

Accelerating Diffusion Sampling with Optimized Time Steps

Shuchen Xue¹, Zhaoqiang Liu^{2*}, Fei Chen³, Shifeng Zhang³, Tianyang Hu³, Enze Xie³, Zhenguo Li³
¹ Academy of Mathematics and Systems Science, Chinese Academy of Sciences
² University of Electronic Science and Technology of China
³ Huawei Noah's Ark Lab

Abstract

Diffusion probabilistic models (DPMs) have shown remarkable performance in high-resolution image synthesis, but their sampling efficiency is still to be desired due to the typically large number of sampling steps. Recent advancements in high-order numerical ODE solvers for DPMs have enabled the generation of high-quality images with much fewer sampling steps. While this is a significant development, most sampling methods still employ uniform time steps, which is not optimal when using a small number of steps. To address this issue, we propose a general framework for designing an optimization problem that seeks more appropriate time steps for a specific numerical ODE solver for DPMs. This optimization problem aims to minimize the distance between the ground-truth solution to the ODE and an approximate solution corresponding to the numerical solver. It can be efficiently solved using the constrained trust region method, taking less than 15 seconds. Our extensive experiments on both unconditional and conditional sampling using pixel- and latent-space DPMs demonstrate that, when combined with the state-of-the-art sampling method UniPC, our optimized time steps significantly improve image generation performance in terms of FID scores for datasets such as CIFAR-10 and ImageNet, compared to using uniform time steps.¹

1. Introduction

Diffusion Probabilistic Models (DPMs), as referenced in various studies [15, 43, 46], have garnered significant attention for their success in a wide range of generative tasks. These tasks include, but are not limited to, image synthesis [11, 16, 40], video generation [3, 17], text-to-image conversion [36, 39, 41], and speech synthesis [24, 45]. Central to the operation of DPMs is a dual-process mecha-

nism: a forward diffusion process incrementally introduces noise into the data, while a reverse diffusion process is employed to reconstruct data from this noise. Despite their superior performance in generation tasks when compared to alternative methodologies like Generative Adversarial Networks (GANs) [13] or Variational Autoencoders (VAEs) [21], DPMs require a considerable number of network function evaluations (NFEs) [15], which is a notable computational limitation for their broader application.

The scholarly pursuit to enhance the sampling efficiency of DPMs bifurcates into two distinct categories: methods that require additional training of the DPMs [33–35, 42, 47, 49, 50, 52] and those that do not [2, 18, 19, 31, 44, 53, 55]. This paper focuses on the latter training-free methods, aiming to improve the sampling efficiency of DPMs. Contemporary training-free samplers utilize efficient numerical schemes to solve the diffusion Stochastic Differential Equations (SDEs) or Ordinary Differential Equations (ODEs). While introducing stochasticity in diffusion sampling has been shown to achieve better quality and diversity [53], ODE-based sampling methods are superior when the sampling steps are fewer. For example, Song et al. [44] provide an empirically efficient solver DDIM. Zhang and Chen [54] and Lu et al. [31] point out the semi-linear structure of diffusion ODEs, and develop higher-order ODE samplers based on it. Zhao et al. [55] further improve these samplers in terms of NFEs by integrating the mechanism of the predictor-corrector method. Due to its better efficiency, sampling from diffusion ODE has been predominantly adopted. Quality samples can be generated within as few as 10 steps. However, state-of-the-art ODE-based samplers are still far from optimal. How to further improve efficiency is an active field of research.

In this paper, we focus on another important but less explored aspect of diffusion sampling, the *discretization scheme for time steps (sampling schedule)*. Our main contributions are summarized as follows:

- Motivated by classic assumptions on score approximation adopted in relevant theoretical works on DPMs, we propose a general framework for constructing an optimiza-

*Corresponding author. The work was done when the author worked at Huawei Noah's Ark Lab.

¹Code is available at <https://github.com/scxue/DM-NonUniform>

tion problem with respect to time steps, in order to minimize the distance between the ground-truth solution and an approximate solution to the corresponding ODE. The parameters for the optimization problem can be easily obtained for any explicit numerical ODE solver. In particular, we allow the numerical ODE solver to have varying orders in different steps. The optimization problem can be efficiently solved using the constrained trust region method, requiring less than 15 seconds.

