

Improving Denoising Diffusion Models via Simultaneous Estimation of Image and Noise

Zhenkai Zhang

ZHENKAI.ZHANG@STUDENT.UNIMELB.EDU.AU

Krista A. Ehinger

KRIS.EHINGER@UNIMELB.EDU.AU

Tom Drummond

TOM.DRUMMOND@UNIMELB.EDU.AU

School of Computing and Information Systems, The University of Melbourne

Editors: Berrin Yanıkoğlu and Wray Buntine

Abstract

This paper introduces two key contributions aimed at improving the speed and quality of images generated through inverse diffusion processes. The first contribution involves reparameterizing the diffusion process in terms of the angle on a quarter-circular arc between the image and noise, specifically setting the conventional $\sqrt{\alpha} = \cos(\eta)$. This reparameterization eliminates two singularities and allows for the expression of diffusion evolution as a well-behaved ordinary differential equation (ODE). In turn, this allows higher order ODE solvers such as Runge-Kutta methods to be used effectively. The second contribution is to directly estimate both the image (\mathbf{x}_0) and noise (ϵ) using our network, which enables more stable calculations of the update step in the inverse diffusion steps, as accurate estimation of both the image and noise are crucial at different stages of the process. Together with these changes, our model achieves faster generation, with the ability to converge on high-quality images more quickly, and higher quality of the generated images, as measured by metrics such as Fréchet Inception Distance (FID), spatial Fréchet Inception Distance (sFID), precision, and recall.

Keywords: Diffusion Model; Estimating Image and Noise; ODE; Faster Converge; Better Quality

1. Introduction

Diffusion models [Sohl-Dickstein et al. \(2015\)](#); [Ho et al. \(2020\)](#), which have emerged as powerful tools, are capable of generating synthetic images that exhibit both high quality and diversity. These models leverage the concept of gradual image refinement through a diffusion process, where noise is sequentially transformed into realistic images. The ability to generate visually appealing and realistic images has found applications in various fields, such as data augmentation, computer vision, and image synthesis.

Significant progress has been made in recent years in advancing diffusion models, either by improving the empirical performance [Nichol and Dhariwal \(2021b\)](#); [Song et al. \(2020a\)](#); [Song and Ermon \(2020\)](#) or by extending the capabilities of the models from a theoretical perspective [Lu et al. \(2022a\)](#); [Lu et al. \(2022d\)](#); [Song et al. \(2021\)](#); [Song et al. \(2020b\)](#), leading to improvements in image quality and generation capabilities. However, the current diffusion models face a limitation in terms of time inefficiency during the inference process, particularly in models that utilize noise as the target object [Ho et al. \(2020\)](#); [Nichol and Dhariwal \(2021b\)](#); [Song et al. \(2020a\)](#). This inefficiency is due to the substantial amount of

time required in the initial stages of sampling, where the model progresses from pure noise to low-quality images. Meanwhile, models that use images as the target object [Bansal et al. \(2022\)](#), do not encounter the same issues as noise-based diffusion models, as they can achieve a faster transition from pure noise to low-quality images. However, these models still face the challenge of directly estimating the image, which becomes increasingly difficult at the final stage of the diffusion process when the input is dominated by noise. Consequently, this difficulty leads to a lower performance of the final result than the noise-based diffusion model.

To address these limitations, in this paper, we propose a novel approach that combines the advantages of both noise-targeted training and image-targeted training in the diffusion process. By simultaneously incorporating noise and actual images as training objectives, our model aims to overcome the aforementioned limitations and achieve superior performance in terms of both result quality and diversity. Figure 1 depicts our model architecture and concepts.

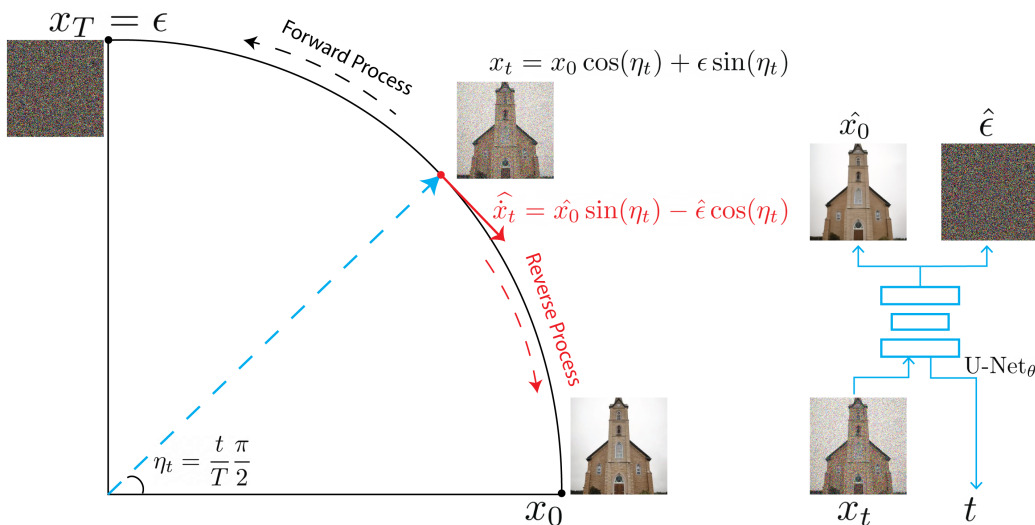


Figure 1: The overview of our model architecture and concept. It depicts how our system diffuses from input images to images with noise, eventually leading to pure noise. The figure also showcases our sampling process, which utilizes gradients to generate images from noise.

