule. As shown in Fig. 3, the cosine schedule assigns smaller weights to the content stage compared with the linear schedule. We would like to note that designing weighting schemes and noise schedules are correlated but not equivalent, as the noise schedules affects both weights and MSE terms.

To summarize, our P2 weighting provides a good inductive bias for learning rich visual concepts, by uplifting weights at the coarse and the content stages, and suppressing weights at the clean-up stage.

### 3.4. Implementation

We set  $k$  as 1 for easy deployment, because  $1/(1 + \text{SNR}(t)) = 1 - \alpha_t$ , as discussed in Sec. 2.1. We set  $\gamma$  as either 0.5 and 1. We empirically observed that  $\gamma$  over 2 suffers noise artifacts in the sample because it assigns almost zero weight to the clean-up stage. We set  $T = 1000$  for all experiments. We implemented the proposed approach on top of ADM [8], which offers well-designed architecture and efficient sampling. We use lighter version of ADM through our experiments. See appendix for more details. Our codes and models will be available.

## 4. Experiment

We start by exhibiting the effectiveness of our new training objective over the baseline objective in Sec. 4.1. Then, we compare with prior literature of various types of generative models in Sec. 4.2. Finally, we conduct analysis studies to further support our method in Sec. 4.3. Samples generated with our models are shown in Fig. 5.

### 4.1. Comparison to the Baseline

**Quantitative comparison.** We trained diffusion models by optimizing training objectives with both baseline and our weighting scheme on FFHQ [20], AFHQ-dog [7], MetFaces [18], and CUB [44] datasets. These datasets contain




Figure 5. Samples generated by our models trained on several datasets (FFHQ, CelebA-HQ, MetFaces, AFHQ-Dogs, Oxford Flowers, CUB Bird) at  $256 \times 256$  resolution. See appendix for more samples.

Dataset	Step	FID-50k↓		KID-50k↓	
		Base	Ours	Base	Ours
FFHQ	1000	7.86	<b>6.92</b>	3.85	<b>3.46</b>
	500	8.41	<b>6.97</b>	4.48	<b>3.56</b>
CUB	1000	9.60	<b>6.95</b>	3.49	<b>2.38</b>
	250	10.26	<b>6.32</b>	4.06	<b>1.93</b>
AFHQ-D	1000	12.47	<b>11.55</b>	4.79	<b>4.10</b>
	250	12.95	<b>11.66</b>	5.25	<b>4.20</b>
Flowers	250	20.01	<b>17.29</b>	16.8	<b>14.8</b>
MetFaces	250	44.34	<b>36.80</b>	22.1	<b>17.6</b>

Table 1. **Quantitative comparison.** Diffusion models trained with our weighting scheme achieve consistent improvement over the baseline at various datasets and sampling steps, in terms of both FID and KID ( $\times 10^3$ ).

approximately 70k, 50k, 1k, and 12k images respectively. We resized and center-cropped data to  $256 \times 256$  pixels, following the pre-processing performed by ADM [8].

Tab. 1 shows the results. Our method consistently exhibits superior performance to the baseline in terms of FID and KID. The results suggest that our weighting scheme imposes a good inductive bias for training diffusion models, regardless of the dataset. Our method outperforms the baseline by a large margin especially on MetFaces, which contains only 1k images. Hence, we assume that wasting the model capacity on learning imperceptible details is very harmful when training with limited data.

**Qualitative comparison.** We observe that diffusion models trained with the baseline objective are likely to suffer color shift artifacts, as shown in Fig. 6. We assume that the baseline training objective unnecessarily focuses on the imperceptible details; therefore, it fails to learn global color

schemes properly. In contrast, our objective encourages the model to learn global and holistic concepts in a dataset.

## 4.2. Comparison to the Prior Literature

We compare diffusion models trained with our method to existing models on FFHQ [20], Oxford flowers [28], and CelebA-HQ [17] datasets, as shown in Tab. 2. We use  $256 \times 256$  resolutions for all datasets. We achieve state-of-the-art FIDs on the Oxford Flowers and CelebA-HQ datasets. While our models are trained with  $T = 1000$ , we already achieve state-of-the-art with reduced sampling steps; 250 and 500 steps respectively. On FFHQ, we achieve a superior result to most models except StyleGAN2 [21], whose architecture was carefully designed for FFHQ. We note that our method brought the diffusion model closer to the state-of-the-art, and scaling model architectures and sampling steps will further improve the performance.

## 4.3. Analysis

In this section, we analyze whether our weighting scheme is robust to the model configurations, number of sampling steps, and sampling schedules.

