

MANIFOLD CONSTRAINT REDUCES EXPOSURE BIAS IN ACCELERATED DIFFUSION SAMPLING

Yuzhe Yao^{1*}, Jun Chen^{2,5}, Zeyi Huang³, Haonan Lin¹, Mengmeng Wang^{4,5†}

Guang Dai⁵, Jingdong Wang⁶

¹Xi'an Jiaotong University, ²Zhejiang Normal University, ³Huawei,

⁴Zhejiang University of Technology, ⁵SGIT AI Lab, State Grid Corporation of China, ⁶Baidu

{yuzheyao.22, junc.change, mengmewang, guang.gdai}@gmail.com
 huangzeyi2@huawei.com, linhaonan@stu.xjtu.edu.cn
 wangjingdong@outlook.com

ABSTRACT

Diffusion models have demonstrated significant potential for generating high-quality images, audio, and videos. However, their iterative inference process entails substantial computational costs, limiting practical applications. Recently, researchers have introduced accelerated sampling methods that enable diffusion models to generate samples with far fewer timesteps than those used during training. Nonetheless, as the number of sampling steps decreases, the prediction errors significantly degrade the quality of generated outputs. Additionally, the exposure bias in diffusion models further amplifies these errors. To address these challenges, we leverage a manifold hypothesis to explore the exposure bias problem in depth. Based on this geometric perspective, we propose a manifold constraint that effectively reduces exposure bias during accelerated sampling of diffusion models. Notably, our method involves no additional training and requires only minimal hyperparameter tuning. Extensive experiments demonstrate the effectiveness of our approach, achieving a FID score of 15.60 with 10-step SDXL on MS-COCO, surpassing the baseline by a reduction of 2.57 in FID.

1 INTRODUCTION

Diffusion models (DMs) Sohl-Dickstein et al. (2015); Ho et al. (2020); Song et al. (2021; 2020); Nichol & Dhariwal (2021) have emerged as powerful generative models for various applications, including image generation (Rombach et al., 2022; Li et al., 2024b), video synthesis (Ho et al., 2022), and audio generation (Kong et al., 2021). DMs are able to generate images of high fidelity and high diversity, when compared to generative adversarial networks (GANs) (Heusel et al., 2017; Goodfellow et al., 2014), variational autoencoders (VAEs) (Kingma & Welling, 2022; Vahdat & Kautz, 2020) and flow models (Lipman et al., 2023; Papamakarios et al., 2017; Kingma et al., 2016). However, DMs necessitate an iterative process to synthesize high-quality results (Ho et al., 2020).

In response to the computational demands of diffusion models, Song et al. (2021) propose an accelerated sampling method. Given a pre-trained diffusion model, during generation, instead of sampling all latent variables on the sampling trajectory, Song et al. (2021) sample along a trajectory shorter than that used during training, significantly accelerating the generation. Nevertheless, Kim et al. (2024) find that during the accelerated generation, the error introduced to the sample at every timestep is proportional to the step size. Therefore, reducing the



Figure 1: Samples generated with 5, 10, 20, 50, and 1k steps (from left to right). Prompt: "A propaganda poster depicting a cat dressed as French emperor Napoleon holding a piece of cheese".

*This work was partly completed during the internship at SGIT AI Lab, State Grid Corporation of China.

†Corresponding Author.

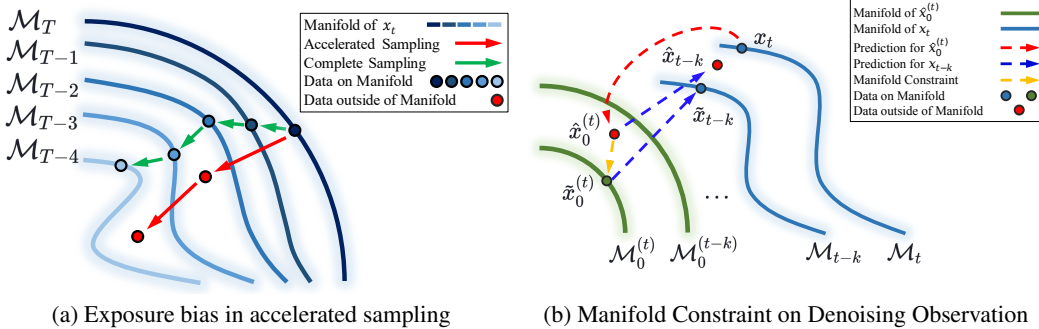


Figure 2: In both (a) and (b), the blue manifolds represent noisy data manifolds. In (a), small step sizes allow diffusion models to transport x_t accurately between manifolds. However, with large step intervals (accelerated sampling), the noisy prediction transports data away from the manifolds, reducing accuracy in the next step. In (b), the green manifolds are manifolds of denoising observations (Equation 6). The yellow arrow represents the manifold constraint, which projects $\hat{x}_0^{(t)}$ toward $\mathcal{M}_0^{(t)}$, allowing error correction in noise estimation and bringing x_{t-k} closer to \mathcal{M}_{t-k} .

number of sampling steps, while improving efficiency, can introduce substantial noise and degrade the sample quality (Rombach et al., 2022). To illustrate how the prediction errors affect sample quality, we utilized Stable Diffusion XL (Podell et al., 2024) to generate images with varying sampling steps. The results, shown in Figure 1, demonstrate that as the number of sampling steps decreases, the generated images exhibit a noticeable loss of detail and increased distortion of objects.

We propose that another issue contributes to the degradation of sample quality during accelerated generation: exposure bias (Schmidt, 2019). This training-inference discrepancy, commonly observed in auto-regressive models, also affects diffusion models (Ning et al., 2023; Li & van der Schaar, 2024; Li et al., 2024a; Ning et al., 2024; Ren et al., 2024). During training, diffusion models are only exposed to ground truth inputs; however, when the input is significantly corrupted (e.g., during accelerated sampling), the models may generate increasingly inaccurate predictions. This cumulative effect is illustrated in Figure 2a.

Li et al. (2024a) evaluate exposure bias using the concept of *pixel variance* within a sample and propose a simple yet effective method based on it. Since pixel variance differs significantly from the *sample variance* used in prior works (Ho et al., 2020; Song et al., 2021; Nichol & Dhariwal, 2021; Dhariwal & Nichol, 2021), we are particularly interested in its implications, especially in accelerated sampling.

To explore this, we collected the pixel variance for different sampling step settings in the generation of Figure 1 and visualized it in Figure 3. The results show that with larger step sizes, the pixel variance deviates significantly from that observed with smaller step sizes. Inspired by prior works (Song et al., 2021; Chung et al., 2022; Chen et al., 2023; Yu et al., 2023; Humayun et al., 2024; Su et al., 2024), we aim to explore the relation between exposure bias and pixel variance of a sample under the manifold hypothesis.

In this work, we focus on addressing exposure bias in the accelerated sampling process of diffusion models. Our contributions are summarized as follows:

- We identify exposure bias as a key challenge in accelerated sampling for diffusion models and introduce the manifold hypothesis, offering a geometric perspective on this issue.




Figure 3: Variation in pixel variance for a single sample across sampling settings: 5, 10, 20, 50, and 1000 steps.

- Based on the manifold hypothesis, we propose a manifold constraint on the denoising observation in diffusion models. This method is training-free, nearly hyperparameter-free, and effectively reduces exposure bias by correcting prediction errors during accelerated sampling.
- Extensive experiments across various image generation tasks and base models demonstrate consistent and substantial improvements with our approach.

2 RELATED WORK

Diffusion Models represent a significant advancement in generative modeling. Initially introduced by Sohl-Dickstein et al. (2015), these models iteratively refine data through denoising. Key enhancements were made by Ho et al. (2020), who incorporated a weighted variational bound in training. Song et al. (2021) introduced a more efficient class of iterative implicit probabilistic models with the same training objective as DDPMs. Further developments include Nichol & Dhariwal (2021), who optimized model architecture. Additionally, Ho & Salimans (2022) offers a method to jointly train a conditional and an unconditional diffusion model. Additionally, works on accelerating sampling methods (Song et al., 2023; Zhou et al., 2023; Chen et al., 2024; Lu et al., 2022a;b; Zhang & Chen, 2023; Liu et al., 2022) have contributed to their practical use in various domains.

Exposure Bias refers to the discrepancy between inputs during training and inference in sequential prediction models (Schmidt, 2019; Wang & Sennrich, 2020; Zhang et al., 2024). This issue also arises in diffusion models, where ground truth inputs are used during training, but previous predictions are fed into the model during inference (Li & van der Schaar, 2024; Ning et al., 2024; 2023; Ren et al., 2024; Yao et al., 2024; Everaert et al., 2024). This discrepancy can lead to cumulative errors in the generation process. Various strategies have been proposed to mitigate exposure bias in diffusion models. Ning et al. (2023) suggest perturbing training data to reflect potential input variations during inference, while Li & van der Schaar (2024) analyze error propagation and recommend using the upper bound of cumulative error as a regularization term during training. Although these approaches enhance robustness, they necessitate model retraining. In contrast, Li et al. (2024a) observe that the pixel variance of noisy samples is not always aligned with the schedule during inference, prompting them to adjust the timestep accordingly. Additionally, Ning et al. (2024) identify noise prediction error as a key source of exposure bias and propose scaling noise estimation to address magnitude discrepancies, although direction errors remain unaddressed.

3 BACKGROUND

Diffusion Models In the forward diffusion process, Gaussian noise is added to samples from the data distribution $q(x_0)$, over a series of time steps $t \in \{0, 1, \dots, T\}$, resulting in increasingly noisy versions of the data:

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}), \quad (1)$$

where β_t is a variance schedule that controls the amount of noise added at each step. The overall forward process is described as:

$$q(\mathbf{x}_{1:T}|\mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1}). \quad (2)$$

In the reverse process, diffusion models denoise \mathbf{x}_T back to the original data distribution:

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \boldsymbol{\Sigma}), \quad (3)$$

where $\boldsymbol{\mu}_\theta$ is the mean predicted by a neural network with parameters θ , and $\boldsymbol{\Sigma}$ is the predicted variance. Ho et al. (2020) found that in practice, using a fixed variance schedule $\boldsymbol{\Sigma}_t$ leads to simpler and more stable training. Diffusion models can be trained by minimizing the Kullback-Leibler divergence:

$$\mathcal{L} = \mathbb{E}_{q(\mathbf{x}_{0:T})} \left[\sum_{t=1}^T \text{KL}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) \| p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)) \right]. \quad (4)$$

An alternative, often simpler, loss function involves predicting the added noise directly:

$$\mathcal{L}_{simple} = \mathbb{E}_{t, \mathbf{x}_0, \epsilon_t} \left[\|\epsilon_t - \epsilon_\theta^{(t)}(\mathbf{x}_t)\|^2 \right], \quad (5)$$

where ϵ_θ denotes the noise estimation network in diffusion model.

After training, diffusion models can generate data by sampling from the standard normal prior $p(\mathbf{x}_T)$, iteratively refining these samples through the reverse process defined in Equation 3. When employing the DDIM sampler (Song et al., 2021), the model proceeds as follows at each timestep t .

First, the model predicts the *denoising observation* using noise estimation:

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

where $\hat{\mathbf{x}}_0^{(t)}$ represents the estimated clean image at timestep t , \mathbf{x}_t is the noisy image at timestep t , and $\epsilon_\theta^{(t)}(\mathbf{x}_t)$ is the model’s prediction of noise.