Overall, our framework proposes several practical and theoretical contributions,

- Our approach develops a new and efficient parameterization technique that results in an innovative noise scheduler. In Section 5.1 of the ablation experiments, we describe that this approach improves the overall quality of generated images while reducing the number of steps needed to produce high-quality results.

- Our proposed model significantly improves over traditional models that only estimate images or noise separately. Our model is trained to simultaneously predict both noise and images, which improves the estimate of the gradient at each timestep. We use gradients as a loss function to update the model, which enhances controllability during image generation. Additionally, we utilize gradient information to optimize the generation process during the reversed process. Our model obtains higher quality and more realistic results by calculating the gradient information of the current state and using that information to update the steps more accurately and stable, and the results are shown in Section 5.

2. Background and Related Work

In this section, a brief introduction is provided regarding the relevant work that our model is built upon.

The concept of diffusion models [Sohl-Dickstein et al. \(2015\)](#); [Ho et al. \(2020\)](#) is rooted in non-equilibrium thermodynamics, wherein the diffusion and reverse processes act as Markov chains. Unlike other approaches to generative modelling such as Variational Autoencoders (VAEs, [Kingma and Welling \(2013\)](#)), normalizing flows ([Rezende and Mohamed \(2015\)](#)), or Generative Adversarial Networks (GANs, [Goodfellow et al. \(2020\)](#)), the purpose of the diffusion model is to describe the distribution of the data by gradually reducing the noise or uncertainty of the input, and the latent space has the same dimensionality as the input space.

2.1. Noise-based diffusion models

The majority of diffusion models [Kong et al. \(2020\)](#); [Song et al. \(2020a\)](#); [Nichol and Dhariwal \(2021a\)](#); [Roman et al. \(2021\)](#); [Rombach et al. \(2021\)](#) are trained by adding Gaussian noise to the images to generate noisy samples, and feeding them into the U-Net network [Ronneberger et al. \(2015\)](#) for predicting the added noise. This process is effective because incorporating noise prediction into the training procedure can effectively disentangle the noise from the underlying signal, allowing for a more accurate and robust generation of samples, and alleviating the burden of explicitly modelling complex data patterns. The first widely followed and used model is the Denoising diffusion probabilistic models (DDPMs, [Ho et al. \(2020\)](#)), latent variable models that utilise a Markov chain concept. These models have learned Gaussian transitions that begin with a Normal distribution $p(x_T) \sim \mathcal{N}(x_T; \mathbf{0}, \mathbf{I})$ in the following form:

$$p_\theta(x_0) := \int p_\theta(x_{0:T}) dx_{1:T}, \quad p_\theta(x_{0:T}) := p(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t) \quad (1)$$

where x_1, x_2, \dots, x_T keep the same shape with x_0 , θ means the weights of neural networks, commonly using the U-net structure, and the $p_\theta(x_{0:T})$ is defined as the reverse process. The parameters θ are trained to match the data distribution $q(x_0)$ by maximizing a variational lower bound by optimizing the negative log-likelihood.

$$L_{VLB} = \max_{\theta} \mathbb{E}_{q(x_{0:T})} \left[\log \frac{q(x_{1:T}|x_0)}{p_\theta(x_{0:T})} \right] \geq \max_{\theta} [-\mathbb{E}_{q(x_0)} [\log p_\theta(x_0)]] \quad (2)$$

where the $q(x_{1:T}|x_0)$ follows a Markov chain that systematically introduces Gaussian noise to the data, following a pre-determined variance schedule $\beta_t \in (0, 1)_{t=1}^T$, named as the forward process or diffusion process. [Ho et al. \(2020\)](#) defines β_t as a sequence of linearly increasing constants from $\beta_1 = 0.0001$ to $\beta_T = 0.02$.

$$q(x_{1:T}|x_0) := \prod_{t=1}^T q(x_t|x_{t-1}), \text{ where } q(x_t|x_{t-1}) \sim \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t \mathbf{I}) \quad (3)$$

Then, by employing the reparameterization trick [Rezende and Mohamed \(2015\)](#), the conditional distribution $q(x_t|x_{t-1})$ can be expressed as a closed form.

$$x_t = \sqrt{1 - \beta_t} \mathbf{x}_{t-1} + \sqrt{\beta_t} \epsilon_{t-1}, \text{ where } \epsilon_{t-1} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

Let $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$,

$$x_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (4)$$

Among them, the linear design of the β_t in [Ho et al. \(2020\)](#) causes most of the values in $\bar{\alpha}_t$ to be concentrated near 0, as shown in the [Figure 3](#), thus causing inefficiencies in the sampling process.

Combining Equation (4), we can simplify the objective function defined by Equation (2) to,

$$\mathbb{E}_{\mathbf{x}_0, \epsilon} [C \|\epsilon_\theta(\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, t) - \epsilon\|], \text{ where } C \text{ is an constant.} \quad (5)$$

However, DDPM [Ho et al. \(2020\)](#) suffers from limitations in inference due to its Markov chain-based framework, leading to computational burden and extended inference times.