- We construct the optimization problem with respect to time steps for DPM-Solver++ [32] and UniPC [55], which are recently proposed high-order numerical solvers for DPMs. Extensive experimental results show that our method consistently improves the performance for both unconditional sampling and conditional sampling in both pixel and latent space diffusion models. For example, when combined with UniPC and only using 5 neural function evaluations (NFEs), the time steps found by our optimization problem lead to an FID of 11.91 for unconditional pixel-space CIFAR-10 [23], an FID of 10.47 on conditional pixel-space ImageNet [10] 64x64, and an FID of 8.66 on conditional latent-space ImageNet 256x256.

2. Related Work

Training-based Methods Training-based methods, e.g., knowledge distillation [33–35, 42], learning-to-sample [50], integration with GANs [49, 52], and learning the consistency of diffusion ODE [47], have the potential to sampling for one or very few steps to enhance the efficiency. However, they lack a plug-and-play nature and require substantial extra training, which greatly limits their applicability across diverse tasks. We mainly focus on the training-free methods in this paper. Nevertheless, our investigation for the optimal sampling strategy can be orthogonally combined with training-based methods to further boost the performance.

Adaptive Step Size Adaptive step size during diffusion sampling has also been explored. It’s widely used in numerical solutions of ordinary differential equations the errors of the method. Jolicoeur-Martineau et al. [18] designed an adaptive step size via a lower-order and a higher-order method. Lu et al. [31] also proposes an adaptive step schedule using a similar method. Gao et al. [12] propose fast sampling through the restricting backward error schedule based on dynamically moderating the long-time backward error. However, the empirical performances of these methods in a few steps can still be improved.

Learning to Schedule Most related to our work is a line of research that also aims to find the optimal time schedule. Watson et al. [50] use dynamic programming to dis-

cover the optimal time schedule with maximum Evidence Lower Bound (ELBO). Wang et al. [48] leverage reinforcement learning method to search a sampling schedule. Liu et al. [30] design a predictor-based search algorithm to optimize both the sampling schedule and decide which model to sample from at each step given a set of pre-trained diffusion models. Xia et al. [51] propose to train a timestep aligner to align the sampling schedule. Li et al. [27] propose to utilize the evolution algorithm to search over the space of sampling schedules in terms of FID score. In comparison, our method has negligible computational cost compared to the learning-based method.

3. Preliminary

In this section, we provide introductory discussions about diffusion models and the commonly-used discretization schemes for time steps.

3.1. Diffusion Models

In the regime of the continuous SDE, DPMs [15, 22, 43, 46] construct noisy data through the following linear SDE:

$$d\mathbf{x}_t = f(t)\mathbf{x}_t dt + g(t)d\mathbf{w}_t, \quad (1)$$

where $\mathbf{w}_t \in \mathbb{R}^D$ represents the standard Wiener process, and $f(t)\mathbf{x}_t$ and $g(t)$ denote the drift and diffusion coefficients respectively. In addition, for any $t \in [0, T]$, the distribution of \mathbf{x}_t conditioned on \mathbf{x}_0 is a Gaussian distribution with mean vector $\alpha(t)\mathbf{x}_0$ and covariance matrix $\sigma^2(t)\mathbf{I}$, i.e., $\mathbf{x}_t|\mathbf{x}_0 \sim \mathcal{N}(\alpha(t)\mathbf{x}_0, \sigma^2(t)\mathbf{I})$, where positive functions $\alpha(t)$ and $\sigma(t)$ are differentiable with bounded derivatives, and are denoted as α_t and σ_t for brevity. Let q_t denote the marginal distribution of \mathbf{x}_t , and q_{0t} be the distribution of \mathbf{x}_t conditioned on \mathbf{x}_0 . The functions α_t and σ_t are chosen such that q_T closely approximate a zero-mean Gaussian distribution with covariance matrix $\tilde{\sigma}^2\mathbf{I}$ for some $\tilde{\sigma} > 0$, and the signal-to-noise-ratio (SNR) α_t^2/σ_t^2 is strictly decreasing with respect to t . Moreover, to ensure that the SDE in (1) have the same transition distribution q_{0t} for \mathbf{x}_t conditioned on \mathbf{x}_0 , $f(t)$ and $g(t)$ need to be dependent on α_t and σ_t in the following form:

$$f(t) = \frac{d \log \alpha_t}{dt}, \quad g^2(t) = \frac{d\sigma_t^2}{dt} - 2\frac{d \log \alpha_t}{dt} \sigma_t^2. \quad (2)$$

Anderson [1] establishes a pivotal theorem that the forward process (1) has an equivalent reverse-time diffusion process (from T to 0) as in the following equation, so that the generating process is equivalent to solving the diffusion SDE [15, 46]:

$$d\mathbf{x}_t = [f(t)\mathbf{x}_t - g^2(t)\nabla_{\mathbf{x}} \log q_t(\mathbf{x}_t)] dt + g(t)d\bar{\mathbf{w}}_t, \quad (3)$$

where $\bar{\mathbf{w}}_t$ represents the Wiener process in reverse time, and $\nabla_{\mathbf{x}} \log q_t(\mathbf{x})$ is the score function.

Moreover, Song et al. [46] also show that there exists a corresponding deterministic process whose trajectories share the same marginal probability densities $q_t(\mathbf{x})$ as those of (3), thus laying the foundation of efficient sampling via numerical ODE solvers [31, 55]:

$$d\mathbf{x}_t = \left[f(t)\mathbf{x}_t - \frac{1}{2}g^2(t)\nabla_{\mathbf{x}} \log q_t(\mathbf{x}_t) \right] dt. \quad (4)$$

We usually train a score network $\mathbf{s}_\theta(\mathbf{x}, t)$ parameterized by θ to approximate the score function $\nabla_{\mathbf{x}} \log q_t(\mathbf{x})$ in (3) by optimizing the denoising score matching loss [46]:

$$\mathcal{L} = \mathbb{E}_t \left\{ \omega(t) \mathbb{E}_{\mathbf{x}_0, \mathbf{x}_t} \left[\left\| \mathbf{s}_\theta(\mathbf{x}_t, t) - \nabla_{\mathbf{x}} \log q_{0t}(\mathbf{x}_t | \mathbf{x}_0) \right\|_2^2 \right] \right\}, \quad (5)$$

where $\omega(t)$ is a weighting function. In practice, two methods are commonly used to reparameterize the score network [22]. The first approach utilizes a *noise prediction network* ϵ_θ such that $\epsilon_\theta(\mathbf{x}, t) = -\sigma_t \mathbf{s}_\theta(\mathbf{x}, t)$, and the second approach employs a *data prediction network* \mathbf{x}_θ such that $\mathbf{x}_\theta(\mathbf{x}, t) = (\sigma_t^2 \mathbf{s}_\theta(\mathbf{x}, t) + \mathbf{x}) / \alpha_t$. In particular, for the data prediction network \mathbf{x}_θ , based on score approximation and (2), the reverse ODE in (4) can be expressed as²

$$\frac{d\mathbf{x}_t}{dt} = \frac{d \log \sigma_t}{dt} \mathbf{x}_t + \frac{\alpha_t d \log \frac{\alpha_t}{\sigma_t}}{dt} \mathbf{x}_\theta(\mathbf{x}_t, t), \quad (6)$$

with the initial vector \mathbf{x}_T being sampled from $\mathcal{N}(\mathbf{0}, \tilde{\sigma}^2 \mathbf{I})$.

3.2. Discretization Schemes

As mentioned above, we can numerically solve a reverse ODE such as (6) to perform the task of sampling from diffusion models. Different numerical methods are then employed to approximately solve the ODE, for which the time interval $[\epsilon, T]$ needs to be divided into N sub-intervals for some positive integer N using $N + 1$ time steps $T = t_0 > t_1 > \dots > t_N = \epsilon$.³ Let $\lambda_t = \log(\alpha_t / \sigma_t)$ denote the one half of the log-SNR. Popular discretization schemes for the time steps include the i) uniform- t scheme, ii) uniform- λ scheme [31], and the discretization strategy used in [19], which we refer to as the EDM scheme.