Next, the latent variable for timestep $t - 1$ is predicted:

$$\hat{\mathbf{x}}_{t-1} = \sqrt{\bar{\alpha}_{t-1}} \hat{\mathbf{x}}_0^{(t)} + \sqrt{1 - \bar{\alpha}_{t-1} - \sigma_t^2} \epsilon_\theta^{(t)}(\mathbf{x}_t) + \sigma_t \epsilon'_t, \quad \epsilon'_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (7)$$

where $\alpha_t := 1 - \beta_t$, $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$, and σ_t controls the stochasticity of the generation. The term ϵ'_t is sampled from a standard normal distribution, introducing stochasticity into the generation process.

Prediction Error in Accelerated Sampling Given a diffusion model trained with a forward process of T steps, during generation, instead of sampling all latent variables $\{\mathbf{x}_T, \mathbf{x}_{T-1}, \dots, \mathbf{x}_1\}$ iteratively, Song et al. (2021) propose to sample along a trajectory of length T' : $\{\mathbf{x}_{\tau_{T'}}, \mathbf{x}_{\tau_{T'-1}}, \dots, \mathbf{x}_{\tau_1}\}$, where τ is defined as an increasing subsequence of $\{0, 1, \dots, T\}$. Since $T' < T$, the proposed method significantly accelerates the sampling.

However, the prediction error tends to increase as the sampling interval grows. Kim et al. (2024) derive that the error introduced to the latent variable \mathbf{x}_t during the transition from time t to $t - \Delta t$:

$$\mathbf{e}_t = \frac{\Delta t}{t} \mathcal{O}(\Delta t), \quad (8)$$

where Δt denotes the timestep interval. Equation 8 shows an inherent trade-off between accuracy and efficiency in the sampling process in diffusion models.

Exposure Bias in Diffusion Models During training, diffusion models are trained on noisy data sampled from the real data distribution \mathbf{x}_0 :

$$q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I}). \quad (9)$$

On the contrary, models are only exposed to their previously prediction $\hat{\mathbf{x}}_t$ during inference. The input discrepancy, $\mathbf{x}_t - \hat{\mathbf{x}}_t$, disrupts the model’s prediction in subsequent timesteps.

The exposure bias is exacerbated in accelerated sampling, where errors at each timestep are significantly amplified, as shown in Equation 8. Therefore, we identify the combined effect of exposure bias and enlarged prediction error as the main challenge in accelerated diffusion model sampling.

4 METHOD

Inspired by previous work (Song & Ermon, 2019; Chung et al., 2022; Chen et al., 2024; He et al., 2024), we are interested in understanding the exposure bias problem under the geometric view of diffusion models. In this section, we first provide the geometric view of the exposure bias problem. Then, based on the geometrical analysis, we introduce *Manifold Constraint on Denoising Observation* (MCDO), a training-free and almost hyperparameter tuning-free method that mitigates the exposure bias in the accelerated sampling of diffusion models.

4.1 GEOMETRIC VIEW OF EXPOSURE BIAS

To better understand the training-inference discrepancy in diffusion models, we borrow a geometric view of the data manifold.

Assumption 1 (Chung et al., 2022) *Let $\mathcal{M}_0 \subset \mathbb{R}^n$ be the clean data manifold. The distribution of the ground truth noisy data $p(\mathbf{x}_t)$ is concentrated on an $(n - 1)$ -dimensional manifold \mathcal{M}_t .*

Under Assumption 1, for a diffusion model trained with T sampling steps, its ground truth noisy samples obtained with Equation 9 form a series of manifolds $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_T$. The model is trained to smoothly map the data from \mathcal{M}_t to \mathcal{M}_{t-1} (Chung et al., 2022). During the inference, pure noises from \mathcal{M}_T are gradually mapped towards the target manifold \mathcal{M}_0 . The perfect generation process can be formulated as follows:

$$\mathcal{M}_T \xrightarrow{\Lambda_T} \dots \xrightarrow{\Lambda_3} \mathcal{M}_2 \xrightarrow{\Lambda_2} \mathcal{M}_1 \xrightarrow{\Lambda_1} \mathcal{M}_0, \quad (10)$$

where $\Lambda_i, i \in \{1, 2, \dots, T\}$, denotes a series of smooth maps.

However, the perfect mapping is only applicable to data concentrated on the manifold.

Assumption 2 (Extended from He et al. (2024)) *Define the distance from a data point to a given manifold as: $d(\mathbf{x}, \nu, \mathcal{M}) := \inf_{\mathbf{z} \in \mathcal{M}} \|\mathbf{x} - \nu \mathbf{z}\|$, for $\nu > 0$. For a given timestep t , diffusion models only make accurate noise predictions for data concentrated on the manifold: $\mathbb{X}_t := \{\mathbf{x}_t \in \mathbb{R}^n \mid d(\mathbf{x}_t, 1, \mathcal{M}_t) < r_t\}$, where $r_t > 0$ is the radius.*

While prediction errors are inevitable, it is possible that, for a given timestep $t + 1$, data from \mathcal{M}_{t+1} is mapped to points away from \mathcal{M}_t such that $d(\mathbf{x}_t, 1, \mathcal{M}_t) > r_t$. With Assumption 2, models' noise estimation for data $\hat{\mathbf{x}}_t \notin \mathbb{X}_t$ will degrade. Based on Equation 8, we argue that the large deviation from the manifold is the main reason for exposure bias in the accelerated sampling of diffusion models. We illustrate this problem in Figure 2a.

4.2 RELATION BETWEEN DISTANCE TO MANIFOLD AND THE PIXEL VARIANCE OF DATA

Under Assumptions 1 and 2, we relate the pixel variance of a sample $\hat{\mathbf{x}}_t$, to its L^2 distance to \mathcal{M}_t .

For a prediction $\hat{\mathbf{x}}_t$, let $\mathbf{x}_t^* = \operatorname{argmin}_{\mathbf{z}_t \in \mathcal{M}_t} \|\hat{\mathbf{x}}_t - \mathbf{z}_t\|$. The distance between $\hat{\mathbf{x}}_t$ and \mathcal{M}_t satisfies:

$$d(\hat{\mathbf{x}}_t, 1, \mathcal{M}_t) = \|\hat{\mathbf{x}}_t - \mathbf{x}_t^*\| \geq \|\|\mathbf{x}_t^*\| - \|\hat{\mathbf{x}}_t\|\|. \quad (11)$$

Given that $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, all pixels in \mathbf{x}_T follows a normal distribution, such that $\mathbb{E}[\frac{1}{n} \sum_{i=0}^n \mathbf{x}_T] = 0$. According to Equation 4, we have $\mathbb{E}[\frac{1}{n} \sum_{i=0}^n \boldsymbol{\epsilon}_t] = 0$, when the diffusion models are well-trained. As a consequence, the expectation of pixel mean of a single noisy sample approaches 0 as well: $\mathbb{E}[\frac{1}{n} \sum_{i=0}^n \mathbf{x}_t] = 0$, according to Equations 6 and 7 (more details in Appendix A.1). Then the L^2 norm of $\hat{\mathbf{x}}_t$ and its pixel variance satisfy:

$$\operatorname{Var}(\hat{\mathbf{x}}_t) = \frac{1}{n} \sum_{i=1}^n (\hat{\mathbf{x}}_{t_i} - \frac{1}{n} \sum_{j=1}^n \hat{\mathbf{x}}_{t_j})^2 \approx \frac{1}{n} \sum_{i=1}^n \hat{\mathbf{x}}_{t_i}^2 = \frac{1}{n} \|\hat{\mathbf{x}}_t\|^2, \quad (12)$$

where n is the data dimension. We show the evolution of L^2 norm of a sample as well as its pixel variance in Figure 4. It is shown that the curve of the L^2 norm and pixel variance are highly overlapped for most of the time. Therefore, we use Equation 12 to approximate the L^2 norm of noisy data with its pixel variance.

Plugging Equations 9 and 12 into Equation 11, we have:

$$d(\hat{\mathbf{x}}_t, 1, \mathcal{M}_t) \gtrsim \left| \|\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}_t\| - \sqrt{n \operatorname{Var}(\hat{\mathbf{x}}_t)} \right|. \quad (13)$$

According to Wegner (2024), for data $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\boldsymbol{\epsilon} \in \mathbb{R}^n$, one has:

$$|\mathbb{E}[\|\boldsymbol{\epsilon}\|] - \sqrt{n}| \leq \frac{1}{\sqrt{n}}. \quad (14)$$

Given that the dimension n of the latent space is usually large for diffusion models (e.g., $n = 12288$ for LDM-4 trained on CelebA-HQ), with high probability, $\|\boldsymbol{\epsilon}_t\| \approx \sqrt{n}$.

In addition, since ϵ_t is a random noise independent of \mathbf{x}_0 , when n is large, ϵ_t and \mathbf{x}_0 are almost orthogonal with high probability (Wegner, 2024). Therefore, we have:

$$\mathbb{E}_{\mathbf{x}_0 \sim q(\mathbf{x}_0), \epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\|\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon_t\|^2] = \mathbb{E}_{\mathbf{x}_0 \sim q(\mathbf{x}_0), \epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\|\sqrt{\bar{\alpha}_t} \mathbf{x}_0\|^2 + \|\sqrt{1 - \bar{\alpha}_t} \epsilon_t\|^2]. \quad (15)$$

Plugging Equations 12, 14, and 15 into Equation 13, the following inequality holds with high probability:

$$d(\hat{\mathbf{x}}_t, 1, \mathcal{M}_t) \gtrsim \left| \sqrt{\bar{\alpha}_t \|\mathbf{x}_0\|^2 + n(1 - \bar{\alpha}_t)} - \sqrt{n \text{Var}(\hat{\mathbf{x}}_t)} \right|, \quad (16)$$

where equality holds if and only if $\hat{\mathbf{x}}_t$, \mathbf{x}_0^* and the origin are collinear.

Equation 16 indicates that the deviation in $\text{Var}(\hat{\mathbf{x}}_t)$ could potentially make $d(\hat{\mathbf{x}}_t, 1, \mathcal{M}_t) > r_t$, resulting in exposure bias problem under Assumption 2. This is consistent with Figures 1 and 3.

From this perspective, the timestep-shifting strategy proposed by Li et al. (2024a), which determines timestep t_s within range w : $t_s = \text{argmin}_{t'} \left(\sqrt{1 - \bar{\alpha}_{t'}} - \sqrt{\text{Var}(\hat{\mathbf{x}}_{t'})} \right)$, $t' \in \{t - w/2, \dots, t + w/2\}$, can be interpreted as locating the manifold \mathcal{M}_i , $i \in \{t - w/2, \dots, t + w/2\}$, where the lower bound of samples-to-manifold distance is potentially reduced according to Equation 16.

As this strategy provides valuable insights, the inaccessibility of the L^2 norm of the clean data \mathbf{x}_0 in Equation 16 presents an opportunity for further refinement of the approach.

Inspired by prior studies (Bao et al., 2022; Chen et al., 2023; 2024), we delve into the model’s predictions for noisy data and the corresponding denoising observations. Specifically, when the model is given the ground truth input at $t + 1$, the analytical form of $\hat{\mathbf{x}}_0^{(t)}$ and $\hat{\mathbf{x}}_t$ can be expressed as (see Appendix A.2):

$$\hat{\mathbf{x}}_0^{(t)} = \mathbf{x}_0 + d_t \mathbf{e}_\theta^{(t)}, \hat{\mathbf{x}}_t = \mathbf{x}_t + n_t \mathbf{e}_\theta^{(t)}, \quad (17)$$

Here, \mathbf{x}_0 denotes the clean data, \mathbf{x}_t the ground truth noisy data, $\mathbf{e}_\theta^{(t)}$ the noise estimation error. The scaling factors d_t and n_t vary with t . We find that while n_t remains within a small range near 0, d_t is hundreds or even thousands of times larger than n_t when t is large (see Figure 5).