**Model configuration matters?** Previous experiments are conducted using our default model for fair comparisons. Here, we show that P2 weighting is effective regardless of the model configurations. Tab. 3 shows that our method achieves consistently superior performance to the baseline for various configurations. We investigated for following variations: replacing the BigGAN [3] residual block with the residual block of Ho et al. [14], removing self-attention at  $16 \times 16$ , using two BigGAN [3] residual blocks, and training our default model with a learning rate of  $2.5 \times 10^{-5}$ . Our default model contains a single BigGAN [3] residual block



Figure 6. **Qualitative comparison.** Uncurated samples generated during training. The number of images seen for training displayed on the top. We observed that the baseline suffer color shift problem at early iterations of the training (FFHQ) or even at the convergence (MetFaces). The baseline weighting scheme fails to focus on global coherence and waste model capacity on the imperceptible details.

Dataset	Method	Type	FID↓
FFHQ	BigGAN <sub>ICLR'19</sub> [3]	GAN	12.4
	UNet GAN <sub>CVPR'20</sub> [34]	GAN	10.9
	StyleGAN2 <sub>CVPR'20</sub> [21]	GAN	<b>3.73</b>
	NVAE <sub>NeurIPS'20</sub> [41]	VAE	26.02
	VDVAE <sub>ICLR'21</sub> [5]	VAE	33.5
	VQGAN <sub>CVPR'21</sub> [11]	GAN+AR	9.6
	D2C <sub>NeurIPS'21</sub> [35]	Diff	13.04
	Baseline (500 step)	Diff	8.41
	P2 (500 step)	Diff	6.97
	P2 (1000 step)	Diff	<u>6.92</u>
Oxford Flower	PGGAN <sub>ICLR'18</sub> [17]	GAN	64.40
	StyleGAN1 <sub>CVPR'19</sub> [20]	GAN	64.70
	MSG-GAN <sub>CVPR'20</sub> [16]	GAN	<u>19.60</u>
	Baseline (250step)	Diff	20.01
P2 (250step)	Diff	<b>17.29</b>	
CelebA -HQ	PGGAN <sub>ICLR'18</sub> [17]	GAN	8.03
	GLOW <sub>NeurIPS'18</sub> [22]	Flow	68.93
	ALAE <sub>CVPR'20</sub> [31]	GAN	19.21
	NVAE <sub>NeurIPS'20</sub> [41]	VAE	29.76
	VAEBM <sub>ICLR'21</sub> [45]	VAE+EM	20.38
	VQGAN <sub>CVPR'21</sub> [11]	GAN+AR	10.70
LSGM <sub>NeurIPS'21</sub> [42]	VAE+Diff	<u>7.22</u>	
P2 (500step)	Diff	<b>6.91</b>	

Table 2. **Comparison to prior literature.** FFHQ results reproduced from [11, 30, 35], Oxford Flower from [16], and CelebA-HQ from from [42]. Except for GANs, we achieve superior results.

and is trained with a learning rate  $2 \times 10^{-5}$ . Our weighting scheme consistently improves FID and KID by a large margin, across various model configurations. Our method is especially effective when the self-attention is removed ((c)), indicating that P2 encourages learning global dependency.

**Sampling step matters?** We trained our models on 1000 diffusion steps following the convention of previous studies. However, it requires more than 10 min to gener-

	FID-10k↓		KID-10k↓	
	Base	Ours	Base	Ours
(a)	46.80	41.93 (-4.87)	22.6	20.5 (-2.1)
(b)	47.62	47.37 (-0.25)	23.4	22.7 (-0.7)
(c)	49.56	43.09 (-6.47)	24.3	20.6 (-3.7)
(d)	45.45	42.06 (-3.39)	21.1	18.9 (-2.2)
(e)	46.34	39.51 (-6.83)	23.0	17.4 (-5.6)

Table 3. **Comparison among various model configurations.** (a) Our default configuration (b) No BigGAN block (c) Self-attention only at bottleneck ( $8 \times 8$  resolution) (d) Two residual blocks (e) Learning rate  $2.5e^{-5}$ . Models trained on MetFaces dataset for 800k images. Samples generated with 250 steps.

ate a high-resolution image with a modern GPU. Nichol et al. [27] have shown that their sampling strategy maintains performance even when reducing the sampling steps. They also observed that using the DDIM [37] sampler was effective when using 50 or fewer sampling steps.

