

# Blue noise for diffusion models

XINGCHANG HUANG, MPI Informatics, VIA Center, Germany  
CORENTIN SALAÜN, MPI Informatics, Germany  
CRISTINA VASCONCELOS, Google DeepMind, UK  
CHRISTIAN THEOBALT, MPI Informatics, VIA Center, Germany  
CENGİZ ÖZTIRELI, Google Research, University of Cambridge, UK  
GURPRIT SINGH, MPI Informatics, VIA Center, Germany

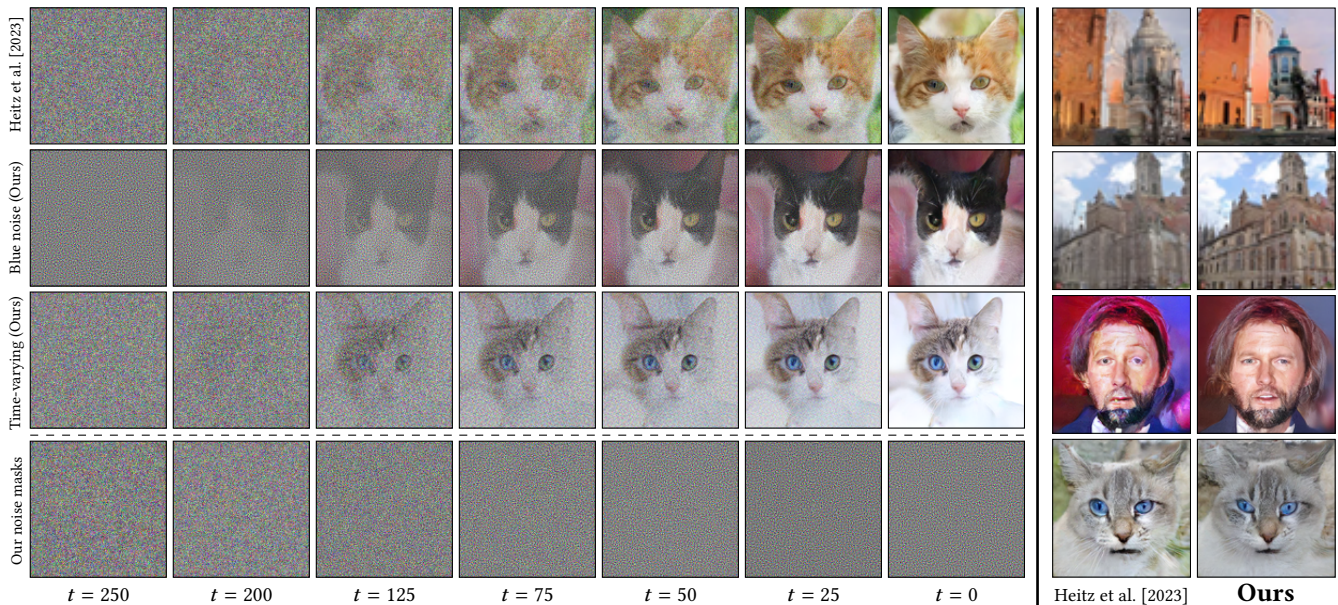


Fig. 1. In conventional diffusion-based generative modeling, data is corrupted by adding Gaussian (random) noise (as illustrated in the top row). Here, we explore the alternative approach of using correlated Gaussian noise in diffusion-based generative modeling. Different correlation can be used such as blue noise (second row) or a time-varying correlations (third row). Bottom row shows the corresponding noise mask for time-varying example at different time steps. Rightmost two columns show visual comparisons on generated images (from the same initial noise) between Heitz et al. [2023] (using only Gaussian noise) and Ours (using time-varying noise). Our generated images are more natural-looking and detailed with less artifacts.

Most of the existing diffusion models use Gaussian noise for training and sampling across all time steps, which may not optimally account for the frequency contents reconstructed by the denoising network. Despite the diverse applications of correlated noise in computer graphics, its potential for improving the training process has been underexplored. In this paper, we introduce a novel and general class of diffusion models taking correlated noise within and across images into account. More specifically, we propose a time-varying noise model to incorporate correlated noise into the training process, as well as a method for fast generation of correlated noise mask. Our model is built upon deterministic diffusion models and utilizes blue noise to help improve the generation quality compared to using Gaussian white (random) noise only. Further, our framework allows introducing correlation across images within a single mini-batch to improve gradient flow. We

perform both qualitative and quantitative evaluations on a variety of datasets using our method, achieving improvements on different tasks over existing deterministic diffusion models in terms of FID metric. Code will be available at <https://github.com/xchhuang/bndm>.

CCS Concepts: • **Computing methodologies** → **Neural networks; Computer graphics.**

Additional Key Words and Phrases: Blue noise, Diffusion models, Generative modeling

## ACM Reference Format:

Xingchang Huang, Corentin Salaün, Cristina Vasconcelos, Christian Theobalt, Cengiz Öztireli, and Gurprit Singh. 2024. Blue noise for diffusion models. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers '24 (SIGGRAPH Conference Papers '24)*, July 27-August 1, 2024, Denver, CO, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3641519.3657435>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).  
SIGGRAPH Conference Papers '24, July 27-August 1, 2024, Denver, CO, USA  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0525-0/24/07.  
<https://doi.org/10.1145/3641519.3657435>

## 1 INTRODUCTION

Since the groundbreaking work by Sohl-Dickstein et al. [2015]; Ho et al. [2020]; Song and Ermon [2019], there has been extensive research on diffusion models. These models have demonstrated superior performance in terms of generative quality and training stability compared to Generative Adversarial Networks (GANs) [Dhariwal and Nichol 2021] for image synthesis. Additionally, diffusion models can be trained to perform various tasks such as text-to-image synthesis, image inpainting, image super-resolution, and image editing.

Typically, a diffusion model consists of two processes: forward and backward. In the forward process, the model gradually adds noise to an original data point (e.g., an image), transforming it into a random noise pattern. In the backward process, the model learns to reconstruct the original data from this noise using a denoising neural network. The denoising network initially focuses on reconstructing the coarse shape and structure (low-frequency components) in the early time steps. As the time steps decrease, the denoising network progressively refines the details (high-frequency components). This behavior indicates that diffusion models generate data in a coarse-to-fine manner and have a hidden relationship with frequency components.

Despite these advancements, there is limited research on the relationship between this behavior and the noise used during the forward and backward processes. Most existing diffusion models rely solely on Gaussian noise, also known as uncorrelated Gaussian noise or Gaussian white noise, as its frequency power spectrum spans all frequencies (similar to the *white* color). While correlated noise has not been thoroughly examined in diffusion models, there has been some relevant research that has delved into this domain. For example, Rissanen et al. [2023] propose a diffusion process inspired by heat dissipation to explicitly control frequencies. Similarly, Voleti et al. [2022] suggest using non-isotropic noise as a replacement for isotropic Gaussian noise in score-based diffusion models. Despite their potential, both methods face limitations regarding the quality of the generated images, which could explain their limited adoption in mainstream models.