This substantial magnitude of d_t suggests that $\hat{\mathbf{x}}_0^{(t)}$ is predominantly influenced by the scaled noise prediction error term $d_t \mathbf{e}_\theta^{(t)}$ at most timesteps. As a consequence, we propose to rectify the noise estimation error $\mathbf{e}_\theta^{(t)}$ by incorporating information from the denoising observation $\hat{\mathbf{x}}_0^{(t)}$.

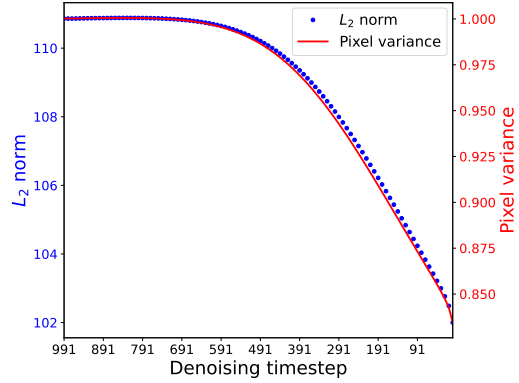


Figure 4: Evolution of L^2 norm and pixel variance of the noisy data. We plot the average statistics of 6,400 samples in 100-steps generation on CelebA-HQ.

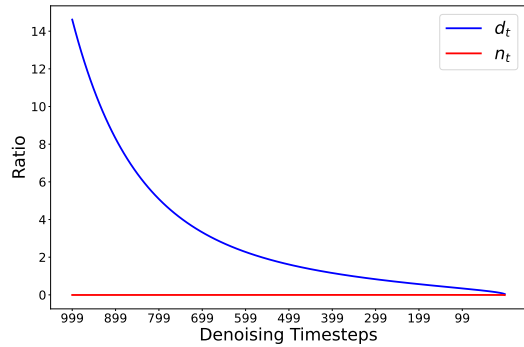


Figure 5: Comparison between error scale on denoising observation (d_t) and on noisy data (n_t).

4.3 MANIFOLD CONSTRAINT ON DENOISING OBSERVATION

Chung et al. (2022) view the diffusion process as an interpolation between \mathcal{M}_0 and \mathcal{M}_T . Inspired by this, we substitute $\hat{\epsilon}_\theta^{(t)}(\mathbf{x})$ in Equation 7 with Equation 6:

$$\left(\frac{\sqrt{1-\bar{\alpha}_{t-1}-\sigma_t^2}\sqrt{\bar{\alpha}_t}-\sqrt{\bar{\alpha}_{t-1}}}{\sqrt{1-\bar{\alpha}_t}}\hat{\mathbf{x}}_0^{(t)}\right)=\frac{\sqrt{1-\bar{\alpha}_{t-1}-\sigma_t^2}}{\sqrt{1-\bar{\alpha}_t}}\mathbf{x}_t-\hat{\mathbf{x}}_{t-1}+\sigma_t\epsilon'_t, \epsilon'_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (18)$$

Equation 18 suggests that $\hat{\mathbf{x}}_0^{(t)}$ can be interpreted as interpolating between two noisy data. Based on this, we assume that the denoising observation $\hat{\mathbf{x}}_0^{(t)}$ also concentrates on a series of manifolds.

Assumption 3 (Extended from Chung et al. (2022)) Denote $\mathbb{X}_0^{(t)}$ as the set of denoising observations of \mathbf{x}_t , $\mathbf{x}_t \in \mathbb{X}_t$. There exists a series of manifolds: $\mathcal{M}_0^{(T)}, \mathcal{M}_0^{(T-1)}, \dots, \mathcal{M}_0^{(1)}, \mathcal{M}_0^{(0)}$, where $\mathbb{X}_0^{(t)}$, for $t \in \{T, T-1, \dots, 0\}$, concentrates on: $d(\hat{\mathbf{x}}_0^{(t)}, 1, \mathcal{M}_0^{(t)}) < r_0^{(t)}, \forall \hat{\mathbf{x}}_0^{(t)} \in \mathbb{X}_0^{(t)}$. Where $r_0^{(t)} > 0$ is the radius.

Under Assumption 3, when a noisy data $\hat{\mathbf{x}}_t$ is distant from \mathcal{M}_t , its denoising observation, calculated using Equation 6, could potentially be mapped away from $\mathcal{M}_0^{(t)}$. For clarity, we denote denoising observations close to $\mathcal{M}_0^{(t)}$ as $\hat{\mathbf{x}}_0^{(t)}$, and those distant from $\mathcal{M}_0^{(t)}$ as $\hat{\mathbf{x}}_0^{\prime(t)}$.

For a denoising observation $\hat{\mathbf{x}}_0^{\prime(t)} \notin \mathbb{X}_0^{(t)}$. Let $\hat{\mathbf{x}}_0^{*(t)} := \operatorname{argmin}_{\mathbf{z}_0^{(t)} \in \mathcal{M}_0^{(t)}} \|\hat{\mathbf{x}}_0^{\prime(t)} - \mathbf{z}_0^{(t)}\|$.

As $\hat{\mathbf{x}}_0^{\prime(t)}$ is a linear combination of \mathbf{x}_t and $\epsilon_\theta^{(t)}$ according to Equation 6, the expectation of pixel mean for both $\hat{\mathbf{x}}_0^{\prime(t)}$ and $\hat{\mathbf{x}}_0^{*(t)}$ is 0. We substitute $\hat{\mathbf{x}}_t$ and \mathbf{x}_t^* in Equation 11 with $\hat{\mathbf{x}}_0^{\prime(t)}$ and $\hat{\mathbf{x}}_0^{*(t)}$:

$$d(\hat{\mathbf{x}}_0^{\prime(t)}, 1, \mathcal{M}_0^{(t)}) \geq \|\hat{\mathbf{x}}_0^{\prime(t)}\| - \|\hat{\mathbf{x}}_0^{*(t)}\| \approx \sqrt{n} \left| \sqrt{\operatorname{Var}(\hat{\mathbf{x}}_0^{\prime(t)})} - \sqrt{\operatorname{Var}(\hat{\mathbf{x}}_0^{*(t)})} \right|. \quad (19)$$

Equation 19 indicates that the deviated $\hat{\mathbf{x}}_0^{\prime(t)}$ can be constrained to approach $\mathcal{M}_0^{(t)}$, when its pixel variance is scaled to $\operatorname{Var}(\hat{\mathbf{x}}_0^{*(t)})$.

Given that $\hat{\mathbf{x}}_0^{*(t)}$ is inaccessible in generation process, we introduce the *reference pixel variance*: $v_t = \mathbb{E}_{\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), y \sim \mathbb{Y}} \operatorname{Var}(\hat{\mathbf{x}}_0^{(t)} | \mathbf{x}_T, y)$, where \mathbb{Y} denotes the set of conditions. We approximate: $\operatorname{Var}(\hat{\mathbf{x}}_0^{*(t)}) \approx v_t, \forall \hat{\mathbf{x}}_0^{*(t)} \in \mathcal{M}_0^{(t)}$. In practice, the pixel variance used for estimating v_t is collected during generation with large sampling steps (e.g., $T = 1000$), such that the denoising observations are close to their corresponding manifolds. The pre-computation process is detailed in Algorithm 1.

During accelerated sampling, we first align the pixel variance of $\hat{\mathbf{x}}_0^{\prime(t)}$ to approach v_t :

$$c_t = \operatorname{Var}(\hat{\mathbf{x}}_0^{\prime(t)}), \quad \tilde{\mathbf{x}}_0^{(t)} = \sqrt{\frac{v_t}{c_t}} \left[\hat{\mathbf{x}}_0^{\prime(t)} - \operatorname{Mean}(\hat{\mathbf{x}}_0^{\prime(t)}) \right] + \operatorname{Mean}(\hat{\mathbf{x}}_0^{\prime(t)}). \quad (20)$$

Then, we correct the noise estimation, by inverting Equation 6:

$$\tilde{\epsilon}_\theta^{(t)}(\mathbf{x}_t) = \frac{\mathbf{x}_t - \sqrt{\bar{\alpha}_t} \tilde{\mathbf{x}}_0^{(t)}}{\sqrt{1-\bar{\alpha}_t}}. \quad (21)$$

By substituting Equations 20 and 21 into Equation 7, we obtain a refined prediction for \mathbf{x}_{t-1} :

$$\tilde{\mathbf{x}}_{t-1} = \sqrt{\bar{\alpha}_{t-1}} \tilde{\mathbf{x}}_0^{(t)} + \sqrt{1-\bar{\alpha}_{t-1}-\sigma_t^2} \tilde{\epsilon}_\theta^{(t)}(\mathbf{x}_t) + \sigma_t \epsilon'_t, \epsilon'_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (22)$$

The correction is applied for $t \in \{T, \dots, t_{thre}\}$, where t_{thre} is a timestep threshold. The proposed method is illustrated in Figure 2b and detailed in Algorithm 2.

We refer to the proposed method as *Manifold Constraint on Denoising Observation* (MCDO), as its key idea is to constrain the denoising observation to approach its corresponding manifold.

A distinction between reducing the lower bound of $d(\hat{\mathbf{x}}_t, 1, \mathcal{M}_t)$ by scaling $\text{Var}(\hat{\mathbf{x}}_t)$ and that of $d(\hat{\mathbf{x}}_0^{(t)}, 1, \mathcal{M}_0^{(t)})$ by correcting $\text{Var}(\mathbf{x}_0^{(t)})$ is that the former modifies \mathbf{x}_t , while the latter refines the noise estimation when \mathbf{x}_t is preserved (see Appendix A.3).

Estimating v_t requires few samples (e.g., $N = 20$ for SDXL on MS-COCO). Once computed, v_t can be directly used to correct $\text{Var}(\hat{\mathbf{x}}_0^{(t)})$ during accelerated sampling *without additional computational overhead*. Since v_t depends only on the pre-trained model, when accurately estimated, the only tunable hyperparameter is t_{thre} , which often works well at $t_{thre} = 0$. Thus, DDIM-MCDO requires minimal hyperparameter tuning.

5 EXPERIMENTS

5.1 EXPERIMENTS SETUP

We conduct experiments on high-resolution datasets across various tasks. Specifically, we evaluate MS-COCO (Lin et al., 2014) with SDXL (Podell et al., 2024), and CelebA-HQ (Karras et al., 2018), ImageNet-256 (Rusakovsky et al., 2015), and LSUN-Bedroom (256×256) with LDM-4 (Rombach et al., 2022).

For text-to-image (T2I) generation, we generate 30k samples using MS-COCO prompts to compute FID (Heusel et al., 2017) and CLIP scores (Radford et al., 2021). For other tasks, 50k images are generated to evaluate FID, with sFID and IS reported for ImageNet. We compare our method with state-of-the-art training-free approaches (Ning et al., 2024; Li et al., 2024a), following Li et al. (2024a) to implement TS-DDIM and conducting a hyperparameter search based on their range and qualitative results. For Ning et al. (2024), a coarse-to-fine search is used for the uniform schedule λ . We visualize exposure bias reduction using Ning et al. (2024)’s metric (see Appendix A.4). Results with DPM-Solver++ (Lu et al., 2022b) and Stable Diffusion 3 (Esser et al., 2024) are in Appendices A.12 and A.13.

5.2 MAIN RESULTS ON TEXT-TO-IMAGE GENERATION

Results of T2I generation (Table 1, Figure 6) show that with few sampling steps, image details and object shapes degrade significantly, suggesting deviation from the data manifold. The proposed manifold constraint improves image quality, indicating more accurate inference.