By addressing the limitations of DDPM, Denoising Diffusion Implicit Models (DDIM, [Song et al. \(2020a\)](#)) improve inference by utilizing non-Markovian methods and implicit models to reduce the computational burden and enable faster and more efficient inference, without changing the training process. [Song et al. \(2020a\)](#) proposes that comparing the FID score of DDPM, the DDIM model can obtain lower fid values when the number of steps is below 1000, but the DDPM model achieves better results when the number of steps is equal to 1000.

The forward process of Equation (3) from DDPM is that the x_t only rely on x_{t-1} based on Markovian, and the DDIM derives such forward process from Bayes' rule,

$$q_\sigma(x_t|x_{t-1}, x_0) = \frac{q_\sigma(x_{t-1}|x_t, x_0)q_\sigma(x_t|x_0)}{q_\sigma(x_{t-1}|x_0)}$$

where the hyper-parameter σ is utilized to regulate the stochastic nature of the forward process, then $q_\sigma(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ defines as $\mathcal{N}(\sqrt{\bar{\alpha}_{t-1}} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_{t-1} - \sigma_t^2} \frac{\mathbf{x}_t - \sqrt{\bar{\alpha}_t} \mathbf{x}_0}{\sqrt{1 - \bar{\alpha}_t}}, \sigma_t^2 \mathbf{I})$. For the $\sigma_t = 0$ setting in DDIM, the forward process is deterministic except for $t = 1$, given the values of \mathbf{x}_{t-1} and \mathbf{x}_0 .

Then, the reverse process from sample \mathbf{x}_t to the generated sample \mathbf{x}_{t-1} can be defined as (setting the $\sigma = 0$)

$$\mathbf{x}_{t-1} = \sqrt{\bar{\alpha}_{t-1}} \hat{\mathbf{x}}_0 + \sqrt{1 - \bar{\alpha}_{t-1}} \epsilon_\theta(\mathbf{x}_t), \text{ where } \hat{\mathbf{x}}_0 = \frac{\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(\mathbf{x}_t)}{\sqrt{\bar{\alpha}_t}} \quad (6)$$

Meanwhile, there are various methods [Song et al. \(2020c\)](#); [Lu et al. \(2022b\)](#); [Lu et al. \(2022c\)](#) that share a similar motivation with noise-based diffusion models, but they rely on score matching during the reverse process. However, such type of diffusion models are known to have certain limitations, such as difficulty in initial stage learning and limited control over sample generation.

2.2. Image-based diffusion models

Some models use a different approach to train generative models. Instead of predicting noise, they train the models to predict images directly [Bansal et al. \(2022\)](#); [Hoogeboom and Salimans \(2022\)](#). This method simplifies the training process by focusing on learning patterns and structures within input images. By predicting images, these models take advantage of the inherent advantages of images, such as spatial coherence and semantic information, to create high-quality samples. This approach can speed up the learning process, as the initial prediction stages provide meaningful visual information for generating samples. The aim of using image prediction is to capture the essence of the underlying data distribution and generate visually appealing and coherent outputs.

Cold Diffusion [Bansal et al. \(2022\)](#), as the first model to introduce this concept into diffusion models, defines Equation (3) as the broader concept of "degradation", which is defined as \mathbf{x}_0 by operation D with the timestep t , denoted as $\mathbf{x}_t = D(\mathbf{x}_0, t)$.

Correspondingly, the reverse process is also defined as the "restoration", which is an operation to approximately invert operation D , denoted as R , implemented by a neural network with θ as the parameter.

$$R_\theta(\mathbf{x}_t, t) \simeq \mathbf{x}_0.$$

Therefore this objective function becomes,

$$\min_{\theta} \mathbb{E} [\|R_\theta(D(\mathbf{x}_0, t) - \mathbf{x}_0)\|] \quad (7)$$

The sampling for Cold Diffusion can be defined as

$$\mathbf{x}_{t-1} = D(\hat{\mathbf{x}}_0, t - 1).$$

When using noise-based degradation $D(\mathbf{x}_0, t) = \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon$, the sampling process becomes deterministic, which is similar to Equation (6) defined in DDIM [Song et al. \(2020a\)](#),

$$\mathbf{x}_{t-1} = \sqrt{\bar{\alpha}_{t-1}}R_\theta(\mathbf{x}_t, t) + \sqrt{1 - \bar{\alpha}_{t-1}}\hat{\epsilon}(\mathbf{x}_t, t), \quad \hat{\epsilon}(\mathbf{x}_t, t) = \frac{\mathbf{x}_t - \sqrt{\bar{\alpha}_t}R_\theta(\mathbf{x}_t, t)}{\sqrt{1 - \bar{\alpha}_t}} \quad (8)$$

Although image-based diffusion models are useful, they have certain limitations. Specifically, they struggle to capture long-range dependencies and can be computationally complex. Hence, we introduce a novel architecture to eliminate the limitations of noise-based or image-based models while retaining the advantages of both.

3. Methodology

Through examining various categories of diffusion models, it becomes apparent that there are numerous opportunities for enhancing and improving these models. Firstly, proposing a more effective noise scheduler in Section 3.1. Secondly, pursuing the simultaneous estimation of both images and noise in Section 3.2. Lastly, utilizing gradient descent to improve the stability of the sampling process in Section 3.3.