In this paper, we propose a new diffusion model that supports a diffusion process with time-varying noise. Our goal is to utilize correlated noise, such as blue noise ([Ulichney 1987]), to enhance the generative process. Blue noise is characterized by a power spectrum with no energy in its low-frequency region. Our focus is on blue noise masks ([Ulichney 1999]), which provide noise profiles with blue noise properties. We propose using these blue noise *masks* to design a time-varying noise for diffusion-based generative modeling. Generating such correlated noise masks for diffusion is a time-consuming process, as it may require generating thousands to millions of masks on the fly. To address this issue, we propose an efficient method to generate Gaussian blue noise masks on the fly for both low-dimensional and high-dimensional images. In summary, our contributions are as follows:

- We propose a framework that investigates the impact of correlated noise and correlation across training images on generative diffusion models.
- Based on our framework, we introduce a deterministic diffusion process with time-varying noise, allowing control over the correlation introduced in the model at each step.
- We overcome the computational challenges of generating correlated noise masks by introducing a real-time mask generation approach.
- By interpolating Gaussian noise and Gaussian blue noise using our proposed time-varying noise model, our model outperforms existing deterministic models like IADB [Heitz et al. 2023] or DDIM [Song et al. 2021a] in various image generation tasks.

## 2 RELATED WORK

*Blue noise.* Blue noise is a type of noise characterized by high frequency signals and the absence of low frequencies. It has found numerous applications in computer graphics. One such application is the utilization of blue noise masks, originally introduced by Ulichney [1993], for image dithering to enhance their perceptual quality. Blue noise masks are also employed in Monte Carlo rendering to improve the distribution of error, as demonstrated by Georgiev and Fajardo [2016] and Heitz and Belcour [2019]. The relationship between blue noise and denoising in rendering has been further explored by Chizhov et al. [2022] and Salaün et al. [2022], revealing that combining blue noise with a low pass filter can reduce perceptual errors. To leverage the advantageous denoising properties of blue noise masks, we propose to use them as additive noise to corrupt the data for diffusion-based generative modeling.

*Diffusion models.* There are various formulations for image generation by diffusion models, including stochastic [Ho et al. 2020; Song and Ermon 2019; Song et al. 2021b] and deterministic ones [Song et al. 2021a; Heitz et al. 2023]. Diffusion models also extend beyond image generation to video generation [Ho et al. 2022] and 3D content generation [Poole et al. 2023]. More comprehensive reviews can be found in the surveys by Cao et al. [2024] and Po et al. [2023].

Diffusion models are known to be slow to train, as well as slow in the generative process. How to speed up the generative process to generate images in a few steps becomes an increasingly important research topic [Lu et al. 2022; Liu et al. 2023a,b; Salimans and Ho 2022; Karras et al. 2022, 2023; Song et al. 2023; Luo et al. 2023]. Orthogonal to reducing the number of inference steps, some work focus on developing a more general framework that can support various types of noise addition [Jolicœur-Martineau et al. 2023] or image corruption operations [Bansal et al. 2024]. Some work explicitly take frequency of the image content into account to model the generative process in a coarse-to-fine manner [Rissanen et al. 2023; Phung et al. 2023]. However, there exists limited work researching the how the frequency of noise used in image corruption can make an impact on the denoising process in diffusion-based generative modeling. To understand this problem, we propose a framework utilizing correlated noise to improve the denoising process.

## 3 OUR METHOD

A generative model based on diffusion comprises two key processes: a forward process and a backward process. In the forward process, noise, denoted as  $\epsilon$ , is introduced to corrupt an initial image,  $\mathbf{x}_0$ , by scaling it with a factor determined by a discrete-time parameter,

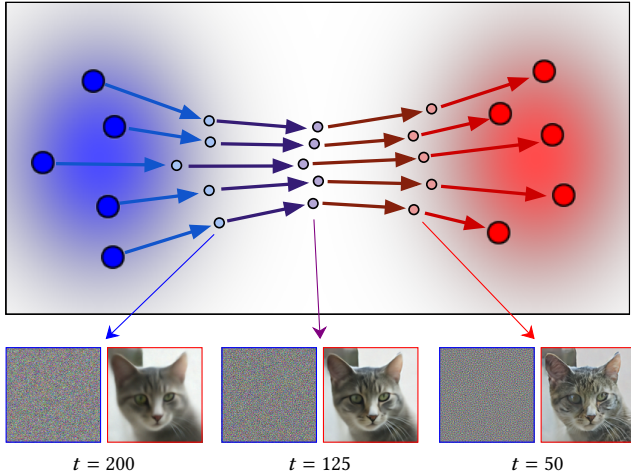


Fig. 2. Schema of diffusion process using our time-varying noise. The diffusion transforms initial noise distribution (blue) into the target data distribution (red). Five examples are shown with the intermediates diffusion steps between the two distributions. For one of the data we illustrate the intermediates time steps with the current expect result and noise. The evolution of the noise from random to blue noise is visible as well as the quality of the expected result.

$t$ . Here,  $\mathbf{x}_0$  represents a real image sampled from the training data distribution, denoted as  $p_0$ . The time step,  $t$ , ranges from 0 to  $T - 1$ , where  $T$  is the total number of discrete time steps. The corrupted image, along with the corresponding time step  $t$ , is then used as input to train a neural network,  $f_{\theta}(\mathbf{x}_t, t)$ .

In the backward process, the trained network is employed to denoise pure noise and generate new images. Figure 2 illustrates this process. Starting from Gaussian noise (blue distribution), the image iteratively passes through the network to eventually yield a fully denoised image, aligning with the target distribution (red distribution). Intermediate steps of the process involve a mixture of noise and image. Three examples are visible in the figure. As more time steps are executed (with  $t$  closer to 0), the image quality improves, and more details emerge. In this example, the intermediate noise transitions from Gaussian noise to Gaussian blue noise following a time-varying schedule following Section 3.2.

This section explores correlations across two different axes: across pixels in the noise and across images within a mini-batch. To demonstrate the impact of correlations between noise masks and images, we build a deterministic diffusion process with time-varying noise following the work by Heitz et al. [2023] namely, the IADB method. For the sake of simplicity and fair one-to-one comparison, our method was developed on top of IADB while preserving its characteristics and hyperparameters other than the ones described as new on our method. But our method is general enough that could be potentially explored on top of other existing generative diffusion process.

For IADB, the forward and backward processes and the objective function are defined as the following:

$$\mathbf{x}_t = \alpha_t \epsilon + (1 - \alpha_t) \mathbf{x}_0 \quad (1)$$

$$\mathbf{x}_{t-1} = \mathbf{x}_t + (\alpha_t - \alpha_{t-1}) f_{\theta}(\mathbf{x}_t, t) \quad (2)$$

$$\mathcal{L}_{IADB} = \sum_t (f_{\theta}(\mathbf{x}_t, t) - (\mathbf{x}_0 - \epsilon))^2 \quad (3)$$

With  $\mathbf{x}_0$  a target image,  $\epsilon \sim \mathcal{N}(\mathbf{0}, I)$  a random Gaussian noise and  $\alpha_t$  and  $\alpha_{t-1}$  two blending coefficient. The network model is referred as  $f_{\theta}$  and take 2 input parameters : a corrupted image  $\mathbf{x}_t$  and the timestep  $t$ . A stochastic formulation of IADB also exist but this work will focus on the deterministic variant for its stability.