To implement MCDO, we use $N = 20, T = 1000$ for pre-computation. Samples are conditioned on different prompts from MS-COCO. While a larger N could improve v_t estimation, we keep it fixed for simplicity. For TS-DDIM (Li et al., 2024a) we use $t_c = 300, w = 4$ for cutoff timestep and window size. A uniform epsilon scale: $\lambda = 1.008$ is applied for DDIM-ES (Ning et al., 2024).

Algorithm 1 Pre-Computation for v_t

```

1: Initialize  $n_t = \text{list}()$ 
2: for  $i = 1, 2, \dots, N$  do
3:    $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), y \sim \mathbb{Y}$ 
4:   for  $t = T, \dots, 2, 1$  do
5:      $\hat{\mathbf{x}}_0^{(t)} = \frac{\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}_\theta^{(t)}(\mathbf{x}_t, y)}{\sqrt{\bar{\alpha}_t}}$ 
6:      $n_t.append(\text{Var}(\hat{\mathbf{x}}_0^{(t)}))$ 
7:      $\vec{\mathbf{x}}_t = \sqrt{1 - \bar{\alpha}_{t-1} - \sigma_t^2} \boldsymbol{\epsilon}_\theta^{(t)}(\mathbf{x}_t, y)$ 
8:      $\boldsymbol{\epsilon}'_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
9:      $\mathbf{x}_{t-1} = \sqrt{\bar{\alpha}_{t-1}} \hat{\mathbf{x}}_0^{(t)} + \vec{\mathbf{x}}_t + \sigma_t \boldsymbol{\epsilon}'_t$ 
10:  end for
11: end for
12: for  $t = T, \dots, 2, 1$  do
13:    $v_t = \text{Mean}(n_t)$ 
14: end for

```

Algorithm 2 MCDO Sampling

```

1:  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), y \sim \mathbb{Y}$ 
2: for  $i = T', T' - 1, \dots, 1$  do
3:    $\hat{\mathbf{x}}_0^{(\tau_i)} = \frac{\mathbf{x}_{\tau_i} - \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}_\theta^{(\tau_i)}(\mathbf{x}_{\tau_i}, y)}{\sqrt{\bar{\alpha}_{\tau_i}}}$ 
4:   if  $\tau_i \geq t_{thre}$  then
5:      $\mu_{\tau_i} = \text{Mean}(\hat{\mathbf{x}}_0^{(\tau_i)})$ 
6:      $c_{\tau_i} = \text{Var}(\hat{\mathbf{x}}_0^{(\tau_i)})$ 
7:      $\tilde{\mathbf{x}}_0^{(\tau_i)} = \sqrt{\frac{v_{\tau_i}}{c_{\tau_i}}} (\hat{\mathbf{x}}_0^{(\tau_i)} - \mu_{\tau_i}) + \mu_{\tau_i}$ 
8:      $\hat{\boldsymbol{\epsilon}}_\theta^{(\tau_i)} = \frac{\mathbf{x}_{\tau_i} - \sqrt{\bar{\alpha}_{\tau_i}} \tilde{\mathbf{x}}_0^{(\tau_i)}}{\sqrt{1 - \bar{\alpha}_{\tau_i}}}$ 
9:   else
10:     $\hat{\boldsymbol{\epsilon}}_\theta^{(\tau_i)} = \hat{\boldsymbol{\epsilon}}_\theta^{(\tau_i)}(\mathbf{x}_{\tau_i}, y), \tilde{\mathbf{x}}_0^{(\tau_i)} = \hat{\mathbf{x}}_0^{(\tau_i)}$ 
11:   end if
12:    $\vec{\mathbf{x}}_{\tau_i} = \sqrt{1 - \bar{\alpha}_{\tau_i-1} - \sigma_\tau^2} \hat{\boldsymbol{\epsilon}}_\theta^{(\tau_i)}$ 
13:    $\boldsymbol{\epsilon}'_\tau \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
14:    $\tilde{\mathbf{x}}_{\tau_i-1} = \sqrt{\bar{\alpha}_{\tau_i-1}} \tilde{\mathbf{x}}_0^{(\tau_i)} + \vec{\mathbf{x}}_{\tau_i} + \sigma_\tau \boldsymbol{\epsilon}'_\tau$ 
15: end for

```

Table 1: Results on Text-to-Image generation (1024×1024) on MS-COCO with SDXL.

Model	Method	Steps	FID↓	CLIP Score↑	t_{thre}
SDXL	DDIM	10	18.17	31.58	–
	DDIM-ES Ning et al. (2024)	10	18.91	31.59	–
	TS-DDIM Li et al. (2024a)	10	23.82	31.31	–
	DDIM-MCDO	10	15.60	31.75	0



(a) "Three people on horse back at a rural road intersection."



(b) "A baby next to a stuffed bear of some sort."

Figure 6: Images generated using 10 steps SDXL with DDIM sampler. From left to right: SDXL, TS-DDIM (Li et al., 2024a), LDM-ES (Ning et al., 2024), and DDIM-MCDO (ours).

5.3 RESULTS ON UNCONDITIONAL GENERATION

Considering that both LDM-ES (Ning et al., 2024) and DDIM-MCDO aim to reduce noise estimation error, we investigate their orthogonality. We conduct experiments combining DDIM-MCDO with LDM-ES using the same hyperparameters, denoted as DDIM-MCDO[†].¹

We compare DDIM-MCDO and DDIM-MCDO[†] with state-of-the-art methods (Ning et al., 2024; Li et al., 2024a) on the CelebA-HQ dataset (Karras et al., 2018) using LDM-4. Results under the recommended setting ($\eta = 0.0$) (Rombach et al., 2022) are shown in Table 2, with λ representing the *epsilon scaling factor* from Ning et al. (2024). For TS-DDIM (Li et al., 2024a), we set $t_c = 100$, $w = 60$ for 20 and 10 steps sampling, and $t_c = 200$, $w = 30$ for 5 steps. Qualitative results are in Appendix A.6.

DDIM-MCDO significantly enhances performance across all three accelerated generation settings, attributed to its correction of both magnitude

Table 2: Results on CelebA-HQ 256×256 with DDIM Sampler. $\eta = 0.0$. Bold indicates the best results, and underlined indicates the second-best.

Method	T'	FID↓	λ
DDIM	20	10.59	–
TS-DDIM (Li et al., 2024a)	20	8.29	–
LDM-ES (Ning et al., 2024)	20	18.39	1.001
LDM-ES (Ning et al., 2024)	20	11.73	0.999
DDIM-MCDO	20	<u>9.49</u>	–
DDIM-MCDO [†]	20	10.59	1.001
DDIM-MCDO [†]	20	10.53	0.999
DDIM	10	21.08	–
TS-DDIM (Li et al., 2024a)	10	15.71	–
LDM-ES (Ning et al., 2024)	10	18.56	1.002
DDIM-MCDO	10	<u>15.63</u>	–
DDIM-MCDO [†]	10	13.36	1.002
DDIM	5	54.99	–
TS-DDIM (Li et al., 2024a)	5	50.69	–
LDM-ES (Ning et al., 2024)	5	49.40	1.005
DDIM-MCDO	5	<u>30.76</u>	–
DDIM-MCDO [†]	5	27.39	1.005

¹†DDIM-MCDO combined with LDM-ES using the same hyperparameters.

and direction errors in noise estimation. It is worth noting that as T' decreases, the DDIM-MCDO yields more improvement, compared to other methods. Notably, DDIM-MCDO†outperforms either component method alone, both exceeding the baseline. Additionally, at $T' = 20$, LDM-ES (Ning et al., 2024) shows sensitivity to hyperparameters (e.g., a large FID difference when λ varies from 0.999 to 1.001), but performance stabilizes when combined with DDIM-MCDO.

To maintain consistency with (Ning et al., 2024), we report results under the DDPM setting ($\eta = 1.0$) in Table 8 (see Appendix A.8). Additionally, we provide quantitative results for LSUN-Bedroom 256×256 (Yu et al., 2015) in Table 3, with corresponding qualitative results available in Appendix A.9.

For implementing DDIM-MCDO, variance statistics from 64 samples are collected during 500 steps of DDIM sampling. In all experiments on CelebA-HQ and LSUN-Bedroom, we set the threshold timestep t_{thre} to 0. Given that DDIM-MCDO consistently improves performance across all experiments, we consider it a hyperparameter tuning-free.

Table 3: Results on LSUN-Bedroom 256×256 , with LDM-4 and DDIM sampler, $\eta = 0.0$, $t_{thre} = 0$.

Method	T'	FID↓
DDIM	20	4.26
DDIM-MCDO	20	4.20
DDIM	10	9.33
DDIM-MCDO	10	6.25
DDIM	6	25.94
DDIM-MCDO	6	12.34
DDIM	5	44.91
DDIM-MCDO	5	19.31
DDIM	4	78.70
DDIM-MCDO	4	33.79

5.4 RESULTS ON CLASS-CONDITIONAL GENERATION

We conduct class-conditional image generation on ImageNet-256 (Russakovsky et al., 2015). To show the hyperparameter tuning-free feature of DDIM-MCDO, we employ the same hyperparameters in Section 5.3: $N = 64$, $t_{thre} = 0$. The pixel variance of 64 samples from different classes (1 sample per class) is collected during a 500-step DDIM sampling. Quantitative and qualitative results are shown in Table 4 and Appendix A.10, respectively.

Table 4: Results on ImageNet-256, with LDM-4 and DDIM sampler. $\eta = 0.0$, guidance scale $s = 3.0$, $t_{thre} = 0$.

Method	T'	FID↓	sFID↓	IS↑
DDIM	20	11.01	7.69	69.44
DDIM-MCDO	20	11.29	6.03	69.47
DDIM	10	9.78	12.35	66.30
DDIM-MCDO	10	10.22	7.13	67.95
DDIM	5	16.81	36.07	49.92
DDIM-MCDO	5	8.30	10.89	62.15

5.5 ABLATION STUDY

In this section, we investigate how the number of constraints affects image quality by varying the hyperparameter t_{thre} from 1000 to 0. We evaluate the FID, sFID, and IS scores of 50,000 images generated with each t_{thre} on ImageNet-256 (Russakovsky et al., 2015) using LDM-4 and 5-step DDIM samplers (Song et al., 2021).

Results in Table 5 show consistent improvements as t_{thre} decreases. While increasing discrepancy in $\mathbf{x}_0^{(t)}|_y$ at lower t may introduce errors in pixel variance correction, the benefits of DDIM-MCDO in later stages outweigh this issue, leading to better results than its early-stopping version (e.g., $t_{thre} = 200$).

Table 5: Result on ImageNet-256 with LDM-4. $T = 5$, $\eta = 0.0$, $s = 3.0$

Method	FID↓	sFID↓	IS↑	t_{thre}
DDIM-MCDO	16.81	36.07	49.92	1000
	13.45	26.72	54.64	800
	10.67	19.14	58.29	600
	9.38	15.32	60.29	400
	8.48	11.76	61.82	200
	8.30	10.89	62.15	0

6 CONCLUSION

In this paper, we address exposure bias in accelerated sampling of diffusion models. By analyzing this bias through the manifold hypothesis, we propose a manifold constraint that aligns deviated denoising observations with the manifold, enhancing both magnitude and direction accuracy in noise estimation. Our experiments on several large-resolution datasets demonstrate significant effectiveness. Furthermore, the method is training-free and requires minimal hyperparameter tuning, making it a simple plug-and-play module for other sampling methods.

ACKNOWLEDGMENTS

This work is funded by the National Natural Science Foundation of China (62403429), Zhejiang Provincial Natural Science Foundation of China under Grant (LQN25F030018, LQN25F030008).