3.1. Noise Scheduler

In the forward process, DDPM, DDIM and noise-based Cold Diffusion all rely on Equation (4) with the square root parameterization. This definition leads to two singularities at $t = 0$ and $t = T$ ($\frac{dx_t}{d\bar{\alpha}} \rightarrow \infty$) and also causes more steps to be required in the reverse diffusion process.

We propose changing the parameterization to:

$$\mathbf{x}_t = \cos(\eta_t)\mathbf{x}_0 + \sin(\eta_t)\epsilon. \text{ where } \eta_t = \frac{t}{T} \frac{\pi}{2}, t \in \{0, 1, 2, \dots, T\}. \quad (9)$$

where mapping the $\sqrt{\bar{\alpha}_t}$ to $\cos(\eta_t)$ and $\sqrt{1 - \bar{\alpha}_t}$ to $\sin(\eta_t)$.

The benefit of our new parameterization approach is to avoid the problem of singularities that existed in previous methods. This has the benefit of allowing us to express the reverse diffusion processes as an ordinary differential equation (ODE) in continuous time, as shown in Fig 1. In turn, this enables us to replace the commonly used Euler step method with higher order ODE solvers such as Runge-Kutta methods.

Recalling the parameterization equation in DDIM defined in Equation (4), there are two singularities at $t = 0$ and $t = T$ when computing the gradient for this formula,

$$\mathbf{x}_t \bar{\alpha}_t = (\sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon)\bar{\alpha}_t = \frac{1}{2\sqrt{\bar{\alpha}_t}}\mathbf{x}_0 - \frac{1}{2\sqrt{1 - \bar{\alpha}_t}}\epsilon,$$

$$\text{When } t = 0, \bar{\alpha}_t = 0, \mathbf{x}_t \bar{\alpha}_t|_{\bar{\alpha}_t=0} \rightarrow +\infty,$$

$$\text{When } t = T, \bar{\alpha}_t = 1, \mathbf{x}_t \bar{\alpha}_t|_{\bar{\alpha}_t=1} \rightarrow -\infty.$$

However, our new parameterization formula defined in Equation (9) can remove the singularities,

$$x_t \eta_t = (\cos(\eta_t)\mathbf{x}_0 + \sin(\eta_t)\epsilon)\eta_t = -\sin(\eta_t)\mathbf{x}_0 + \cos(\eta_t)\epsilon,$$

$$\text{When } t = 0, \eta_t = 0, \mathbf{x}_t \eta_t|_{\eta_t=0} = \epsilon,$$

$$\text{When } t = T, \eta_t = \frac{\pi}{2}, \mathbf{x}_t \eta_t|_{\eta_t=\frac{\pi}{2}} = -\mathbf{x}_0.$$

With the new parameterization defined in Equation (9), we can back-derive an equivalent schedule for $\bar{\alpha}_t$ and β_t .

$$\begin{aligned} \bar{\alpha}_t &= \cos^2(\eta_t) \\ \beta_t &= 1 - \frac{\cos^2(\eta_t)}{\cos^2(\eta_{t-1})} \end{aligned}$$

In Figure 2 and Figure 3, comparing the traditional linear noise schedule Ho et al. (2020) and JSD schedule, which is defined as $\beta_t = \frac{1}{t}$, $t \in \{1, 2, 3, \dots, T\}$, we found that our schedule makes the forward process smoother and does not cause the problem of information loss due to sudden drops as in the linear scheme. In practice, we also found that using this schedule would be beneficial in balancing the contribution of image loss and noise loss, in Section 3.2.




Figure 2: The β_t schedule changes with increasing diffusion steps.




Figure 3: The $\bar{\alpha}_t$ changes with increasing diffusion steps.

3.2. Simultaneous Estimation of Image and Noise

In the objective function, DDPM Ho et al. (2020) and DDIM Song et al. (2020a) train the generative models by estimating the noise, and Cold Diffusion Bansal et al. (2022) fits the generative model by predicting the image directly.

Noise-based diffusion models offer certain advantages, particularly in the later stages of sampling where the process becomes simpler and clearer due to the subtraction of noise. However, starting from pure noise during the initial stages of sampling can pose learning challenges because of the presence of noise-dominated data. In contrast, image-based diffusion models exhibit the opposite characteristics. In the initial stages of sampling, they begin with actual images, which facilitates the learning process by providing meaningful information from the start. However, during the later stages of sampling, the complexity increases as the model needs to manipulate and transform the images to generate new samples. Therefore, by simultaneously predicting both noise and images, we aim to overcome the limitations of each approach and achieve a more controllable sampling process. By utilizing joint prediction, we can harness the benefits offered by both noise-based and image-based diffusion models. This approach provides improved control and flexibility throughout the sampling process, as the Figure 4 shows.

As shown in Figure 4, during the initial steps of the inference stage, the image reconstruction error is lower when estimating both the image and noise, compared to estimating the noise only and using it to derive the image. In the final steps of the inference stage,

both image and noise reconstruction errors are lower when both are estimated together, compared to models that estimate image only or noise only.