### 3.1 Correlated noise

In a deterministic diffusion process noise masks are used as initialization of the backward process to generate images and at each training step to corrupt target images. Mask generation during training is a critical factor and must meet specific requirements. The process must be stochastic to produce different masks at each iteration, reducing overfitting and increasing diversity in the generated results. Mask generation also required to be of fast computation as it is employed at every training step.

Gaussian noise naturally meet this requirements but it is not the case for all correlated mask method. In particular, IADB uses mask generated from a multivariate Gaussian distribution with zero mean and identity covariance matrix. A method to create correlated noise such as blue noise requires a non-identity covariance matrix. The covariance matrix of the blue noise mask can be estimated from a collection of masks. We employed simulated annealing using objective function from Ulichney [1993] to generate ten thousand blue noise masks. While this method yields high-quality masks, it does require significant optimization time. Then the blue noise correlation matrix  $\Sigma$  can be computed by averaging the respective correlation matrix of the example masks.

To create a noise mask with a specified covariance matrix  $\Sigma$ , the Cholesky decomposition is applied to  $\Sigma$ , resulting in a lower-triangular matrix  $L$  ( $LL^T = \Sigma$ ). Finally, the random vector is multiplied with  $L$  to produce the desired noise mask efficiently:

$$\mathbf{b} = L\epsilon \quad (4)$$

where  $\epsilon \sim \mathcal{N}(\mathbf{0}, I)$  is a unit-variance Gaussian distribution. Figure 3 shows one realization of Gaussian blue noise mask generated using Eq. (4). Each row or column in the  $L$  represent a pixel index of the noise mask. Each cell in the  $L$  represents the correlation strength between the pixels in the noise mask. Positive values correspond to bright cells that represent positive correlations. Similarly, negative values correspond to dark cells representing negative correlations. For each pixel, only its adjacent pixels have values far from zero while other non-adjacent pixels have values close to zero, as indicated by the white and black lines. Note that we use the name Gaussian blue noise for  $\mathbf{b}$  in this paper, but is different from the Gaussian Blue Noise method from Ahmed et al. [2022].

*Higher-dimension.* Matrix-vector multiplication are computationally expensive when the  $L$  matrix is high-dimensional. Increasing






Fig. 3. To generate on-the-fly Gaussian blue noise masks  $\mathbf{b}$  (leftmost), we pre-compute a lower triangular matrix  $L$ . We then multiply this matrix with Gaussian noise  $\epsilon$  to obtain  $\mathbf{b} = L\epsilon$ . For Gaussian noise of size  $64 \times 64$ , the lower triangular matrix has a size of  $64^2 \times 64^2$ . Here we show a  $64 \times 64$  zoom-in version of the matrix to better visualize the positive and negative correlations shown as the white and black lines, respectively.




Fig. 4. Visualization of linearly interpolated noises from Gaussian noise to Gaussian blue noise at resolution  $64^2$  (top row), and the corresponding frequency power spectra (bottom row).

directly the matrix size to generate higher-dimensional noise becomes slower than generating (uncorrelated) Gaussian noise using a modern machine learning framework such as PyTorch [Paszke et al. 2017]. As noise generation is used at every training step, the overhead should remain minimal. Therefore, for generating higher-dimensional noise masks, we adapt Kollig and Keller [2002] method for producing our blue noise masks in which a padding of a set of lower dimensional masks is applied.

More specifically, to generate a batch of high-dimensional Gaussian blue noise mask, a larger batch is generated at resolution  $64^2$  using Eq. (4) with  $L \in \mathbb{R}^{64^2 \times 64^2}$ . Then the  $64^2$  masks are padded into larger tiles to get Gaussian blue noise masks at higher dimensions. Thus, the computation overhead for generating a  $128^2$  resolution Gaussian blue noise is negligible ( $\sim 0.0002$  seconds). Figure 1 shows examples of Gaussian blue noise with resolution  $128^2$  generated by padding. Using padding for higher-dimensional mask creates seams between the padded tiles. This artifact is practically not visible and is compensated by the low overhead of the method. We provide the masks at different resolutions, with corresponding frequency power spectra for Gaussian blue noise in the Supplemental document Sec. 3, demonstrating that the property of blue noise is preserved at different resolutions using our padding method.

### 3.2 Diffusion model with time-varying noise

With a single matrix  $L$  only a single correlation can be generated. For a diffusion model, controlling the amount of correlation introduced

#### Algorithm 1. Pseudocode for forward method

```

1: function forward( $L_w, L_b, \gamma_t$ )           ← White and blue noise matrices  $L$ ,
2:    $\mathbf{b} \leftarrow \text{get\_noise}(L_w, L_b, \gamma_t)$        blending coefficient  $\gamma_t$ 
3:    $\alpha_t \sim \mathcal{U}(0, 1)$ 
4:    $\mathbf{x}_0 \sim p_0$ 
5:    $\mathbf{x}_t \leftarrow \alpha_t \mathbf{b} + (1 - \alpha_t) \mathbf{x}_0$            ← Eq. (6)
6:   return  $\mathbf{x}_t$ 

```

#### Algorithm 2. Pseudocode for backward method

```

1: function backward()
2:    $\mathbf{x} \sim \mathcal{N}(0, I)$ 
3:   for  $t \leftarrow T$  to 1 do
4:      $\alpha_t \leftarrow \text{get\_alpha}(t)$ 
5:      $\alpha_{t-1} \leftarrow \text{get\_alpha}(t-1)$ 
6:      $\gamma_t \leftarrow \text{get\_gamma}(t)$ 
7:      $\gamma_{t-1} \leftarrow \text{get\_gamma}(t-1)$ 
8:      $\mathbf{x} \leftarrow \mathbf{x} + (\alpha_t - \alpha_{t-1}) f'_\theta(\mathbf{x}, t)$ 
9:        $+ (\gamma_t - \gamma_{t-1}) f''_\theta(\mathbf{x}, t)$            ← Eq. (7)
10:  return  $\mathbf{x}$ 

```

#### Algorithm 3. Pseudocode for time-varying get\_noise method

```

1: function get_noise( $L_w, L_b, \gamma_t$ )
2:    $\epsilon \sim \mathcal{N}(0, I)$ 
3:    $L_t \leftarrow \gamma_t L_w + (1 - \gamma_t) L_b$            ← Eq. (5)
4:    $\mathbf{b} \leftarrow L_t \epsilon$                              ← Eq. (4)
5:   return  $\mathbf{b}$ 

```

within the model at each time step is necessary. A time-varying  $L$  can be compute from 2 fixed matrices encoding 2 correlation types:

$$L_t = \gamma_t L_w + (1 - \gamma_t) L_b, \quad (5)$$

where  $L_w$  and  $L_b$  represent 2 different matrices and  $\gamma_t$  the blending coefficient. Based on it, the forward process is defined as:

$$\mathbf{x}_t = \alpha_t (L_t \epsilon) + (1 - \alpha_t) \mathbf{x}_0 \quad (6)$$

With this forward process, Gaussian noise and Gaussian blue noise are interpolated based on the time step  $t$ . More generally, this model supports smooth interpolation any two type of noises based on  $\gamma_t$ . Figure 4 shows an example of linear interpolating from Gaussian noise to Gaussian blue noise. The corresponding frequency power spectra, computed by Discrete Fourier Transform, show that the energy in the low-frequency region is decreasing from left to right.