REFERENCES

- Fan Bao, Chongxuan Li, Jun Zhu, and Bo Zhang. Analytic-dpm: an analytic estimate of the optimal reverse variance in diffusion probabilistic models. In *International Conference on Learning Representations*, 2022.
- Defang Chen, Zhenyu Zhou, Jian-Ping Mei, Chunhua Shen, Chun Chen, and Can Wang. A geometric perspective on diffusion models, 2023.
- Defang Chen, Zhenyu Zhou, Can Wang, Chunhua Shen, and Siwei Lyu. On the trajectory regularity of ODE-based diffusion sampling. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=H86WzfH5N1>.
- Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=nJJjv0JDJju>.
- Hyungjin Chung, Jeongsol Kim, Geon Yeong Park, Hyelin Nam, and Jong Chul Ye. CFG++: Manifold-constrained classifier free guidance for diffusion models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=E77uvbOTtp>.
- Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=FPnUhsQJ5B>.
- Martin Nicolas Everaert, Athanasios Fitsios, Marco Bocchio, Sami Arpa, Sabine Süsstrunk, and Radhakrishna Achanta. Exploiting the signal-leak bias in diffusion models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 4025–4034, 2024.
- Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014.
- Yutong He, Naoki Murata, Chieh-Hsin Lai, Yuhta Takida, Toshimitsu Uesaka, Dongjun Kim, Wei-Hsiang Liao, Yuki Mitsufuji, J Zico Kolter, Ruslan Salakhutdinov, and Stefano Ermon. Manifold preserving guided diffusion. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=o3BxOLoxml>.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Advances in Neural Information Processing Systems*, pp. 6626–6637, 2017.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance, 2022. URL <https://arxiv.org/abs/2207.12598>.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Jonathan Ho, Tim Salimans, Alexey A. Gritsenko, William Chan, Mohammad Norouzi, and David J. Fleet. Video diffusion models. In *Advances in Neural Information Processing Systems*, pp. 8633–8646, 2022.

- Ahmed Imtiaz Humayun, Ibtihel Amara, Candice Schumann, Golnoosh Farnadi, Negar Rostamzadeh, and Mohammad Havaei. Understanding the local geometry of generative model manifolds, 2024. URL <https://arxiv.org/abs/2408.08307>.
- Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=Hk99zCeAb>.
- Dongjun Kim, Chieh-Hsin Lai, Wei-Hsiang Liao, Naoki Murata, Yuhta Takida, Toshimitsu Uesaka, Yutong He, Yuki Mitsufuji, and Stefano Ermon. Consistency trajectory models: Learning probability flow ODE trajectory of diffusion. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=ymjI8feDTD>.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2022. URL <https://arxiv.org/abs/1312.6114>.
- Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. *Advances in neural information processing systems*, 29, 2016.
- Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis, 2021. URL <https://arxiv.org/abs/2009.09761>.
- Mingxiao Li, Tingyu Qu, Ruicong Yao, Wei Sun, and Marie-Francine Moens. Alleviating exposure bias in diffusion models through sampling with shifted time steps. In *The Twelfth International Conference on Learning Representations*, 2024a. URL <https://openreview.net/forum?id=ZSD3Ml0Ke6>.
- Yangming Li and Mihaela van der Schaar. On error propagation of diffusion models. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=RtAct1E2zS>.
- Zhimin Li, Jianwei Zhang, Qin Lin, Jiangfeng Xiong, Yanxin Long, Xincheng Deng, Yingfang Zhang, Xingchao Liu, Minbin Huang, Zedong Xiao, et al. Hunyuan-dit: A powerful multi-resolution diffusion transformer with fine-grained chinese understanding. *arXiv preprint arXiv:2405.08748*, 2024b.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Proceedings of the European Conference on Computer Vision*, pp. 740–755, 2014.
- Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.
- Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on manifolds. In *International Conference on Learning Representations*, 2022.
- 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*, 2022a.
- 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 preprint arXiv:2211.01095*, 2022b.
- Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pp. 8162–8171. PMLR, 2021.
- Mang Ning, Enver Sangineto, Angelo Porrello, Simone Calderara, and Rita Cucchiara. Input perturbation reduces exposure bias in diffusion models. *arXiv preprint arXiv:2301.11706*, 2023.

- Mang Ning, Mingxiao Li, Jianlin Su, Albert Ali Salah, and Itir Onal Ertugrul. Elucidating the exposure bias in diffusion models. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=xEJMojl1SpX>.
- George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. *Advances in neural information processing systems*, 30, 2017.
- Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. In *International Conference on Learning Representations*, 2024.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Zhiyao Ren, Yibing Zhan, Liang Ding, Gaoang Wang, Chaoyue Wang, Zhongyi Fan, and Dacheng Tao. Multi-step denoising scheduled sampling: Towards alleviating exposure bias for diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 4667–4675, 2024.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Fei-Fei Li. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- Florian Schmidt. Generalization in generation: A closer look at exposure bias. In Alexandra Birch, Andrew Finch, Hiroaki Hayashi, Ioannis Konstas, Thang Luong, Graham Neubig, Yusuke Oda, and Katsuhito Sudoh (eds.), *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pp. 157–167, Hong Kong, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-5616. URL <https://aclanthology.org/D19-5616>.
- Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics, 2015.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *International Conference on Learning Representations*, 2020. URL [dblp.org/rec/journals/corr/abs-2010-02502](https://arxiv.org/abs/2010.02502).
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems*, pp. 11895–11907, 2019.
- Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *International Conference on Machine Learning*, pp. 32211–32252, 2023.
- Xingzhe Su, Daixi Jia, Fengge Wu, Junsuo Zhao, Changwen Zheng, and Wenwen Qiang. Unbiased image synthesis via manifold guidance in diffusion models, 2024. URL <https://arxiv.org/abs/2307.08199>.
- Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical variational autoencoder. *Advances in neural information processing systems*, 33:19667–19679, 2020.
- Chaojun Wang and Rico Sennrich. On exposure bias, hallucination and domain shift in neural machine translation. *arXiv preprint arXiv:2005.03642*, 2020.
- Sven-Ake Wegner. Lecture notes on high-dimensional data, 2024. URL <https://arxiv.org/abs/2101.05841>.

- Yuzhe Yao, Feng Tian, Jun Chen, Haonan Lin, Guang Dai, Yong Liu, and Jingdong Wang. Timestep-aware correction for quantized diffusion models. *arXiv preprint arXiv:2407.03917*, 2024.
- 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.
- Yaodong Yu, Sam Buchanan, Druv Pai, Tianzhe Chu, Ziyang Wu, Shengbang Tong, Benjamin David Haeffele, and Yi Ma. White-box transformers via sparse rate reduction. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=THf18hdVxH>.
- Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator. In *International Conference on Learning Representations*, 2023.
- Yizhe Zhang, Jiatao Gu, Zhuofeng Wu, Shuangfei Zhai, Joshua Susskind, and Navdeep Jaitly. Planner: generating diversified paragraph via latent language diffusion model. *Advances in Neural Information Processing Systems*, 36, 2024.
- Zhenyu Zhou, Defang Chen, Can Wang, and Chun Chen. Fast ode-based sampling for diffusion models in around 5 steps. *arXiv preprint arXiv:2312.00094*, 2023.

A APPENDIX

A.1 ZERO MEAN ASSUMPTION AND ERROR ANALYSIS

The noise added to clean data during the training phase follows normal distribution: $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, which implies $\mathbb{E}[\epsilon_t] = 0$. According to Equation 5, assuming $\mathbb{E}[\epsilon_\theta^t] = 0$ is reasonable for well-trained diffusion models. According to Equation 7, the predicted noisy data at timestep $t - 1$ can

be expressed as $\hat{\mathbf{x}}_{t-1} = \sqrt{\bar{\alpha}_t} \frac{\hat{\mathbf{x}}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta^{(t)}(\mathbf{x}_t)}{\sqrt{\bar{\alpha}_t}} + \sqrt{1 - \bar{\alpha}_{t-1} - \sigma_t^2} \epsilon_\theta^{(t)}(\mathbf{x}_t) + \sigma_t \epsilon'_t$, $\epsilon'_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

Then as long as $\mathbb{E}[\frac{1}{n} \sum_{i=0}^n \mathbf{x}_{t_i}] = 0$, one has $\mathbb{E}[\frac{1}{n} \sum_{i=0}^n \hat{\mathbf{x}}_{t-1_i}] = 0, \forall t \in [1, T]$.

We illustrate the evolution of the average absolute pixel mean ($\mathbb{E}[\frac{1}{n} \sum_{i=0}^n \cdot]$) for the noisy data prediction, denoising observation, and noise estimation in Figure 9. The statistics were collected from 6,400 samples generated using a 100-step LDM on CelebA-HQ. It can be observed that the average absolute pixel mean values for $\hat{\mathbf{x}}_t$, $\hat{\mathbf{x}}_0^{(t)}$ and $\epsilon_\theta^{(t)}$ all remain close to zero throughout the generation process. This suggests that the pixel mean for all variables exhibits only minor deviations from 0 across timesteps.

In Equations 12 and 19, the square root of the pixel variance for $\hat{\mathbf{x}}_t$ and $\hat{\mathbf{x}}_0^{(t)}$ is used to approximate their L^2 norm: $\sqrt{n \text{Var}(\hat{\mathbf{x}}_t)} \rightarrow \|\hat{\mathbf{x}}_t\|$, $\sqrt{n \text{Var}(\hat{\mathbf{x}}_0^{(t)})} \rightarrow \|\hat{\mathbf{x}}_0^{(t)}\|$. We present the evolution of the L^2 norm alongside the pixel variance for the denoising observation in Figure 7. The relative error of this approximation is presented in Figure 8. The results demonstrate that throughout the entire generation process, the approximation error remains very small (less than 0.045). This indicates that using the pixel variance of $\hat{\mathbf{x}}_t$ or $\hat{\mathbf{x}}_0^{(t)}$ as an approximation for their L^2 norm introduces only minimal error, which is acceptable for practical applications.

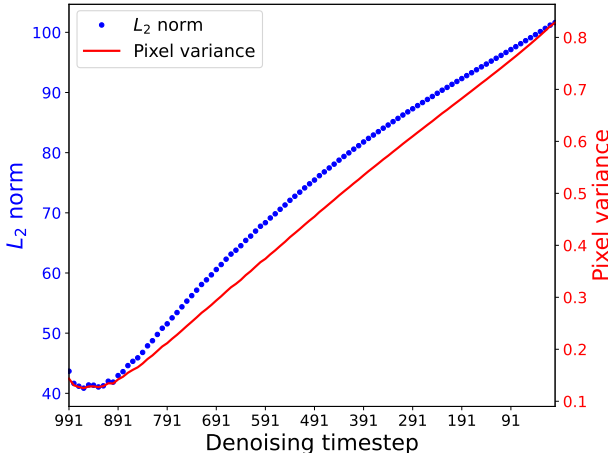


Figure 7: Evolution of the L^2 norm and pixel variance for the denoising observation, based on average statistics from 6,400 samples generated using 100-step LDM on CelebA-HQ.