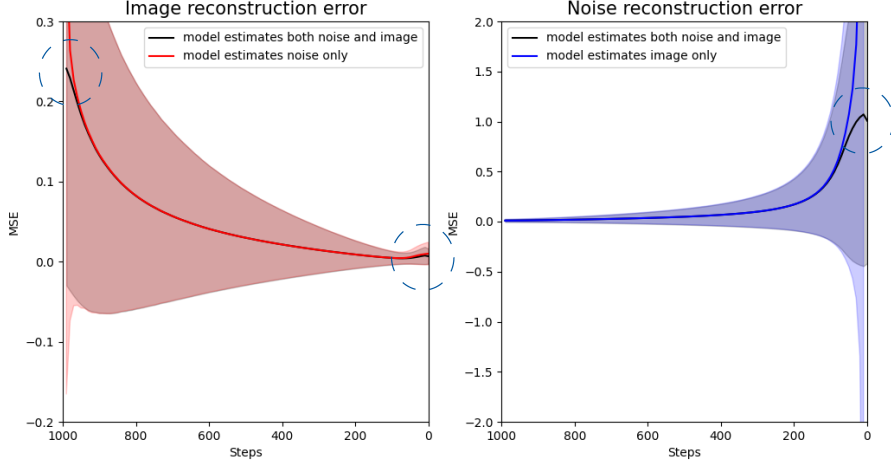


Figure 4: Image reconstruction error (left) and noise reconstruction error (right) at inference for a model that estimates both noise and image, compared to models that estimate only one of these. Lines indicate average MSE and the shaded area indicates the standard variance of MSE.

Hence, by combining Equation (5) and Equation (7), we derive the desired loss function for our model.

$$\min_{\theta} \mathbb{E} [\|\mathbf{R}_{\theta}(\mathbf{x}_t, t) - \mathbf{x}_0\| + \|\epsilon_{\theta}(\mathbf{x}_t, t) - \epsilon\|]$$

3.3. Sampling with Gradient Update

From Figure 1, we can conceptualize the entire diffusion process as the progress of starting from an initial image \mathbf{x}_0 and progressively reducing the level of the image along a curve until obtaining pure noise ϵ . Hence, we can consider discovering an image from noise as an iterative optimization process that implements gradient descent to seek the solution. We take the gradient of Equation (9) to get the ground-truth gradient.

$$\dot{\mathbf{x}}_t = \eta'_t [-\sin(\eta_t)\mathbf{x}_0 + \cos(\eta_t)\epsilon]. \quad (10)$$

Then, the estimated gradient is,

$$\hat{\mathbf{x}}_t = \eta'_t [-\sin(\eta_t)\mathbf{R}_{\theta}(\mathbf{x}_t, t) + \cos(\eta_t)\epsilon_{\theta}(\mathbf{x}_t, t)]. \quad (11)$$

In practice, we have observed that incorporating the gradient of the loss as part of the objective function significantly improves the performance of the model. Then, we obtain the final objective function,

$$\min_{\theta} \mathbb{E} [\|\mathbf{R}_{\theta}(\mathbf{x}_t, t) - \mathbf{x}_0\| + \|\epsilon_{\theta}(\mathbf{x}_t, t) - \epsilon\| + \gamma\|\hat{\mathbf{x}} - \dot{\mathbf{x}}\|] \quad (12)$$

where the γ is used to control the weight of gradient loss.

With the gradient update approach, our sampling process is,

$$\mathbf{x}_{t-1} = \mathbf{x}_t - \Delta_t \hat{\mathbf{x}}_t \quad (13)$$

In practice, we can use techniques like Second Order Runge-Kutta (RK2) and Fourth Order Runge-Kutta (RK4) Hairer et al. (1993), to extend the gradient function defined in Equation (11) to improve performance.

4. Experiments Setup

Baseline Models. In our experiments, we use the noise-based diffusion models DDPM Ho et al. (2020) and DDIM Song et al. (2020a) and the image-based diffusion model Cold Diffusion Bansal et al. (2022) as baseline models. By comparing these two types of models, we can validate the improvements in performance that our contributions and model bring to diffusion models. However, for the Cold Diffusion model Bansal et al. (2022), we found that this model only allows for full-step inference. In other words, if our model is trained with a total of 200 steps, we can only set it to 200 steps during the inference process to achieve optimal performance. Comparing to Equation (6) in DDIM Song et al. (2020a) and Equation (8) in Cold Diffusion Bansal et al. (2022), we can conclude that they are from similar ideas, the only difference is that one is to estimate the noise and the other is to predict the image. Models that predict noise tend to perform better than models that predict images when the sampling steps are large. Therefore, we decided not to include it in the experiment for comparison.

Datasets. We primarily evaluated the performance of our model using three datasets of varying scales, CIFAR-10 (32×32 , Krizhevsky et al. (2009)), CelebA (178×218 , Liu et al. (2015)), and LUSH (256×256 , Yu et al. (2015)). For the UNet architecture model, achieving good results often requires input images to have consistent width and height. To maintain consistency with models’ data preprocessing in DDIM Song et al. (2020a), we performed cropping on the CelebA dataset, resizing the original images from 178×218 to 64×64 Song et al. (2020a). Additionally, in the LUSH dataset, we chose for evaluation a widely used subset of outdoor church images which have dimensions 256×256 .

Evaluation Metrics. We used four evaluation metrics to gauge the models’ performances. The FID Heusel et al. (2017) metric is utilized to determine the divergence between the distribution of generated images and the distribution of actual images. Similarly, sFID Nash et al. (2021) is also used to measure this divergence while maintaining spatial information of the images. Additionally, precision indicates the fraction of the generated images that match reality, and recall measures the coverage of the training data manifold by the generator Kynkäänniemi et al. (2019).