Next, the forward process need to be inverted to define the backward process. Based on the definitions of the  $L$  and the forward process, we can derive the backward step as the following:

$$\mathbf{x}_{t-1} = \mathbf{x}_t + (\alpha_t - \alpha_{t-1})(\mathbf{x}_0 - L_t \epsilon) + (\gamma_t - \gamma_{t-1}) \alpha_{t-1} (L_b \epsilon - L_w \epsilon) \quad (7)$$

Detailed derivations can be found in the Supplemental document Sec. 1. Here,  $L_w$  is an identity matrix representing Gaussian (white) noise,  $L_b$  is the matrix defined in Eq. (4). When  $L_b = L_w$ , our model falls back to IADB. When  $L_b \neq L_w$ , we obtain a more general model with time-varying noises.

In IADB, the network is designed to learn only the term  $\mathbf{x}_0 - L_t \epsilon$ , where  $L_t$  is simply an identity matrix in their case. Here we can



train the network to learn both terms in Eq. (7). A brute force way to achieve this would be using two neural networks. However, this is not practical as it will introduce significantly more computation than IADB. We choose to output a 6-channel image, representing the two terms as two 3-channel images, noted as  $f'_\theta(\mathbf{x}_t, t)$  and  $f''_\theta(\mathbf{x}_t, t)$ , respectively. The desired network output becomes  $\mathbf{x}_0 - L_t \epsilon$  and  $\alpha_{t-1}(L_b \epsilon - L_w \epsilon)$ . Therefore, the loss function becomes:

$$\mathcal{L}_{Ours} = \sum_t ((f'_\theta(\mathbf{x}_t, t) - (\mathbf{x}_0 - L_t \epsilon))^2 + \frac{\gamma_t - \gamma_{t-1}}{\alpha_t - \alpha_{t-1}} (f''_\theta(\mathbf{x}_t, t) - \alpha_{t-1}(L_b \epsilon - L_w \epsilon))^2) \quad (8)$$

Note that though our model is trained with time-varying noises, it is still deterministic during the backward process. The backward process starts with an initial Gaussian noise and no additional noises is required in the intermediate time steps. Instead, the network learns to guide the backward process in a time-varying denoising manner.

The procedures of forward, backward and noise generation are summarized in Algorithms 1 to 3. In Algorithm 2, we consider `get_alpha` ( $\alpha$ -scheduler) as a linear function ( $\alpha_t = t/T$ ) following Heitz et al. [2023], but it can be non-linear functions as well. Next, we define `get_gamma` as a general sigmoid-based function in Eq. (9). The weighted term  $(\gamma_t - \gamma_{t-1})/(\alpha_t - \alpha_{t-1})$  in Eq. (8) automatically accounts for the difference between  $\alpha$ -schedulers and  $\gamma$ -scheduler. When  $\gamma_t - \gamma_{t-1}$  is small, the contribution of  $f''_\theta(\mathbf{x}_t, t)$  decreases. This is consistent to the backward process described in Eq. (7), where  $f''_\theta(\mathbf{x}_t, t)$  is less important when  $\gamma_t - \gamma_{t-1}$  is small.


*Noise scheduler.* Inspired by the study from Chen [2023], the scheduler choice have an importance impact in particular with increasing image resolution. We parameterize `get_gamma`, the  $\gamma$ -scheduler, as a sigmoid-based function to control the interpolation between two noises. More specifically, the  $\gamma$ -scheduler is parameterized by 3 parameters: *start*, *end*,  $\tau$  according to Chen [2023]:

$$\text{sigmoid}\left(\frac{\text{start} + (\text{end} - \text{start}) * t/T}{\tau}\right) \quad (9)$$

where  $\text{sigmoid}(x) = 1/(1 + e^{-x})$  and  $t$  is the time step.

Since it is not known how to set *start*, *end* and  $\tau$  in advance, we consider optimizing them in addition to the network parameters where  $\text{start} \in [-3, 0]$ ,  $\text{end} \in (0, 3]$ ,  $\tau \in [0.01, 1000.0]$ . During the initial experiments, we found that *start* and *end* converged stably to around 0 and 3, while  $\tau$  converged to around 0.2 or kept increasing, depending on the image resolution. Meanwhile, we found that optimizing these 3 parameters took extra epochs to converge and made the training of the network more difficult, due to their changes over epochs.

To make the choice of the 3 parameters more practical, we choose to fix *start* = 0, *end* = 3 and set  $\tau$  based on the image resolution:  $\tau = 0.2$  for  $128^2$ ,  $\tau = 1000$  for  $64^2$  images. The curves of the  $\gamma$ -scheduler with different  $\tau$  values are



shown in the inset image.

We summarize the values of the 3 parameters we use in the Supplemental document Sec. 2 for all experiments.

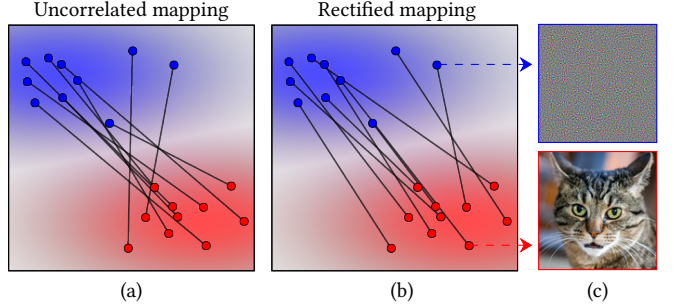


Fig. 5. Visualization of the impact of rectified mapping on mini-batch noise-mapping pairing. Blue and red points respectively represent the randomly selected noise and target image sampled in a given mini-batch from the underlying blue and red distribution. Standard practice (a) consist in a random mapping between the noise and target images. Our rectified mapping (b) improved it by reducing distance between the data pair. One example of noise and target from the mini-batch is visible (c). This example have been generated using our noise mask and mapping algorithm.

*Discussion.* Our time-varying noise model provides more capacity to choose a data-dependent scheduler for  $\gamma_t$  to improve the denoising process. One potential issue is that we need extra epochs to search for optimal parameters for the  $\gamma$ -scheduler. To alleviate this problem, we propose a practical solution, which is to fix  $\tau$  based on our initial optimization. But how to choose the  $\gamma$ -scheduler in a more efficient way requires more study in the future.

### 3.3 Data sample correlation using rectified mapping

The previous paragraphs have demonstrated the use of correlated noise across pixels to enhance the diffusion process. Correlation can also be employed within a single mini-batch to improve the mapping between noise and the target image.