Fig. 7 shows the FID scores of various sampling steps with models trained on the FFHQ. A model trained with our weighting scheme consistently outperforms the baseline by considerable margins. It should be noted that our weighting scheme consistently achieves better performance with half the number of sampling steps required by the baseline.

**Why not schedule sampling steps?** In addition to the consistent improvement across various sampling steps, we sweep over the sampling steps in Tab. 4. Sweeping sampling schedules slightly improves FID and KID but does not reach our improvement. Our method is more effective compared to scheduling sampling steps, as we improve the *model training*, which benefits predictions at all steps.

## 5. Related Work

### 5.1. Diffusion-Based Generative Models

Diffusion models [8, 14, 27, 36] and score-based models [39, 40] are two recent families of generative mod-



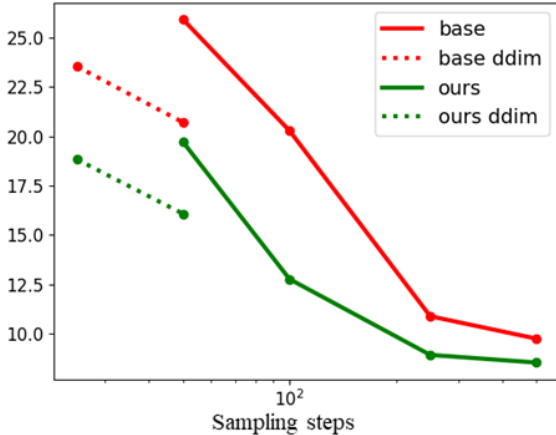


Figure 7. **Reducing sampling steps.** FID as a function of sampling steps. Our method is superior to the baseline regardless of the sampling steps. Samples generated following [27] and DDIM [37]. Models trained on FFHQ dataset.

Method	Schedule	FID-10k↓	KID-10k↓
Base	250 uniform	10.88	5.91
	130-60-60	10.62	5.14
	60-130-60	12.23	7.54
	500 uniform	9.74	4.48
Ours	250 uniform	<b>8.92</b>	<b>4.24</b>

Table 4. **Sweeping sampling schedule.** Schedules expressed as a sequence of integers which are numbers of steps assigned to one-third of the diffusion process. 130-60-60 implies consuming more steps near  $t = 0$ . Modifying the sampling schedule can slightly improve performance, but does not exceed our improvement.

els that generate data using a learned denoising process. Song et al. [40] showed that both families can be expressed with stochastic differentiable equations (SDE) with different noise schedules. We note that score-based models may enjoy different hyperparameters of P2 ( $\gamma$  and  $k$ ) as noise schedules are correlated with weighting schemes (Eq. (5)). Recent studies [8, 32, 40] have achieved remarkable improvements in sample quality. However, they rely on heavy architectures, long training and sampling steps [40], classifier guidance [8], and a cascade of multiple models [32]. In contrast, we improved the performance by simply redesigning the training objective without requiring heavy computations and additional models. Along with the success in the image domain, diffusion models have also shown effectiveness in speech synthesis [4, 25].

## 5.2. Advantages of Diffusion Models

Diffusion models have several advantages over other generative models. First, their sample quality is superior to likelihood-based methods such as autoregressive mod-

els [29, 33], flow models [9, 10], and variational autoencoders (VAEs) [24]. Second, because of the stable training, scaling and applying diffusion models to new domains and datasets is much easier than generative adversarial networks (GANs) [12], which rely on unstable adversarial training.

Moreover, pre-trained diffusion models are surprisingly easy to apply to downstream image synthesis tasks. Recent works [6, 26] have demonstrated that pre-trained diffusion models can easily adapt to image translation and image editing. Compared to GAN-based methods [1, 15, 47], they adapt a single diffusion model to various tasks without task-specific training and loss functions. They also show that diffusion models allow stochastic (one-to-many) generation in those tasks, while GAN-based methods suffer deterministic (one-to-one) generations [15].

## 5.3. Redesigning Training Objectives

Recent works [23, 38, 42] introduced new training objectives to achieve state-of-the-art likelihood. However, their objectives suffer from degradation of sample quality and training instability, therefore rely on importance sampling [38, 42] or sophisticated parameterization [23]. Because the likelihood focuses on fine-scale details, their objectives impede understanding global consistency and high-level concepts of images. For this reason, [42] use different weighting schemes for likelihood training and FID training. Our P2 weighting provides a good inductive bias for perceptually rich contents, allowing the model to achieve improved sample quality with stable training.