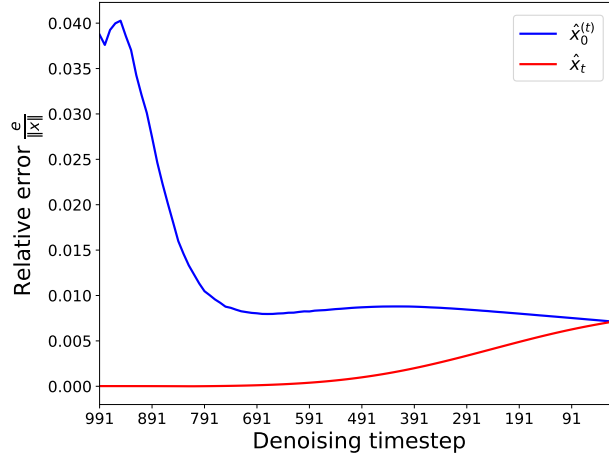


Figure 8: Relative error of using square root of pixel variance to approximate the L^2 norm of \hat{x}_t and $\hat{x}_0^{(t)}$. Statistics are collected from 6,400 samples generated using 100-step LDM on CelebA-HQ.

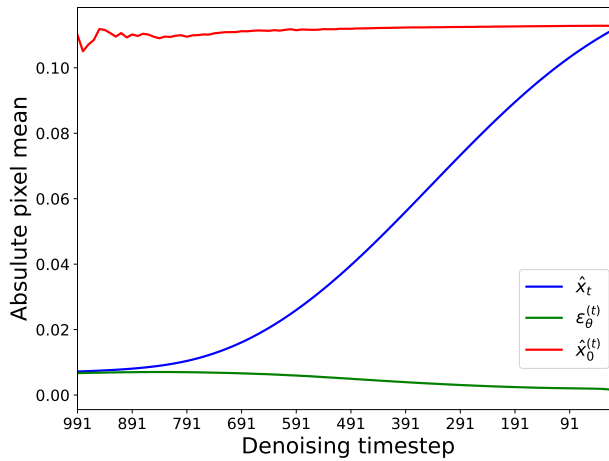


Figure 9: Evolution of the average absolute pixel mean for the noisy data prediction, denoising observation, and noise estimation. Statistics are collected from 6,400 samples generated using 100-step LDM on CelebA-HQ.

A.2 DERIVATION OF NOISY RATIO IN A SINGLE STEP FOR DDIM

In the forward process of DDIM, we have:

$$\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon} = \mathbf{x}_t, \boldsymbol{\epsilon} \in \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (23)$$

In the reverse process, when given the ground truth input \mathbf{x}_{t+1} , the model's prediction of \mathbf{x}_t is:

$$\hat{\mathbf{x}}_t = \sqrt{\bar{\alpha}_t} \hat{\mathbf{x}}_0^{(t+1)} + \sqrt{1 - \bar{\alpha}_t - \sigma_{t+1}^2} \boldsymbol{\epsilon}_\theta^{(t+1)}(\mathbf{x}_{t+1}) + \sigma_{t+1} \boldsymbol{\epsilon}'_{t+1}, \boldsymbol{\epsilon}'_{t+1} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (24)$$

When $\sigma_t = 0$ for all t , the sampling process become deterministic, and Equation 24 is simplified as:

$$\hat{\mathbf{x}}_t = \sqrt{\bar{\alpha}_t} \hat{\mathbf{x}}_0^{(t+1)} + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}_\theta^{(t+1)}(\mathbf{x}_{t+1}). \quad (25)$$

For timestep t , by denoting the prediction error on the noise as $\mathbf{e}_\theta^{(t)}$, we have:

$$\boldsymbol{\epsilon}_\theta^{(t)}(\mathbf{x}_t) = \boldsymbol{\epsilon}_t + \mathbf{e}_\theta^{(t)}. \quad (26)$$

We now derive the error on denoising observation with respect to the error on the noise estimation. By plugging Eq 26 into Equation 6, we have:

$$\begin{aligned} \hat{\mathbf{x}}_0^{(t+1)} &= \frac{\mathbf{x}_{t+1} - \sqrt{1 - \bar{\alpha}_{t+1}} \boldsymbol{\epsilon}_\theta^{(t+1)}(\mathbf{x}_{t+1})}{\sqrt{\bar{\alpha}_{t+1}}} \\ &= \frac{\mathbf{x}_{t+1} - \sqrt{1 - \bar{\alpha}_{t+1}} \boldsymbol{\epsilon}_{t+1}}{\sqrt{\bar{\alpha}_{t+1}}} - \frac{\sqrt{1 - \bar{\alpha}_{t+1}}}{\sqrt{\bar{\alpha}_{t+1}}} \mathbf{e}_\theta^{(t+1)} \\ &= \mathbf{x}_0 - \frac{\sqrt{1 - \bar{\alpha}_{t+1}}}{\sqrt{\bar{\alpha}_{t+1}}} \mathbf{e}_\theta^{(t+1)}. \end{aligned} \quad (27)$$

It is worth noting that in the early stage of generation, e.g., when $t = T$, $\bar{\alpha}_t$ approaches 0. Therefore, according to Equation 27, the error coefficient on the denoising observation: $d_t = \frac{\sqrt{1 - \bar{\alpha}_{t+1}}}{\sqrt{\bar{\alpha}_t}}$ is non-negligible.

As $\hat{\mathbf{x}}_t^{(0)}$ is notably influenced by the amplified noise prediction error term $d_t \mathbf{e}_\theta^{(t)}$ at most timesteps, it is possible to correct the noise estimation error by correcting the denoising observation.

With the same manner, we derive the error on the noisy data prediction with respect to the noise estimation error. By plugging X Equations 26 and 27 into Equation 25, we have:

$$\begin{aligned} \hat{\mathbf{x}}_t &= \sqrt{\bar{\alpha}_t} \left(\mathbf{x}_0 - \frac{\sqrt{1 - \bar{\alpha}_{t+1}}}{\sqrt{\bar{\alpha}_{t+1}}} \mathbf{e}_\theta^{(t+1)} \right) + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}_\theta^{(t+1)}(\mathbf{x}_{t+1}) \\ &= \sqrt{\bar{\alpha}_t} \mathbf{x}_0 - \sqrt{\bar{\alpha}_t} \frac{\sqrt{1 - \bar{\alpha}_{t+1}}}{\sqrt{\bar{\alpha}_{t+1}}} \mathbf{e}_\theta^{(t+1)} + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}_{t+1} + \sqrt{1 - \bar{\alpha}_t} \mathbf{e}_\theta^{(t+1)} \\ &= \left(\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}_{t+1} \right) + \left(\sqrt{1 - \bar{\alpha}_t} - \sqrt{\bar{\alpha}_t} \frac{\sqrt{1 - \bar{\alpha}_{t+1}}}{\sqrt{\bar{\alpha}_{t+1}}} \right) \mathbf{e}_\theta^{(t+1)} \\ &= \mathbf{x}_t + \left(\sqrt{1 - \bar{\alpha}_t} - \sqrt{\bar{\alpha}_t} \frac{\sqrt{1 - \bar{\alpha}_{t+1}}}{\sqrt{\bar{\alpha}_{t+1}}} \right) \mathbf{e}_\theta^{(t+1)}. \end{aligned} \quad (28)$$

We denote the scaling factor for \mathbf{x}_t as n_t , $n_t = \sqrt{1 - \bar{\alpha}_t} - \sqrt{\bar{\alpha}_t} \frac{\sqrt{1 - \bar{\alpha}_{t+1}}}{\sqrt{\bar{\alpha}_{t+1}}}$. The value of both d_t and n_t during denoising process are shown in Figures 10 and 11. It is illustrated that while $\mathbf{e}_\theta^{(t)}$ is scaled down for \mathbf{x}_t , the same error is scaled up for $\mathbf{x}_0^{(t)}$.

Correcting the pixel variance of $\hat{\mathbf{x}}_t$ using a scale factor (directly correcting the pixel variance of the noisy data prediction) will eventually scale the entire term $\mathbf{x}_t + n_t \mathbf{e}_\theta^{(t)}$, which leads to distortion of the \mathbf{x}_t component in $\hat{\mathbf{x}}_t$. When correcting the pixel variance of $\hat{\mathbf{x}}_0^{(t)}$ using the proposed method, one corrects $\mathbf{x}_0 + d_t \mathbf{e}_\theta^{(t)}$. While the $\mathbf{x}_0^{(t)}$ term is also influenced, it is important to note that by recalculating the noise estimation with the corrected denoising observation in Equation 6, the prediction error is reduced while the original \mathbf{x}_t is preserved.




Figure 10: Error scale for denoising observation




Figure 11: Error scale for noisy data

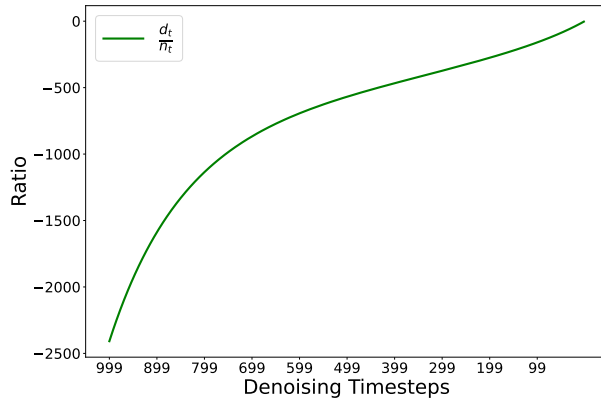


Figure 12: Ratio of error scale

A.3 COMPARISON BETWEEN NOISY DATA AND DENOISING OBSERVATION CORRECTION

According to Equation. 16, adjusting the pixel variance of noisy data \hat{x}_t during the accelerated sampling process to variance collected from sampling process with large enough T , could potentially reduce the exposure bias and improve the generation results. For clarity, we refer to this method as Noisy Data Variance Adjustment (NDVA). We provide the results of NDVA in Tables 6 and 7.

In Table 6, it can be observed that NDVA outperforms the baseline method, and exhibits performance comparable to TS-DDIM when $T = 20$. However, it is notably worse than MCDO as the number of timesteps further decreases.

Table 6: Results on CelebA-HQ dataset.

Method	Steps	FID↓
DDIM	20	10.59
TS-DDIM	20	8.29
DDIM-NVDA	20	9.21
DDIM-MCDO	20	9.49
DDIM	10	21.08
TS-DDIM	10	15.71
DDIM-NVDA	10	16.45
DDIM-MCDO	10	15.63
DDIM	5	54.99
TS-DDIM	5	50.69
DDIM-NVDA	5	47.52
DDIM-MCDO	5	30.76

Table 7: Results on MS-COCO.

Method	Steps	FID↓
DPM-Solver++(2M)	10	15.10
DPM-Solver++(2M) + NVDA	10	15.76
DPM-Solver++(2M) + MCDO	10	14.88
DPM-Solver++(2M)	8	15.57
DPM-Solver++(2M) + NVDA	8	15.76
DPM-Solver++(2M) + MCDO	8	15.32
DPM-Solver++(2M)	6	17.05
DPM-Solver++(2M) + NVDA	6	17.13
DPM-Solver++(2M) + MCDO	6	16.39
DPM-Solver++(2M)	5	19.16
DPM-Solver++(2M) + NVDA	5	19.02
DPM-Solver++(2M) + MCDO	5	18.51

A.4 EXPOSURE BIAS MEASUREMENT

In this section, we employ the exposure bias metric proposed by Ning et al. (2024) to evaluate the single step variance discrepancy reduction when using our method. The statistics are collected during the accelerated diffusion sampling (5, 10, 20 steps) on Stable Diffusion XL. The variance deviation with respect to the fixed schedule $1 - \bar{\alpha}_t$, as well as the standard deviation are illustrated in Figure 13. Although our approach is based on the manifold hypothesis and the concept of pixel variance, the results indicate that the proposed manifold constraint still reduces sample variance error.