Inspired by Rectified flow [Liu et al. 2023a] and InstafLOW [Liu et al. 2023b], correlation can be utilized to rectify the paired noise-image. Figure 5 visualizes a single mini-batch of paired data sample  $\mathbf{x}_0$  (red distribution) and noise  $\mathbf{b}$  (blue distribution) during a training iteration, and our rectified mapping. Previous work (Fig. 5(a)) applies a random mapping between  $\mathbf{x}_0$  and  $\mathbf{b}$ . The noise-data mapping can be improved by applying an in-context stratification before feeding them into the forward process (Fig. 5(b)). This rectified mapping reduces the distance between each noise and its target image, resulting in a more direct trajectory. To find the mapping, we compute the squared distance between noises and images at the individual pixel level using the L2 norm. Then, for each  $\mathbf{b}$ , the  $\mathbf{x}_0$  with the shortest distance that has not yet been used is selected. This improved mapping ensures that a specific image will consistently be associated with the same type of noise during the training process, resulting in a smooth gradient flow across time steps.

## 4 EXPERIMENTS

### 4.1 Implementation details

We use CelebA [Lee et al. 2020], AFHQ-Cat [Choi et al. 2020] and LSUN-Bedroom [Yu et al. 2015] datasets, with different resolutions, for unconditional/conditional image generation. More details on experimental setup can be found in Supplemental document Sec. 2.

Our framework is implemented in Pytorch [Paszke et al. 2017] based on the official implementations from Song et al. [2021a]; Heitz et al. [2023]. We use 2D U-Net [Ronneberger et al. 2015] implemented in diffusers library [von Platen et al. 2022]. More details about the network architecture and training details, including the values of  $\tau$  in Eq. (9), can be found in the Supplemental document Sec. 2. Regarding the hyperparameters in diffusion models, we use  $T = 1000$  for training and  $T = 250$  testing. To optimize the network parameters, we use AdamW optimizer [Loshchilov and Hutter 2017] with learning rate 0.0001. We use 4 NVIDIA Quadro RTX 8000 (48 GB) GPUs to train and test on all datasets.

For evaluation, we use FID [Heusel et al. 2017], Precision and Recall [Kynkäänniemi et al. 2019] to measure the generative quality of all models. The metrics are computed using the implementation from [Stein et al. 2024], with Inception-v3 network [Szegedy et al. 2016] as backbone. We generate 30k images to compute FID, Precision and Recall for all datasets.

### 4.2 Image generation

We compare our method with two existing deterministic diffusion models DDPM Ho et al. [2020], DDIM [Song et al. 2021a] and IADB [Heitz et al. 2023] on unconditional image generation. To ensure equitable comparisons, we employ identical initial Gaussian noise across all methods throughout the generative process. Note that DDIM is trained using the diffusers library [von Platen et al. 2022] with the same training setup compared to IADB and Ours. We also compare with stochastic diffusion models DDPM [Ho et al. 2020] and IHDM [Rissanen et al. 2023] for completeness using the same number of time steps. Our method shows consistent improvement over the two methods.

The results on AFHQ-Cat ( $64^2$ ), LSUN-Church ( $64^2$ ), and Celeba ( $64^2$ ) at the same resolution ( $64^2$ ) are shown in Figure 11. Our method exhibits the blue noise effect starting from approximately  $t = 75$ , which visually distinguishes it from other methods. In terms of the generated images at time step  $t = 0$ , our method produces images with less distortion around the pillars of the building and more detailed content around the windows and doors. In addition to visual comparisons, the quantitative evaluations presented in Table 1 demonstrate consistent improvements of our method over IADB and DDIM for datasets with a resolution of  $64^2$ .

For higher-resolution results, we observe the difference in the generated content starting at approximately  $t = 75$ , as depicted in Fig. 12. Towards the end of the backward process, around  $t = 25$ , we begin to notice the emergence of the blue noise effect as we use  $\tau = 0.2$  (Eq. (9)). To examine the details more closely, please zoom in. Additionally, we offer a Supplemental HTML viewer where the intermediate generated images can be interactively visualized at various time steps during the backward process. In terms of realism, our generated images exhibit improved quality in regions such as

Table 1. Quantitative FID score comparisons among IHDM [Rissanen et al. 2023], DDPM [Ho et al. 2020], DDIM [Song et al. 2021a], IADB [Heitz et al. 2023], and our method across diverse datasets. Notably, our approach exhibits improvements over IADB on every evaluated dataset. While our method is outperformed by DDIM on only one dataset, it’s worth noting that IADB also performs poorly on the same dataset. Additional metrics are provided in the Supplemental document Sec. 3.

FID (↓)	IHDM	DDPM	DDIM	IADB	Ours
AFHQ-Cat ( $64^2$ )	11.02	9.75	9.82	9.19	<b>7.95</b>
AFHQ-Cat ( $128^2$ )	15.40	12.41	10.73	10.81	<b>9.47</b>
CelebA ( $64^2$ )	14.30	8.56	9.26	7.53	<b>7.05</b>
CelebA ( $128^2$ )	36.93	15.06	<b>11.92</b>	20.71	16.38
LSUN-Church ( $64^2$ )	17.76	13.07	16.46	13.12	<b>10.16</b>

hair, mouth, and eyes compared to IADB, as shown in Fig. 12. When compared to DDIM, our method achieves similar visual quality. Quantitative results on the CelebA ( $128^2$ ) dataset, as presented in Table 1, demonstrate that DDIM outperforms IADB and our method. This outcome is attributed to DDIM employing a different expression for  $\alpha_t$  in Eq. (1), as discussed by in Heitz et al. [2023]. Depending on the dataset, DDIM may outperform IADB due to the distinct choice of  $\alpha_t$ . However, this is not a limitation for either IADB or our method.

Additionally, our framework can also consider rectified mapping across images during training. Table 2 provides the FID score with and without data correlation with respect to the number of diffusion steps, tested on AFHQ-Cat ( $64^2$ ). Rectified mapping in the mini-batch achieves lower FID when the number of steps remains low, but a slightly higher FID when increasing the step count.

We provide additional results in the Supplemental document Sec. 3, including detailed timing of the Gaussian blue noise generation and the backward process, and a nearest neighbors test to confirm that our method does not overfit the training data.

*Extension to other diffusion models.* Our method can be extended to DDIM, supported by derivations and preliminary results in Supplemental document Sec. 1, 3. Also, our method can be incorporated into LDM [Rombach et al. 2022] for high-res image generation. As shown in Fig. 9, ours generates more realistic eyes and has better FID ( $11.45 < 12.19$ ) compared with IADB on AFHQ-Cat ( $512^2$ ). Nevertheless, developing a new framework based on other models, necessitates additional effort, which we defer to future work.

Table 2. Comparing the impact of rectified mapping during training on AFHQ-Cat ( $64^2$ ). FID scores (↓) are provided with and without rectified mapping across different step counts. Correlation in the mini-batch results in lower FID at low steps but higher during slow diffusion.

Diffusion step count (t)	1	2	4	16	128	250
Uncorrelated batch mapping	402.4	330.5	130.8	14.3	<b>7.9</b>	<b>7.95</b>
Correlated batch mapping	<b>397.0</b>	<b>321.8</b>	<b>118.3</b>	<b>12.4</b>	8.0	8.2

### 4.3 Conditional image generation

Besides unconditional generation from noise, our model also works for conditional image generation, such as image super-resolution,

by simply concatenating the conditional low-resolution image with the noisy image as input.