Implementation Details. Our proposed method is built upon the DDIM model Song et al. (2020a) as the base model, and we use the same trained model with $T = 1000$ with the objective function from Equation (5) for each dataset. To mitigate the potential confusion caused by varying iterations during model training, in our ablation experiments, we employ models with parameters approximating the number of iterations used in the pre-training

model. This approach ensures that the compared models have a similar number of iterations, reducing the confounding factor of iteration differences in performance evaluation. In our experiments, we utilized pre-trained models from DDPM for LUSH Church and CIFAR10 datasets [Ho et al. \(2020\)](#), as well as a pre-trained model from DDIM for the CelebA dataset [Song et al. \(2020a\)](#), as references.

5. Experiments Results

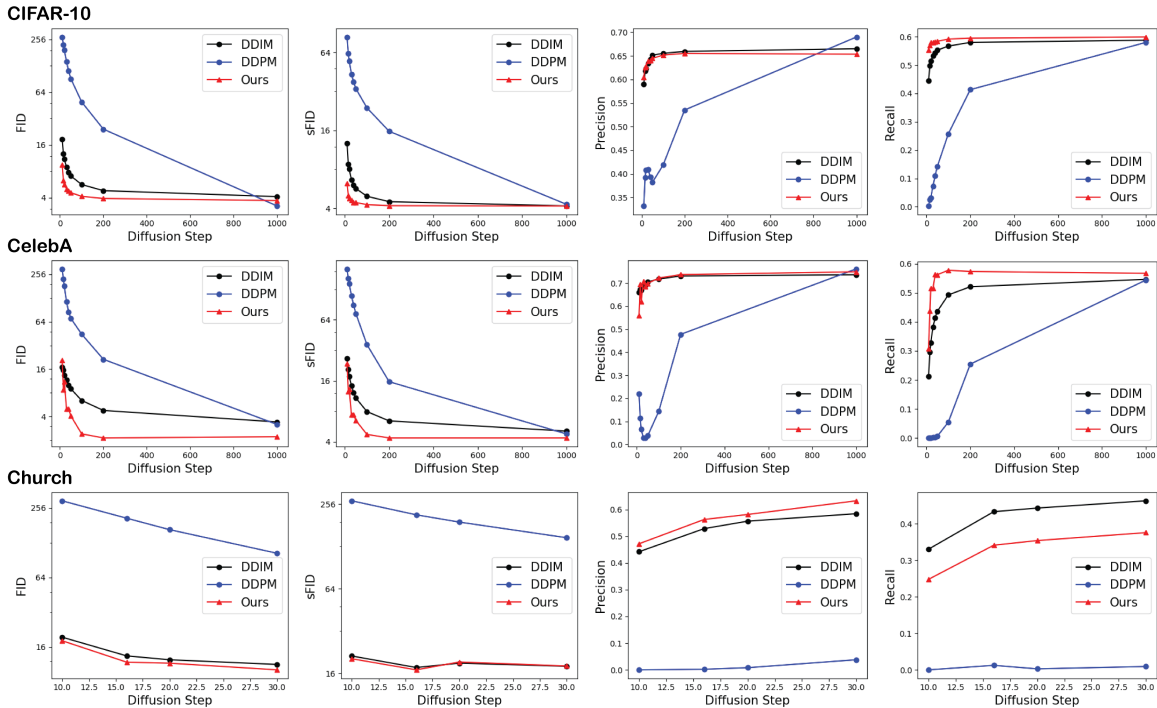


Figure 5: Comparing the performance of DDPM, DDIM and our method in FID \downarrow , sFID \downarrow , Precision \uparrow , and Recall \uparrow evaluation metrics in three different scale datasets, CIFAR-10 (32×32), CelebA (64×64) and LUSH outdoor Church (256×256). **Top:** Measuring the metrics value change with the diffusion steps increase in the CIFAR-10 dataset. **Middle:** CelebA dataset. **Bottom:** LUSH outdoor church.

In Figure 5 and Table 1, we conducted a comparison of sample quality from DDPM [Ho et al. \(2020\)](#), DDIM [Song et al. \(2020a\)](#), and our model trained on CIFAR10, CelebA, and Church datasets. The samples were evaluated based on FID, sFID, precision, and recall, with variation in the number of time steps used for sample generation. In this figure, for the FID score and sFID score, we can find that our model outperforms DDPM and DDIM on all these three datasets, and the improvement is especially noticeable when the steps are between 50 and 200. However, when the number of steps reaches 1000, in the CIFAR-10

dataset, the FID value in DDPM is better than our model, since DDPM tends to get better performance than DDIM when the maximum number of steps for training is the same as the number of steps for inference. For the score of precision and recall, our model can outperform the other two models in both CIFAR-10 and CelebA datasets, which means that the accuracy and diversity of the results predicted by our model are better. For the Church dataset, we observe that its results are similar to those of DDIM. Combined with the experiments in Figure 8, we use these results to show that our model requires less training time to obtain results similar to those of DDIM.




Figure 6: CIFAR-10 (32×32) samples with sampling steps 10, 50, 100, 200.




Figure 7: CelebA (64×64) samples with sampling steps 10, 50, 100, 200.