More specifically, for the uniform- t scheme, we split the interval $[\epsilon, T]$ uniformly and obtain

$$t_n = T + \frac{n}{N}(\epsilon - T) \text{ for } n = 0, 1, \dots, N. \quad (7)$$

Note that since λ_t is a strictly decreasing function of t , it has an inverse function $t_\lambda(\cdot)$. For the uniform- λ scheme,

²Throughout this work, we focus on the case of using a data prediction network \mathbf{x}_θ . For the case of using a noise prediction network ϵ_θ , we can utilize the formula $\mathbf{x}_\theta(\mathbf{x}, t) = \frac{\mathbf{x} - \sigma_t \epsilon_\theta(\mathbf{x}, t)}{\alpha_t}$ to transform it into the data prediction network.

³Practically, it is typical to use end time $t = \epsilon > 0$ instead of $t = 0$ to avoid numerical issues for small t , see [31, Appendix D.1]. The parameter ϵ is assumed to be fixed in our work. For the proper selection of ϵ , the reader may refer to prior works such as [20].

we split the interval $[\lambda_\epsilon, \lambda_T]$ uniformly, and convert the λ values back to t values. That is, we have

$$t_n = t_\lambda \left(\lambda_T + \frac{n}{N}(\lambda_\epsilon - \lambda_T) \right) \text{ for } n = 0, 1, \dots, N. \quad (8)$$

In addition, Karras et al. [19] propose to do a variable substitution $\kappa_t = \frac{\sigma_t}{\alpha_t}$, which is the reciprocal of the square root of SNR. For the EDM scheme, they propose to uniformly discretize the $\kappa_t^{\frac{1}{\rho}}$ for some positive integer ρ , i.e.,

$$t_n = t_\kappa \left(\left(\kappa_T^{\frac{1}{\rho}} + \frac{n}{N}(\kappa_\epsilon^{\frac{1}{\rho}} - \kappa_T^{\frac{1}{\rho}}) \right)^\rho \right) \text{ for } n = 0, 1, \dots, N, \quad (9)$$

where $t_\kappa(\cdot)$ denotes the inverse function of κ_t , which is guaranteed to exist as κ_t is strictly increasing with respect to t . Karras et al. [19] performed some ablation studies and suggested to set the parameter $\rho = 7$ in their experiments.

While the discretization scheme plays an important role in the numerical solvers, the above-mentioned schemes have been widely employed for convenience, leaving a significant room of improvements.

4. Problem Formulation

Recall that $\lambda_t = \log(\alpha_t / \sigma_t)$ is strictly decreasing with respect to t , and it has an inverse function $t_\lambda(\cdot)$. Then, the term $\mathbf{x}_\theta(\mathbf{x}_t, t)$ in (6) can be written as $\mathbf{x}_\theta(\mathbf{x}_{t_\lambda(\lambda)}, t_\lambda(\lambda)) = \hat{\mathbf{x}}_\theta(\hat{\mathbf{x}}_\lambda, \lambda)$. Let $\mathbf{f}(\lambda) = \hat{\mathbf{x}}_\theta(\hat{\mathbf{x}}_\lambda, \lambda)$. For the fixed starting time of sampling T and end time of sampling ϵ , the integral formula [32, Eq. (8)] gives the exact solution to the ODE in (6) at time ϵ :

$$\mathbf{x}_\epsilon = \frac{\sigma_\epsilon}{\sigma_T} \mathbf{x}_T + \sigma_\epsilon \int_{\lambda_T}^{\lambda_\epsilon} e^{\lambda} \mathbf{f}(\lambda) d\lambda.$$