Figure 6 shows comparisons between IADB and Ours for image super-resolution on the LSUN-Church dataset from resolution  $32^2$  to  $128^2$ . Our method outperforms IADB quantitatively according to SSIM [Wang et al. 2004], PSNR and mean squared error (MSE). Our method outperforms IADB in terms of fidelity to the reference, as evidenced by its lower MSE. Visually, IADB tends to introduce excessive details, particularly in the bottom portion of the first image. Our method also effectively preserves straight lines throughout the image. The quantitative results of all image-super resolution experiments can be found in the Supplemental document Sec. 3, showing that our method consistently outperforms IADB.

#### 4.4 Ablation study and analysis

*Combinations of noises.* To confirm that Gaussian blue noise works due to its high-frequency property, we replace Gaussian blue noise by Gaussian red noise, a low-frequency noise visualized in Fig. 10. Red noises are generated using the same method [Ulichney 1993] by simply maximizing the objective function instead of minimizing it. Then we compute the corresponding covariance matrix and lower triangular matrix so that we can generate Gaussian red noise for our framework. As shown in Fig. 7, using Gaussian red noise in our framework failed to recover the fine details due to its low-frequency property. Table 3 shows that replacing Gaussian blue noise by Gaussian red noise dramatically drops the Precision, while the Recall is comparable. This is consistent to the visual observations in Fig. 7 as Precision mainly measures the realism of the generated images.

Table 3. Ablation study on different combinations of noises using our framework on AFHQ-Cat ( $128^2$ ). The last two rows mean blending Gaussian noise with Gaussian red or blue noise using the  $\gamma$ -scheduler with  $\tau = 0.2$ .

Noise model	FID ( $\downarrow$ )	Precision ( $\uparrow$ )	Recall ( $\uparrow$ )
Ours (Gaussian white noise only)	10.81	<b>0.78</b>	0.31
Ours (Gaussian blue noise only)	17.61	0.59	0.18
Ours (Gaussian white+red noise)	13.64	0.67	<b>0.34</b>
Ours (Gaussian white+blue noise)	<b>9.47</b>	<b>0.78</b>	<b>0.34</b>

Another option is to use only Gaussian blue noise, which has been shown in Fig. 1 (second row). The final generated images are less realistic compared to IADB and Ours. The visual quality is also consistent with the quantitative metrics as shown in Table 3. But it is worth mentioning that in the case of using only Gaussian blue noise, we can observe that the content of the image appears faster and cleaner at early time steps compared to other choices, as shown in Fig. 1. We performed an additional experiment called early stopping. This is to show that if we stop at early steps, using only Gaussian blue noise gives better results than using only Gaussian noise. As shown in Supplemental document Fig. 3, the results of using only Gaussian blue noise (second row) stopped at  $t = 200$  are with sharper details than using only Gaussian noise (first row). Quantitative evaluation of early stopping can be found in the Supplemental document Sec. 3. However, as blue noise has no energy in low-frequency region, it is restricting the diffusion process to a limited range of directions. For this reason, it becomes

difficult to refine the intermediately generated content in later time steps and thus results in worse quality, as shown in Table 3. Instead, our method takes blue noise into account from middle or later time steps, when low-frequency components are already visible and the network is more focusing on refining high-frequency details.

*Diffusion with different noise magnitude.* Since using only Gaussian blue noise at all time steps leads to degraded quality, we further conduct an analysis on diffusion at later time steps by explicitly ignoring the early time steps. We compare Gaussian noise and Gaussian blue noise by training a diffusion model (IADB) up to some time steps with certain noise magnitude (e.g., 30%). During the testing phase, ground truth images are provided so we are able to compare the denoised images with the ground truth ones. Running with 100% of noises, the experiment would fall back to the standard diffusion generative process. Based on Fig. 8, by using Gaussian blue noise we generate more detail- and content-preserving images under various noise magnitude, compared to the ones using Gaussian noise. This indicates that Gaussian blue noise is suitable for denoising when low-frequency components become visible. This is consistent to our idea blending Gaussian blue noise from middle or later time steps.

*More ablations.* We further compare different  $\gamma$  values and a cosine-based scheduler [Nichol and Dhariwal 2021] used for the  $\gamma$ -scheduler. Also, we compare different Gaussian blue noise mask sizes used for padding/tiling. More details can be found in Supplemental document Sec. 3.

## 5 CONCLUSION

We have presented a new method for incorporating correlated noise into deterministic generative diffusion models. Our technique involves using a combination of uncorrelated and correlated noise masks generated using matrix-based methods. By investigating different noise correlation, we have uncovered the intricate relationship between noise characteristics and the quality of generated images. Our findings indicate that high-frequency noise is effective at preserving details but struggles with generating low-frequency components, whereas low-frequency noise hinders the generation of complex details. To achieve optimal image quality, we propose selectively using different types of noise in a time-dependent manner, leveraging the strengths of each noise component. To validate the effectiveness of our approach, we conducted extensive experiments using it in conjunction with the well-known method IADB [Heitz et al. 2023]. By keeping the training data and optimization hyperparameters consistent, we consistently observed significant improvements in image quality across various datasets. These results demonstrate the superiority of our approach in enhancing image generation capabilities of deterministic diffusion models.

*Limitations.* Currently, parameter tuning for our  $\gamma$ -scheduler depends on the image resolution. It is also computationally intensive to compute Gaussian blue noise masks while extending our approach to higher resolution models. All data from a specific mini-batch must be on a single GPU for our rectified mapping to function. Further research is needed to expand this to distributed training involving inter-GPU synchronization.



*Future work.* We believe our proposed model will inspire new research directions in designing noise patterns for improving efficiency of generative diffusion models. An interesting future work would be extending our model to interpolate more than two noises to take into account more different types of noises, such as low-pass and band-pass noises. This may provide more degree of freedom to improve the training and sampling efficiency of the diffusion models. Further, we can design more advanced techniques to correlate data samples during training, which is orthogonal to using correlated noise. Extending our framework (e.g., the time-varying noise model) to stochastic models [Ho et al. 2020; Song et al. 2021b] and even fewer-step models [Karras et al. 2022; Song et al. 2023; Luo et al. 2023] would be another interesting future direction. In this way, our framework can be generalized to state-of-the-art denoising diffusion models.

In terms of applications, we tested our model on 2D unconditional and conditional image generation. Interesting future work would include generalizing our model to synthesize other data representation such as video and 3D mesh.

## ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their detailed and constructive comments. Xingchang Huang is supported by Saarbrücken Research Center for Visual Computing, Interaction and Artificial Intelligence, a joint strategic partnership between MPI and Google. Due to the agreement, we do not display the original images from the CelebA dataset [Lee et al. 2020] in Figures 1, 8, 11, 12. The images shown in these figures are all generated with our method.