In Figures 6, 7 and 8, we present the images of CIFAR-10, CelebA, Church sampled in DDPM Ho et al. (2020) model, DDIM Song et al. (2020a) model, and our method using the same number of sampling steps. The results show that, if the sampling steps are 10, the quality of DDPM’s sampling is notably inadequate. DDIM’s sampling quality is considerably better than DDPM, although the image appears slightly blurred. However, our model can achieve clearer outcomes than those of DDIM. Furthermore, we have made an additional noteworthy observation through experiments. As the dimensions of training images increase, the complexity of the training process also escalates. Models trained entirely on noise, like DDIM Song et al. (2020a), and DDPM Ho et al. (2020), require longer training times, e.g., at least 4,432,000 iterations for good performance on the LUSH outdoor Church dataset. However, in our training phase, we improve the performance of the model by simultaneously learning the underlying distribution of the images directly, which can be observed in the Figure 8 that our model only needs 1,135,000 iterations to obtain similar performance to DDIM and DDPM.

In addition, Figure 9 shows the advantage of our model over traditional models in that it can be converted from noise to coarse realistic images much faster. From this figure, the image generated from our model can be clearly seen as the object of a ”horse” by taking

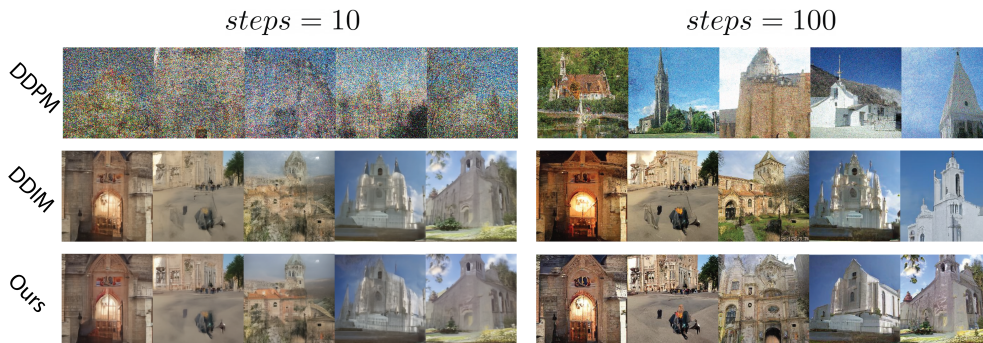


Figure 8: LUSH outdoor Church (256×256) samples with sampling step 10, 100, and **Top:** is DDPM with a pre-trained model with 4432000 iterations, **Middle:** is DDIM with the same pre-trained model with DDPM, **Bottom:** is Our method *without using gradient update* in the sampling process and the model’s iterations are 1135000.



Figure 9: Example of a CIFAR-10 image estimation (\hat{x}_0) over 1000 sampling steps from three models (Ours, DDPM, DDIM). These three results are derived from the same noise input, which leads to different results due to different model parameters and sampling methods.

roughly 150 steps, but the same process in DDIM or DDPM has to spend roughly 400 to 500 steps. This means that during sampling, our model can convert from pure noise to normal images 3 times faster than DDPM and DDIM models.

Moreover, in Figure 10, we presented the results obtained at different sampling steps from our model.

Table 1: To verify the stability of our model, we randomly selected three random seeds and obtained an interval of the FID scores.

Steps	10	20	50	100	1000
CIFAR-10	9.50 ± 0.10	5.62 ± 0.04	4.57 ± 0.02	4.17 ± 0.05	3.70 ± 0.03
CelebA	22.70 ± 0.12	11.00 ± 0.10	4.05 ± 0.05	2.39 ± 0.03	2.20 ± 0.02



Figure 10: Sampling from our model in different sampling step in CIFAR-10 dataset.

Table 2: DDIM+ β^* denotes the first contribution, which proposes a novel noise schedule. DDIM+sin(\cdot) represents the new parameterization method with $\cos(\cdot)\mathbf{x}_0 + \sin(\cdot)\epsilon$. Additionally, the DDIM+ $\hat{\mathbf{x}}_0, \hat{\epsilon}$ indicates the method of estimating noise and image both. Ours represents our final model, which incorporates the contributions (β^* and $\hat{\mathbf{x}}_0, \hat{\epsilon}$) mentioned earlier and utilizes gradient information for image sampling.

Steps	CIFAR-10					CelebA				
	10	16	20	50	100	10	16	20	50	100
DDIM	18.67	12.67	11.06	7.08	5.64	16.92	15.68	13.38	9.00	6.40
DDIM+ β^*	15.04	10.77	9.28	6.00	4.92	23.03	15.23	13.06	5.92	4.24
DDIM+sin(\cdot)	15.97	11.90	9.79	6.22	5.10	24.71	17.18	11.87	5.25	4.22
DDIM+ $\hat{\mathbf{x}}_0, \hat{\epsilon}$	15.03	10.05	8.84	6.25	5.36	14.94	13.01	11.56	8.08	6.11
DDIM+sin(\cdot) + $\hat{\mathbf{x}}_0, \hat{\epsilon}$	13.17	9.77	8.29	5.50	4.50	22.56	13.89	12.80	5.05	3.09
DDIM+ β^* + $\hat{\mathbf{x}}_0, \hat{\epsilon}$	13.69	10.19	8.73	5.80	4.73	23.92	12.58	12.67	4.72	2.83
Ours	9.50	6.32	5.62	4.57	4.17	22.70	8.80	11.00	4.05	2.39

5.1. Ablation Experiments

We show the results of ablation experiments for our proposed three contributions, first is a novel noise schedule, the second is a new parameterization method, which is described in Section 3.1, and the last is a method to estimate both noise and image simultaneously mentioned in Section 3.2. Our ablation experiments are based on modifying the DDIM reference code, and we treat the original DDIM model as a baseline. To compare the performance of each contribution, we use FID as the evaluation metric.