However, $\mathbf{f}(\lambda)$ is a complicated neural network function and we cannot take the above integral exactly. Instead, we need to choose a sequence of time steps $T = t_0 > t_1 > \dots > t_N = \epsilon$ to split the integral as N parts and appropriately approximate the integral in each part separately. In particular, for any $n \in \{1, 2, \dots, N\}$ and a positive integer $k_n \leq n$, a k_n -th order explicit method for numerical ODE typically uses a $(k_n - 1)$ -th degree (local) polynomial $\mathcal{P}_{n; k_n - 1}(\lambda)$ that involves the function values of $\mathbf{f}(\cdot)$ at the points $\lambda_{n-k_n}, \lambda_{n-k_n+1}, \dots, \lambda_{n-1}$ to approximate $\mathbf{f}(\lambda)$ within the interval $[\lambda_{t_{n-1}}, \lambda_{t_n}]$. That is, the polynomial $\mathcal{P}_{n; k_n - 1}(\lambda)$ can be expressed as

$$\mathcal{P}_{n; k_n - 1}(\lambda) = \sum_{j=0}^{k_n - 1} \ell_{n; k_n, j}(\lambda) \mathbf{f}(\lambda_{n-k_n+j}), \quad (10)$$

where $\ell_{n; k_n, j}(\lambda)$ are certain $(k_n - 1)$ -th degree polynomials of λ .

Remark 1. If setting $\mathcal{P}_{n;k_n-1}(\lambda)$ to be the Lagrange polynomial with interpolation points $\lambda_{n-k_n}, \lambda_{n-k_n+1}, \dots, \lambda_{n-1}$, we obtain

$$\ell_{n;k_n,j}(\lambda) := \prod_{i=0, i \neq j}^{k_n-1} \frac{\lambda - \lambda_{n-k_n+i}}{\lambda_{n-k_n+j} - \lambda_{n-k_n+i}}. \quad (11)$$

We note that such an approximation has been essentially considered in the UniPC method proposed in [55], if not using the corrector and a special trick for the case of $k_n = 2$.⁴ In addition, a similar idea was considered in [54] with using local Lagrange polynomials to approximate $e^\lambda \mathbf{f}(\lambda)$.

Or alternatively, $\mathcal{P}_{n;k_n-1}(\lambda)$ may be set as the polynomial corresponding to the Taylor expansion of $\mathbf{f}(\lambda)$ at λ_{n-1} , as was done in the works [31, 32]. That is,

$$\mathcal{P}_{n;k_n-1}(\lambda) = \sum_{j=0}^{k_n-1} \mathbf{f}^{(j)}(\lambda_{n-1}) \frac{(\lambda - \lambda_{n-1})^j}{j!}, \quad (12)$$

where $\mathbf{f}^{(j)}(\lambda_{n-1})$ can be further approximated using the neural function values at $\lambda_{n-k_n}, \lambda_{n-k_n+1}, \dots, \lambda_{n-1}$ [32]. In particular, if letting $h_n = \lambda_{n+1} - \lambda_n$ for any $n \in [N]$, we have for $j = 1$ that $\mathbf{f}^{(j)}(\lambda_{n-1}) \approx \frac{\mathbf{f}(\lambda_{n-1}) - \mathbf{f}(\lambda_{n-2})}{h_{n-2}}$ and for $j = 2$ that $\mathbf{f}^{(j)}(\lambda_{n-1}) \approx \frac{2}{h_{n-2}(h_{n-2} + h_{n-3})} \mathbf{f}(\lambda_{n-1}) - \frac{2}{h_{n-2}h_{n-3}} \mathbf{f}(\lambda_{n-2}) + \frac{2}{h_{n-3}(h_{n-2} + h_{n-3})} \mathbf{f}(\lambda_{n-3})$. Therefore, we observe that Taylor expansion will also give local polynomials $\mathcal{P}_{n;k_n-1}(\lambda)$ that are of the form in (10), although they will be slightly different with those for the case of using Lagrange approximations.