## REFERENCES

- Abdalla GM Ahmed, Jing Ren, and Peter Wonka. 2022. Gaussian blue noise. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–15.
- Arpit Bansal, Eitan Borgnia, Hong-Min Chu, Jie Li, Hamid Kazemi, Furong Huang, Micah Goldblum, Jonas Geiping, and Tom Goldstein. 2024. Cold diffusion: Inverting arbitrary image transforms without noise. *Advances in Neural Information Processing Systems* 36 (2024).
- Hanqun Cao, Cheng Tan, Zhangyang Gao, Yilun Xu, Guangyong Chen, Pheng-Ann Heng, and Stan Z Li. 2024. A survey on generative diffusion models. *IEEE Transactions on Knowledge and Data Engineering* (2024).
- Ting Chen. 2023. On the importance of noise scheduling for diffusion models. *arXiv preprint arXiv:2301.10972* (2023).
- Vassilen Chizhov, Iliyan Georgiev, Karol Myszkowski, and Gurprit Singh. 2022. Perceptual Error Optimization for Monte Carlo Rendering. *ACM Transactions on Graphics* 41, 3, Article 26 (mar 2022), 17 pages. <https://doi.org/10.1145/3504002>
- Yunje Choi, Youngjung Uh, Jaeyun Yoo, and Jung-Woo Ha. 2020. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 8188–8197.
- Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* 34 (2021), 8780–8794.
- Iliyan Georgiev and Marcos Fajardo. 2016. Blue-Noise Dithered Sampling. In *ACM SIGGRAPH 2016 Talks (Anaheim, California) (SIGGRAPH '16)*. Association for Computing Machinery, New York, NY, USA, Article 35, 1 pages. <https://doi.org/10.1145/2897839.2927430>
- Eric Heitz and Laurent Belcour. 2019. Distributing Monte Carlo Errors as a Blue Noise in Screen Space by Permuting Pixel Seeds Between Frames. *Computer Graphics Forum* (2019). <https://doi.org/10.1111/cgf.13778>
- Eric Heitz, Laurent Belcour, and Thomas Chambon. 2023. Iterative  $\alpha$ -(de) blending: A minimalist deterministic diffusion model. In *ACM SIGGRAPH 2023 Conference Proceedings*. 1–8.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* 30 (2017).
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems* 33 (2020), 6840–6851.
- Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. 2022. Video diffusion models. *Advances in Neural Information Processing Systems* 35 (2022), 8633–8646.
- Alexia Jolicœur-Martineau, Kilian Fatras, Ke Li, and Tal Kachman. 2023. Diffusion models with location-scale noise. *arXiv preprint arXiv:2304.05907* (2023).
- Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. 2022. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems* 35 (2022), 26565–26577.
- Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. 2023. Analyzing and improving the training dynamics of diffusion models. *arXiv preprint arXiv:2312.02696* (2023).
- Thomas Kollig and Alexander Keller. 2002. Efficient multidimensional sampling. In *Computer Graphics Forum*, Vol. 21. Wiley Online Library, 557–563.
- Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. 2019. Improved precision and recall metric for assessing generative models. *Advances in Neural Information Processing Systems* 32 (2019).
- Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. 2020. MaskGAN: Towards Diverse and Interactive Facial Image Manipulation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xingchao Liu, Chengyue Gong, and Qiang Liu. 2023a. Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net. <https://openreview.net/pdf?id=XVJTT1nw5z>
- Xingchao Liu, Xiwen Zhang, Jianzhu Ma, Jian Peng, et al. 2023b. InstafLOW: One step is enough for high-quality diffusion-based text-to-image generation. In *The Twelfth International Conference on Learning Representations*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017).
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. 2022. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems* 35 (2022), 5775–5787.
- Simian Luo, Yiqin Tan, Longbo Huang, Jian Li, and Hang Zhao. 2023. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378* (2023).
- Alexander Quinn Nichol and Prafulla Dhariwal. 2021. Improved denoising diffusion probabilistic models. In *International conference on machine learning*. PMLR, 8162–8171.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. (2017).
- Hao Phung, Quan Dao, and Anh Tran. 2023. Wavelet diffusion models are fast and scalable image generators. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10199–10208.
- Ryan Po, Wang Yifan, Vladislav Golyanik, Kfir Aberman, Jonathan T Barron, Amit H Bermano, Eric Ryan Chan, Tali Dekel, Aleksander Holynski, Angjoo Kanazawa, et al. 2023. State of the Art on Diffusion Models for Visual Computing. *arXiv preprint arXiv:2310.07204* (2023).
- Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. 2023. DreamFusion: Text-to-3D using 2D Diffusion. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net. <https://openreview.net/pdf?id=FjNys5c7VyY>
- Severi Rissanen, Markus Heinonen, and Arno Solin. 2023. Generative Modelling with Inverse Heat Dissipation. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net. <https://openreview.net/pdf?id=4PJUBT9f2O1>
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10684–10695.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III* 18. Springer, 234–241.
- Corentin Salaün, Iliyan Georgiev, Hans-Peter Seidel, and Gurprit Singh. 2022. Scalable multi-class sampling via filtered sliced optimal transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 41, 6 (2022). <https://doi.org/10.1145/3550454.3555484>
- Tim Salimans and Jonathan Ho. 2022. Progressive Distillation for Fast Sampling of Diffusion Models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net. <https://openreview.net/forum?id=TidIXpzh0t>
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*. PMLR, 2256–2265.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. 2021a. Denoising Diffusion Implicit Models. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net. <https://openreview.net/>

- forum?id=St1gjarCHLP
- Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. 2023. Consistency Models. In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA (Proceedings of Machine Learning Research, Vol. 202)*. PMLR, 32211–32252. <https://proceedings.mlr.press/v202/song23a.html>
- Yang Song and Stefano Ermon. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems* 32 (2019).
- Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2021b. Score-Based Generative Modeling through Stochastic Differential Equations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net. <https://openreview.net/forum?id=PXTIG12RRHS>
- George Stein, Jesse Cresswell, Rasa Hosseinzadeh, Yi Sui, Brendan Ross, Valentin Villecroze, Zhaoyan Liu, Anthony L. Caterini, Eric Taylor, and Gabriel Loaiza-Ganem. 2024. Exposing flaws of generative model evaluation metrics and their unfair treatment of diffusion models. *Advances in Neural Information Processing Systems* 36 (2024).
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2818–2826.
- Robert Ulichney. 1987. *Digital Halftoning*. MIT Press.
- Robert Ulichney. 1993. Void-and-cluster method for dither array generation. In *Electronic imaging*. <https://api.semanticscholar.org/CorpusID:120266955>
- Robert Ulichney. 1999. The void-and-cluster method for dither array generation. *SPIE MILESTONE SERIES MS 154* (1999), 183–194.
- Vikram Voleti, Christopher Pal, and Adam Oberman. 2022. Score-based denoising diffusion with non-isotropic gaussian noise models. *arXiv preprint arXiv:2210.12254* (2022).
- Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, and Thomas Wolf. 2022. Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers>.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
- Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. 2015. LSUN: Construction of a Large-scale Image Dataset using Deep Learning with Humans in the Loop. *arXiv preprint arXiv:1506.03365* (2015).



Fig. 6. Image super-resolution comparisons between IADB (SSIM/PSNR=0.57/19.46) and Ours (SSIM/PSNR=0.59/20.00) on LSUN-Church ( $32^2 \rightarrow 128^2$ ). The mean squared error w.r.t the reference is visible in the upper corner with the relative error to IADB. Our method achieves lower error and more plausible details with less hallucination.



Fig. 7. Qualitative image generation comparisons between Gaussian red noise and Gaussian blue noise using our method on AFHQ-Cat ( $128^2$ ). Our method with Gaussian blue noise creates more high-frequency details while using Gaussian red noise introduces visible artifacts.