## 6. Conclusion

We proposed perception prioritized weighting, a new weighting scheme for the training objective of the diffusion models. We investigated how the model learns visual concepts at each noise level during training, and divided diffusion steps into three groups. We showed that even the simplest choice (P2) improves diffusion models across datasets, model configurations, and sampling steps. Designing a more sophisticated weighting scheme may further improve the performance, which we leave as future work. We believe that our method will open new opportunities to boost the performance of diffusion models.

**Acknowledgements:** This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) [NO.2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University)], LG AI Research, Samsung SDS, AIRS Company in Hyundai Motor and Kia through HMC/KIA-SNU AI Consortium Fund, and the BK21 FOUR program of the Education and Research Program for Future ICT Pioneers, Seoul National University in 2022.



## References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4432–4441, 2019. 8
- [2] Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*, 2018. 3
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 1, 6, 7
- [4] Nanxin Chen, Yu Zhang, Heiga Zen, Ron J Weiss, Mohammad Norouzi, and William Chan. Wavegrad: Estimating gradients for waveform generation. *arXiv preprint arXiv:2009.00713*, 2020. 8
- [5] Rewon Child. Very deep vaes generalize autoregressive models and can outperform them on images. *arXiv preprint arXiv:2011.10650*, 2020. 7
- [6] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938*, 2021. 8
- [7] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *CVPR*, 2020. 5
- [8] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis. *arXiv preprint arXiv:2105.05233*, 2021. 1, 2, 3, 4, 5, 6, 7, 8
- [9] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014. 8
- [10] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016. 8
- [11] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12873–12883, 2021. 1, 3, 7
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 1, 2, 8
- [13] 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. *Advances in neural information processing systems*, 30, 2017. 3
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *arXiv preprint arXiv:2006.11239*, 2020. 1, 2, 4, 6, 7
- [15] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 8
- [16] Animesh Karnewar and Oliver Wang. Msg-gan: Multi-scale gradients for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7799–7808, 2020. 7
- [17] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 2, 6, 7
- [18] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *NeurIPS*, 2020. 1, 3, 5
- [19] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *arXiv preprint arXiv:2106.12423*, 2021. 3
- [20] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 1, 2, 3, 5, 6, 7
- [21] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. 6, 7
- [22] Diederik P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *arXiv preprint arXiv:1807.03039*, 2018. 1, 7
- [23] Diederik P Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *arXiv preprint arXiv:2107.00630*, 2021. 2, 8
- [24] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 8
- [25] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*, 2020. 1, 4, 8
- [26] Chenlin Meng, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. 8
- [27] Alex Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. *arXiv preprint arXiv:2102.09672*, 2021. 1, 2, 4, 7, 8
- [28] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729. IEEE, 2008. 2, 6
- [29] Aaron van den Oord, Nal Kalchbrenner, Oriol Vinyals, Lasse Espeholt, Alex Graves, and Koray Kavukcuoglu. Conditional image generation with pixelcnn decoders. *arXiv preprint arXiv:1606.05328*, 2016. 8
- [30] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On buggy resizing libraries and surprising subtleties in fid calculation. *arXiv preprint arXiv:2104.11222*, 2021. 3, 7
- [31] Stanislav Pidhorskyi, Donald A Adjeroh, and Gianfranco Doretto. Adversarial latent autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14104–14113, 2020. 7
- [32] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *arXiv preprint arXiv:2104.07636*, 2021. 1, 8
- [33] Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P Kingma. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications.

- arXiv preprint arXiv:1701.05517*, 2017. 8
- [34] Edgar Schonfeld, Bernt Schiele, and Anna Khoreva. A u-net based discriminator for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8207–8216, 2020. 7
- [35] Abhishek Sinha, Jiaming Song, Chenlin Meng, and Stefano Ermon. D2c: Diffusion-denoising models for few-shot conditional generation. *arXiv preprint arXiv:2106.06819*, 2021. 7
- [36] Jascha Sohl-Dickstein, Eric A Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. *arXiv preprint arXiv:1503.03585*, 2015. 1, 2, 7
- [37] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 7, 8
- [38] Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. 2021. 8
- [39] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *NeurIPS*, 2019. 1, 7
- [40] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 1, 4, 7, 8
- [41] Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical variational autoencoder. *NeurIPS*, 2020. 1, 7
- [42] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. *arXiv preprint arXiv:2106.05931*, 2021. 7, 8
- [43] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 2011. 1, 2
- [44] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 5
- [45] Zhisheng Xiao, Karsten Kreis, Jan Kautz, and Arash Vahdat. Vaebm: A symbiosis between variational autoencoders and energy-based models. In *International Conference on Learning Representations*, 2020. 7
- [46] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 3
- [47] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 8