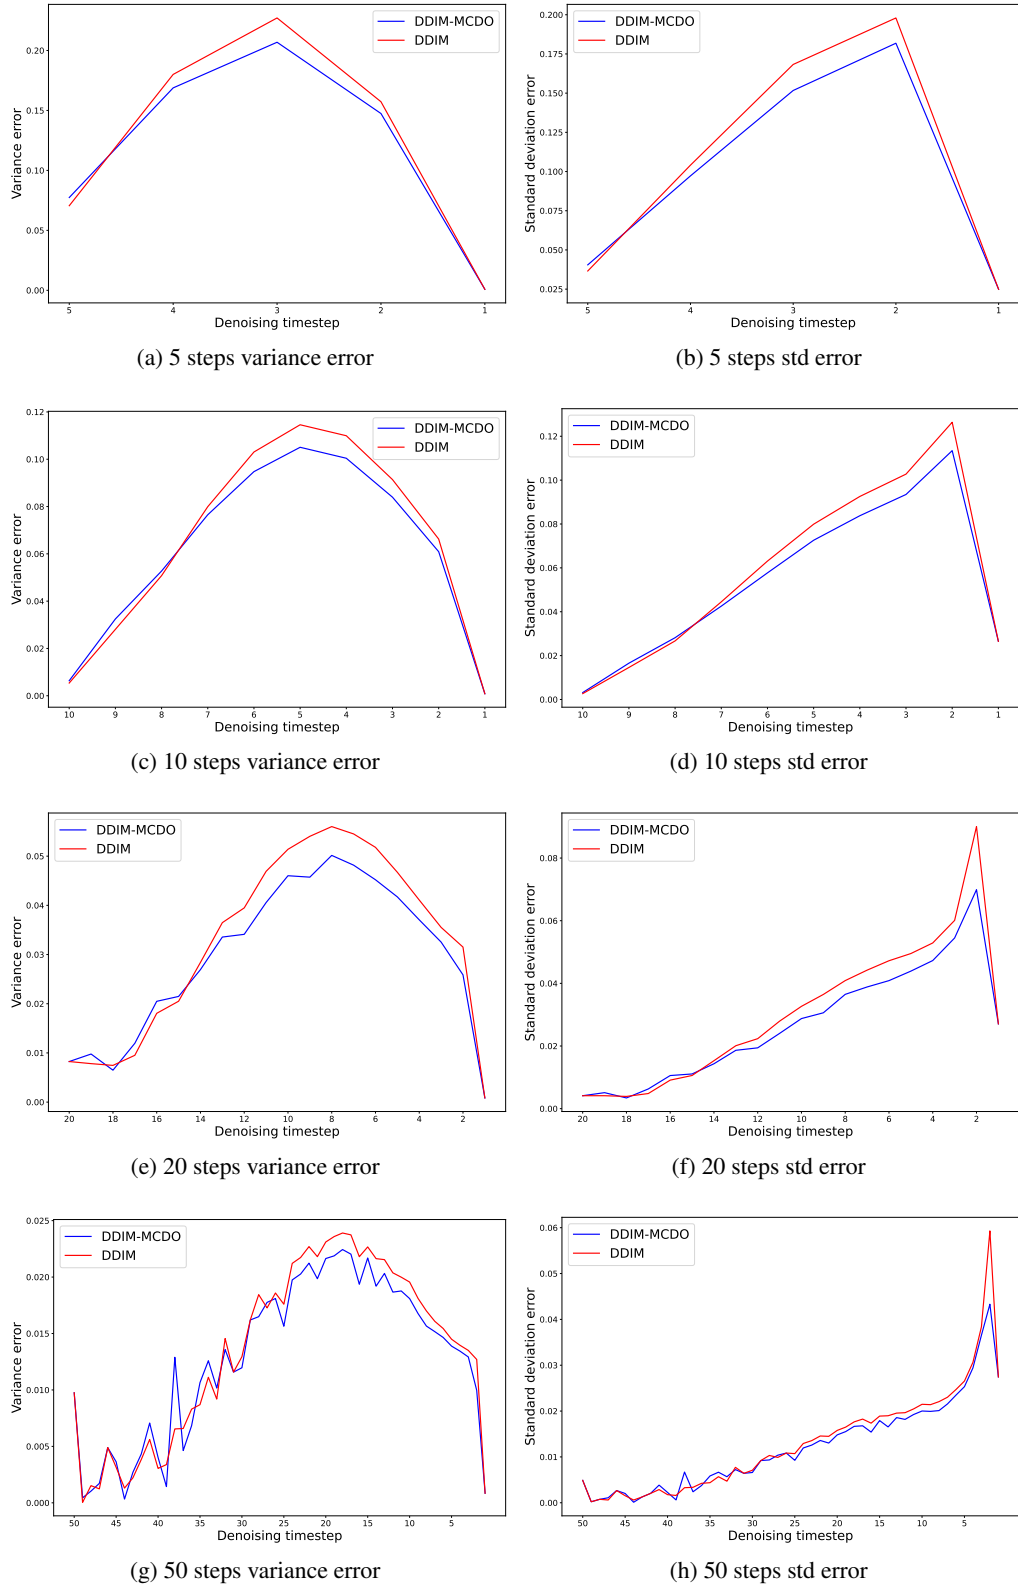


Figure 13: Exposure bias measurement. We plot the variance error and standard deviation error in single-step samplings.

A.5 MORE QUALITATIVE RESULTS ON MS-COCO USING SDXL

We provide more qualitative results on text-to-image generation in Figure 14.



(a) "A bunch of sheep are standing in a snowy field."



(b) "a woman walking along a sidewalk while holding a suitcase"



(c) "A girl holding a foot long hot dog in a restaurant"



(d) "A young man in a hat holding a cup of coffee and a pastry."

Figure 14: Images generated using 10 steps SDXL with DDIM sampler. From left to right: SDXL, TS-DDIM (Li et al., 2024a), LDM-ES (Ning et al., 2024), DDIM-MCDO (ours).

A.6 QUALITATIVE RESULTS ON CELEBA-HQ

We present qualitative results for unconditional generation on CelebA-HQ using LDM-4 in Figures 15, 16, and 17, with sampling steps of 5, 10, and 20, respectively. For the sampling settings, we use the recommended value of $\eta = 0.0$ when the number of steps is low (Rombach et al., 2022).




Figure 15: Samples on CelebA-HQ 256×256 using 5 steps LDM-4. From left to right: DDIM, TS-DDIM (Li et al., 2024a), LDM-ES (Ning et al., 2024) with $\lambda = 1.005$, DDIM-MCDO, DDIM-MCDO† with $\lambda = 1.005$.




Figure 16: Samples on CelebA-HQ 256×256 using 10 steps LDM-4. From left to right: DDIM, TS-DDIM (Li et al., 2024a), LDM-ES (Ning et al., 2024) with $\lambda = 1.002$, DDIM-MCDO, DDIM-MCDO† with $\lambda = 1.002$.



Figure 17: Samples on CelebA-HQ 256×256 using 20 steps LDM-4. From left to right: DDIM, TS-DDIM (Li et al., 2024a), LDM-ES (Ning et al., 2024) with $\lambda = 0.999$, LDM-ES (Ning et al., 2024) with $\lambda = 1.001$, DDIM-MCDO, DDIM-MCDO \dagger with $\lambda = 0.999$, DDIM-MCDO with $\lambda = 1.001$.

A.7 PIXEL-LEVEL CONSTRAINT FOR UNCONDITIONAL GENERATION

Considering that in unconditional generation, the similarity between samples within an inference batch is higher than that in Text-to-Image generation, we propose a more fine-grained correction approach: to estimate the pixel-level variance of the denoising observation: $v_{t,c,h,w} = \text{Var}(\hat{x}_{0_{i,c,h,w}}^{(t)})$, $v_t \in \mathbb{R}^{C,H,W}$, $\hat{x}_0^{(t)} \in \mathbb{R}^{B,C,H,W}$. Where B is the batch size, C is the channel number, H and W are the latent’s height and weight, respectively. We denote this variant as DDIM-MCDO*. When DDIM-MCDO* is coupled with LDM-ES (Ning et al., 2024), it is denoted as DDIM-MCDO^{‡2}. The batch size N used for implementing DDIM-MCDO* is 64 on CelebA-HQ.

A.8 QUANTITATIVE RESULTS ON CELEBA-HQ WITH DDPM SAMPLER

We provide results on CelebA-HQ using LDM-4 and DDPM sampler, where $\eta = 1.0$, in Table 8. We notice that contrast to LDM-ES (Ning et al., 2024), our method yields more performance improvement for DDIM than DDPM. Considering the DDIM setting is recommended under accelerated generation (Rombach et al., 2022), results in Table 8 suggests that the proposed methods might be easily combined with other fast sampling method in scenarios where T' is extremely low.

Table 8: Results on CelebA-HQ 256×256 with LDM-4 and DDPM sampler, $\eta = 1.0$.

Method	T'	FID↓	λ
DDPM	20	29.61	–
LDM-ES (Ning et al., 2024)	20	15.68	1.01
DDPM-MCDO	20	26.83	–
DDPM-MCDO*	20	23.84	–
DDPM-MCDO [†]	20	<u>16.47</u>	1.01
DDPM-MCDO [‡]	20	19.21	1.01
DDPM	10	56.22	–
LDM-ES (Ning et al., 2024)	10	33.36	1.03
DDPM-MCDO	10	47.29	–
DDPM-MCDO*	10	29.76	–
DDPM-MCDO [†]	10	<u>19.87</u>	1.03
DDPM-MCDO [‡]	10	16.92	1.03
DDPM	5	102.1	–
LDM-ES (Ning et al., 2024)	5	73.49	1.04
DDPM-MCDO	5	78.12	–
DDPM-MCDO*	5	<u>54.56</u>	–
DDPM-MCDO [†]	5	56.40	1.04
DDPM-MCDO [‡]	5	51.76	1.04

^{‡2}DDIM-MCDO* combined with LDM-ES using the same hyperparameter.

A.9 QUALITATIVE RESULTS ON LSUN-BEDROOM

We present qualitative results for unconditional generation on LSUN-Bedroom using LDM-4 in Figures 18, 19, and 20, with sampling steps of 4, 5, and 10, respectively. For the sampling settings, we use the recommended value of $\eta = 0.0$ when the number of steps is low (Rombach et al., 2022).



Figure 18: Samples on LSUN-bedroom 256×256 using 4 steps LDM-4.



Figure 19: Samples on LSUN-bedroom 256×256 using 5 steps LDM-4.

A.10 QUALITATIVE RESULTS ON IMAGENET-256

We present qualitative results of class-conditional generation on ImageNet-256 using LDM-4 in Figures 21, 22, and 23, with sampling steps of 5, 10, and 20, respectively. For the sampling settings, we use the recommended value of $\eta = 0.0$ when the number of steps is low (Rombach et al., 2022).




Figure 20: Samples on LSUN-bedroom 256×256 using 10 steps LDM-4.