Fig. 8. Evaluating the impact of noise magnitude on detail enhancement. Our Gaussian blue noise method better preserves fine details even with increased noise magnitude, while maintaining the integrity of the content. With 100% noise, both models fall back to full generative process.



Fig. 9. Latent diffusion model (LDM) based high-res image generation using IADB and Ours on AFHQ-Cat ( $512^2$ ). IADB introduces more artifacts around the eye regions and has worse FID (12.19 > 11.45) compared with Ours.




Fig. 10. Visualization of Gaussian noise, Gaussian blue noise and Gaussian red noise at resolution  $64^2$ . According Power spectrum is shown in the bottom right corner. While Gaussian blue noise show mostly high frequency variations, Gaussian red noise shows only low frequencies.



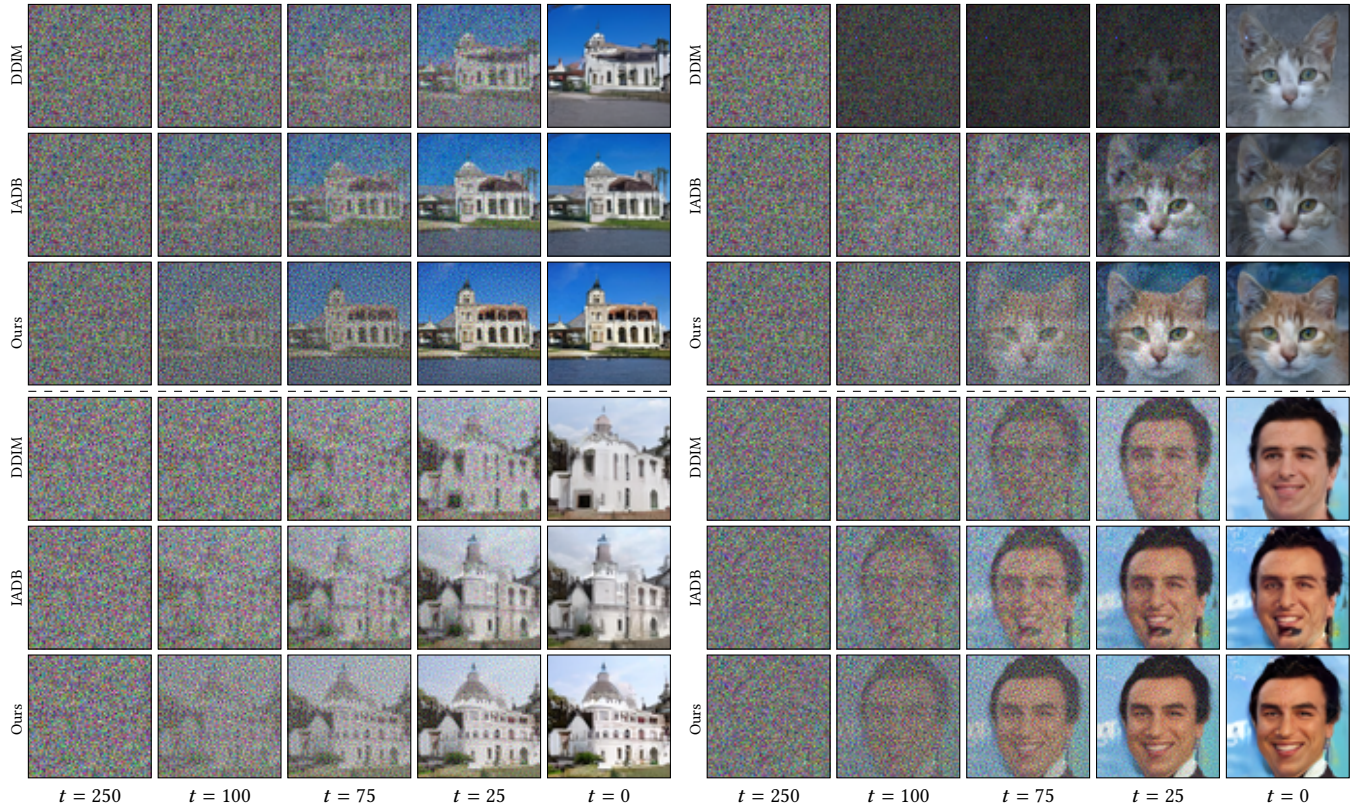


Fig. 11. Comparisons of image generation using DDIM, IADB and Ours trained on LSUN-Church ( $64^2$ ), AFHQ-Cat ( $64^2$ ) and CelebA ( $64^2$ ) datasets. For each example all methods start the diffusion from the same noise. In all cases our method achieve the highest quality result with more realistic images. Quality in details generation can be seen in the windows and doors of the buildings. By looking at the noise at different time steps the evolution from random to blue noise is visible for our method.

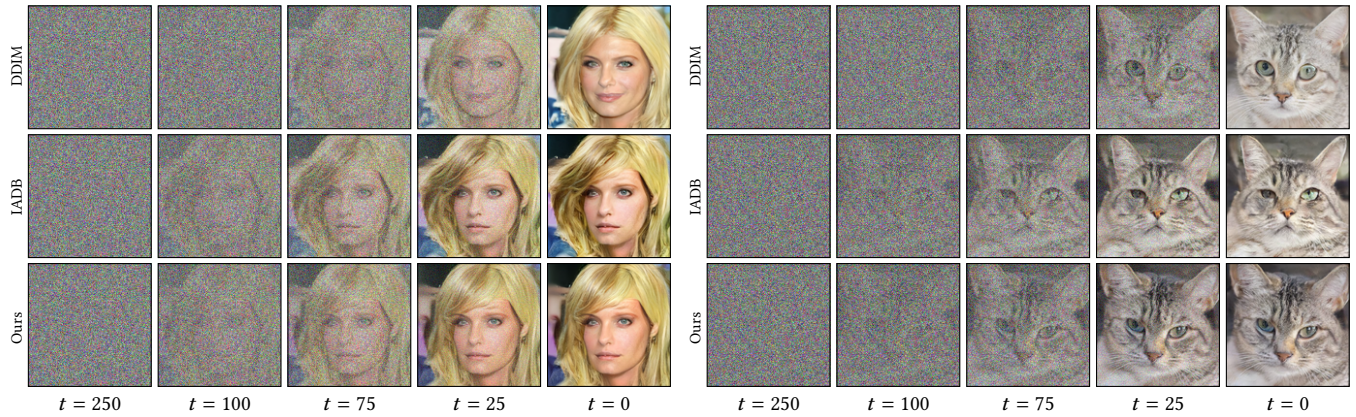


Fig. 12. Image generation comparisons between DDIM, IADB and Ours trained on CelebA ( $128^2$ ) and AFHQ-Cat ( $128^2$ ) datasets, respectively. All methods start with the same initial Gaussian noise during the backward process. Our method generates more realistic content around the hair, eye, mouth regions compared to IADB. Compared to DDIM, our method achieves similar visual quality. The impact of time-varying noise (we use  $\tau = 0.2$  in Eq. (9)) can be seen by comparing IADB and ours starting from around  $t = 75$ .