In Table 2, we present the ablation experiments conducted to demonstrate the three contributions of our proposed approach. The baseline model used for comparison is DDIM Song et al. (2020a). From the result, we found that each contribution results in a lower FID Heusel et al. (2017) value compared to the baseline model, which means the generated images from our model are better than the base model in the CIFAR-10 and CelebA datasets. Mean-

while, we obtained further improvements to the model by combining the three contributions discussed above. From the table, it can be observed that compared to models that only utilize individual contributions, the performance is further enhanced by combining $\sin()$ contribution with $\hat{\mathbf{x}}_0, \hat{\epsilon}$ contribution, as well as by combining β^* contribution with $\hat{\mathbf{x}}_0, \hat{\epsilon}$ contribution. However, since both β^* and $\sin()$ contributions aim at optimizing the efficiency and performance of training and sampling for diffusion models by improving the noise schedule, we do not combine them as a comparison model. Furthermore, in our final model, the use of $\sin()$ makes it more convenient to transform the parameterization equation into an ODE, and it is beneficial to apply the gradient approach as the sampling approach. Thus, our final model consists of $\sin()$, $\hat{\mathbf{x}}_0, \hat{\epsilon}$ and a gradient update-based sampling process, which achieves the best result compared to the aforementioned models.

Meanwhile, by comparing the results of DDIM+ $\hat{\mathbf{x}}_0, \hat{\epsilon}$ with those in DDIM+ β^* and DDIM+ $\sin()$, we find this result tends to be better when the sampling steps are less than 20, but when the steps are greater than 20, the impact of DDIM+ $\hat{\mathbf{x}}_0, \hat{\epsilon}$ is lower than the impact of the other two contributions. Based on our analysis, it is evident that using the conventional linear noise schedule in DDIM+ $\hat{\mathbf{x}}_0, \hat{\epsilon}$ model results in the noise contribution outweighing the contribution from the image in the loss function. This is primarily due to the concentration of training on the noisy regions within the linear noise schedule. This is also demonstrated by the results of our DDIM+ $\beta^*+\hat{\mathbf{x}}_0, \hat{\epsilon}$ model, and DDIM+ $\sin()+\hat{\mathbf{x}}_0, \hat{\epsilon}$ model. When we use the new noise schedule in the DDIM+ $\hat{\mathbf{x}}_0, \hat{\epsilon}$ model, this solves the previous problem that performance deteriorates when steps are increased, through balancing the contribution of noise and image components.

6. Conclusion

In this work, our paper introduces two key contributions to the current diffusion methods, resulting in significant improvements in the quality and diversity of generated results. Through applying these two contributions, our model achieves faster generation, with the ability to converge on high-quality images more quickly, and higher quality of the generated images than DDPM, DDIM, and Cold Diffusion models. Our model increases the controllability of the generation process by learning both noise and image information and using gradient information in the inverse diffusion process.

Acknowledgments

This research was supported by The University of Melbourne’s Research Computing Services and the Petascale Campus Initiative.

References

Arpit Bansal, Eitan Borgnia, Hong-Min Chu, Jie S Li, Hamid Kazemi, Furong Huang, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Cold diffusion: Inverting arbitrary image transforms without noise. *arXiv preprint arXiv:2208.09392*, 2022.

- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- Ernst Hairer, Syvert P Nørsett, and Gerhard Wanner. *Solving ordinary differential equations. 1, Nonstiff problems*. Springer-Vlg, 1993.
- 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.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- Emiel Hooeboom and Tim Salimans. Blurring diffusion models. *ArXiv*, abs/2209.05557, 2022.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. *ArXiv*, abs/2009.09761, 2020.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. *Advances in Neural Information Processing Systems*, 32, 2019.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Cheng Lu, Kaiwen Zheng, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Maximum likelihood training for score-based diffusion odes by high order denoising score matching. In *International Conference on Machine Learning*, pages 14429–14460. PMLR, 2022a.
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *ArXiv*, abs/2206.00927, 2022b.
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *ArXiv*, abs/2211.01095, 2022c.
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *arXiv preprint arXiv:2206.00927*, 2022d.

- Charlie Nash, Jacob Menick, Sander Dieleman, and Peter W Battaglia. Generating images with sparse representations. *arXiv preprint arXiv:2103.03841*, 2021.
- Alex Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. *ArXiv*, abs/2102.09672, 2021a.
- Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021b.
- Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015.
- Robin San Roman, Eliya Nachmani, and Lior Wolf. Noise estimation for generative diffusion models. *ArXiv*, abs/2104.02600, 2021.
- Robin Rombach, A. Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10674–10685, 2021.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020a.
- Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020.
- 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*, 2020b.
- Yang Song, Jascha Narain Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *ArXiv*, abs/2011.13456, 2020c.
- Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in Neural Information Processing Systems*, 34: 1415–1428, 2021.
- Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.