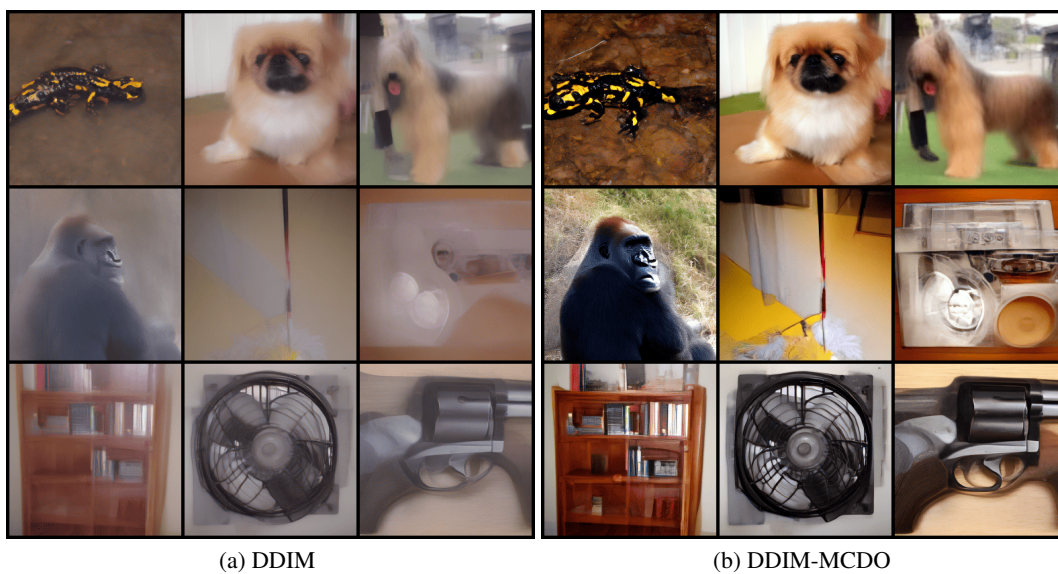


Figure 21: Samples on ImageNet 256×256 using 5 steps LDM-4.

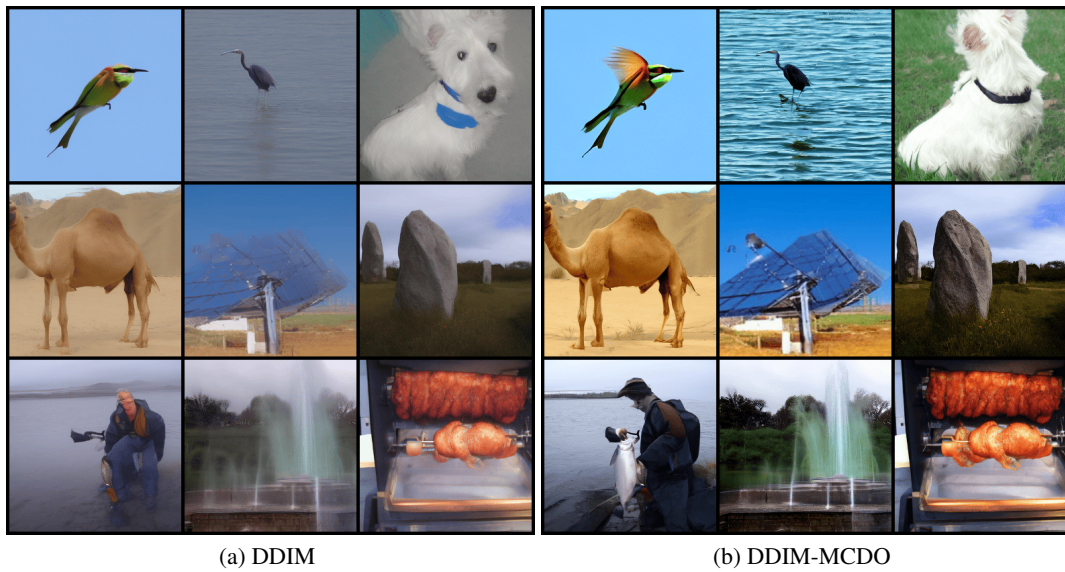


Figure 22: Samples on ImageNet 256×256 using 10 steps LDM-4.




Figure 23: Samples on ImageNet 256×256 using 20 steps LDM-4.

A.11 A MORE FINE-GRAINED MANIFOLD CONSTRAINT

In this section, based on the manifold analysis in Section 4.1, we provide a more fine-grained version of MCOD. Which is denoted as DDIM-MCDO+. The statistics collection process of DDIM-MCDO+ and DDIM-MCDO are shared, so DDIM-MCDO+ can be implemented by adding one more line of code.

During inference, we first correct the pixel variance of the denoising observation $\hat{\mathbf{x}}_0^{(t)}$ like in DDIM-MCDO:

$$\tilde{\mathbf{x}}_0^{(t)}|_{y_2, x_{T_2}} = \sqrt{c_t} \left[\hat{\mathbf{x}}_0^{(t)}|_{y_2, x_{T_2}} - \text{Mean} \left(\hat{\mathbf{x}}_0^{(t)}|_{y_2, x_{T_2}} \right) \right] + \text{Mean} \left(\hat{\mathbf{x}}_0^{(t)}|_{y_2, x_{T_2}} \right). \quad (29)$$

Then, we calculate the transition coefficient d_{t-k} between $\mathcal{M}_0^{(t)}$ and $\mathcal{M}_0^{(t-k)}$:

$$d_{t-k} = \frac{\text{Var}(\mathbf{x}_0^{*(t-k)}|_{y_2, x_{T_2}})}{\text{Var}(\mathbf{x}_0^{*(t)}|_{y_2, x_{T_2}})}. \quad (30)$$

With d_{t-k} , we can further correct the denoising observation, which will be used in predicting \mathbf{x}_{t-k} .

$$\tilde{\mathbf{x}}_0^{(t-k)}|_{y_2, x_{T_2}} = \sqrt{d_{t-k}} \tilde{\mathbf{x}}_0^{(t)}|_{y_2, x_{T_2}}. \quad (31)$$

The final expression of $\tilde{\mathbf{x}}_{t-k}$:

$$\tilde{\mathbf{x}}_{t-k} = \sqrt{\bar{\alpha}_{t-k}} \tilde{\mathbf{x}}_0^{(t-k)} + \sqrt{1 - \bar{\alpha}_{t-k} - \sigma_{t-k}^2} \tilde{\boldsymbol{\epsilon}}_{\theta}(\mathbf{x}_t) + \sigma_{t-k} \boldsymbol{\epsilon}'_{t-k}, \boldsymbol{\epsilon}'_{t-k} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (32)$$

The result on MS-COCO (Lin et al., 2014) dataset with Stable Diffusion XL (Podell et al., 2024) are provided in Figure 24 and Table 9. We find that although DDIM-MCDO+ does not outperform DDIM-MCDO on metrics like FID and CLIP score, the quantitative results of DDIM-MCDO+ might be better.

Table 9: Results on Text-to-Image generation on MS-COCO val2014 with SDXL and DDIM sampler. Image Resolution is 1024×1024 .

Model	Method	Steps	FID↓	CLIP Score↑	t_{thre}
SDXL	DDIM	10	18.17	31.58	–
	DDIM-MCDO	10	15.60 (-2.75)	31.75 (+0.18)	0
	DDIM-MCDO+	10	15.79 (-2.38)	31.72 (+0.14)	100
	DDIM-MCDO+	10	15.98 (-2.19)	31.71 (+0.13)	0



(a) "A beach area with several chairs and umbrellas."



(b) "A dog laying on the beach in the sand."

Figure 24: Comparison between DDIM-MCDO+ (left) DDIM-MCDO (right) using 10 steps SDXL.

A.12 RESULTS WITH DPM-SOLVER++

To further investigate the compatibility of MCOD with high-order solvers, we evaluate the proposed method with DPM-Solver++ (Lu et al., 2022b). The pixel variance of 60 samples, collected from the 1000 steps SDE-DPM-Solver++(2M) sampling process, were used for v_t estimation. We generated 10k images for each evaluation in Table 10. The corresponding qualitative results are provided in Figures 25, 26, 27 and 28.

Table 10: Results on Text-to-Image generation on MS-COCO val2014 with SDXL and DPM-Solver++. Image Resolution is 1024×1024

Method	Steps	FID↓	CLIP Score↑	t_{thre}
DPM-Solver++(2M)	10	15.10	31.74	–
DPM-Solver++(2M)-MCDO	10	14.88	31.85	100
DPM-Solver++(2M)	8	15.57	31.74	–
DPM-Solver++(2M)-MCDO	8	15.32	31.79	100
DPM-Solver++(2M)	6	17.05	31.65	–
DPM-Solver++(2M)-MCDO	6	16.39	31.72	100
DPM-Solver++(2M)	5	19.16	31.52	–
DPM-Solver++(2M)-MCDO	5	18.51	31.62	100



(a) "A busy city street with many different vehicles."



(b) "Two men with glasses sitting at a table using laptops"

Figure 25: Comparison between DPM-Solver++(2M) (Lu et al., 2022b) (left) DPM-Solver++(2M)-MCDO (right) using 10 steps SDXL.



(a) "A public transportation bus near a curb with a bicycle rack."



(b) "Two horses plowing the land as a man directs them."

Figure 26: Comparison between DPM-Solver++(2M) (Lu et al., 2022b) (left) DPM-Solver++(2M)-MCDO (right) using 8 steps SDXL.




(a) "A person walking across the street with a traffic light above."



(b) A man in the ocean surfing holding on to a kite."

Figure 27: Comparison between DPM-Solver++(2M) (Lu et al., 2022b) (left) DPM-Solver++(2M)-MCDO (right) using 6 steps SDXL.



(a) "A bicyclist is riding past a bus that is parked."



(b) "A picture of an altar holding bananas, skeletons, candles and flowers."

Figure 28: Comparison between DPM-Solver++(2M) (Lu et al., 2022b) (left) DPM-Solver++(2M)-MCDO (right) using 5 steps SDXL.

A.13 RESULTS ON STABLE DIFFUSION 3

We further evaluate the scalability on rectified flow model: Stable Diffusion 3 (Esser et al., 2024). Quantitative results, computed using 10k images, are provided as follows.

Table 11: Results on Text-to-Image generation on MS-COCO val2014 with Stable Diffusion 3 (Esser et al., 2024). Image Resolution is 1024×1024 .

Model	Method	Steps	FID↓	CLIP Score↑	t_{thre}
Stable Diffusion 3	DDIM	6	33.31	30.56	–
	DDIM-MCDO	6	27.57	31.18	0
	DDIM	9	22.71	31.44	–
	DDIM-MCDO	9	22.26	31.70	0

A.14 RESULTS WITH DIFFERENT CLASSIFIER-FREE GUIDANCE SCALE

Considering that large classifier-free guidance (CFG) scales might cause exposure bias (Chung et al., 2025), we evaluate the proposed method, as the CFG scale varies. Quantitative results, computed using 10k images, are provided in Table 12.

Table 12: Results on Text-to-Image generation on MS-COCO val2014 with SDXL (Podell et al., 2024). Image Resolution is 1024×1024 .

Method	CFG Scale	FID↓	CLIP Score↑	t_{thre}
DDIM	8	18.51	31.81	–
DDIM-MCDO	8	17.38	32.03	100
DDIM	10	18.99	31.82	–
DDIM-MCDO	10	17.49	32.13	100
DDIM	12	20.62	31.70	–
DDIM-MCDO	12	17.80	32.18	100

The results show that incorporating the proposed method into sampling significantly enhances performance across varying CFG scales. Specifically, while higher CFG scales (e.g., 12) typically exacerbate exposure bias, the DDIM-MCDO method effectively mitigates this issue, reducing the FID score by 2.82.

A.15 QUANTITATIVE COMPARISON WITH 1,000 STEPS RESULTS

We provide quantitative comparison between images generated using 10 steps DDIM, 10 steps DDIM-MCDO, and 1,000 steps DDIM in Figure 29.



Figure 29: Comparison between 10 steps DDIM (left), DDIM-MCDO (middle) and 1000 steps DDIM (right) with SDXL. Prompts used (from top to bottom): "A street light that shows, horse crossing on it", "There are people on the snow bank and snow lift chairs are above.", "Four people gather around a table and smile for the picture.", "A small child with a kite walking on a beach."