Based on (10), we obtain an approximation of \mathbf{x}_ϵ as follows:

$$\mathbf{x}_\epsilon = \frac{\sigma_\epsilon}{\sigma_T} \mathbf{x}_T + \sigma_\epsilon \sum_{n=1}^N \int_{\lambda_{t_{n-1}}}^{\lambda_{t_n}} e^\lambda \mathbf{f}(\lambda) d\lambda \quad (13)$$

$$\approx \frac{\sigma_\epsilon}{\sigma_T} \mathbf{x}_T + \sigma_\epsilon \sum_{n=1}^N \int_{\lambda_{t_{n-1}}}^{\lambda_{t_n}} e^\lambda \mathcal{P}_{n;k_n-1}(\lambda) d\lambda \quad (14)$$

$$= \frac{\sigma_\epsilon}{\sigma_T} \mathbf{x}_T + \sigma_\epsilon \sum_{n=1}^N \int_{\lambda_{t_{n-1}}}^{\lambda_{t_n}} e^\lambda \left(\sum_{j=0}^{k_n-1} \ell_{n;k_n,j}(\lambda) \mathbf{f}(\lambda_{n-k_n+j}) \right) d\lambda \quad (15)$$

$$= \frac{\sigma_\epsilon}{\sigma_T} \mathbf{x}_T + \sigma_\epsilon \sum_{n=1}^N \sum_{j=0}^{k_n-1} w_{n;k_n,j} \mathbf{f}(\lambda_{n-k_n+j}) \quad (16)$$

$$:= \tilde{\mathbf{x}}_\epsilon, \quad (17)$$

where

$$w_{n;k_n,j} = \int_{\lambda_{t_{n-1}}}^{\lambda_{t_n}} e^\lambda \ell_{n;k_n,j}(\lambda) d\lambda. \quad (18)$$

⁴While the iterative formula in [55] appears to be different, we have numerically calculated that (11) leads to the same coefficients for the neural function values as those for [55].

It is worth noting that we would expect that the weights $w_{n;k_n,j}$ satisfy the following

$$\sum_{j=0}^{k_n-1} w_{n;k_n,j} = e^{\lambda_n} - e^{\lambda_{n-1}} \quad (19)$$

for all $n \in \{1, 2, \dots, N\}$ since for the simplest case that $\mathbf{f}(\lambda)$ is a vector of all ones, its approximation $\mathcal{P}_{n;k_n-1}(\lambda)$ should also be a vector of all ones, and then (19) follows from (14) and (16). The condition in (19) can also be easily verified to hold for the case of using Lagrange and Taylor expansion polynomials discussed in Remark 1.

For the vector $\tilde{\mathbf{x}}_\epsilon$ defined in (17), since our goal is to find an estimated vector that is close to the ground-truth solution \mathbf{x}_0 , a natural question is as follows:

Question: Can we design the time steps t_1, t_2, \dots, t_{N-1} (or equivalently, $\lambda_{t_1}, \lambda_{t_2}, \dots, \lambda_{t_{N-1}}$) such that the distance between the ground-truth solution \mathbf{x}_0 and the approximate solution $\tilde{\mathbf{x}}_\epsilon$ is minimized?

Remark 2. The definition of $\tilde{\mathbf{x}}_\epsilon$ in (17) is fairly general in the sense that for any explicit numerical ODE solver with corresponding local polynomials $\mathcal{P}_{n;k_n-1}(\lambda)$ of the form (10), we can easily calculate all the weights $w_{n;k_n,j}$ via (18). In particular, we allow the (local) orders k_n to vary with respect to n to encompass general cases for high-order numerical solvers. For example, for a third-order method, we have $k_n = 1$ for $n = 1$, $k_n = 2$ for $n = 2$, and $k_n = 3$ for $n \geq 3$. The varying k_n can also handle the customized order schedules tested in [55].

5. Analysis and Method

To partially answer the question presented in Section 4, we follow relevant theoretical works [4, 6–9, 25, 26, 37] to make the following assumption on the score approximation.

Assumption 1. For any $t \in \{t_0, t_1, \dots, t_N\}$, the error in the score estimate is bounded in $L^2(q_t)$:

$$\begin{aligned} & \|\nabla_{\mathbf{x}} \log q_t - \mathbf{s}_\theta(\cdot, t)\|_{L^2(q_t)}^2 \\ &= \mathbb{E}_{q_t} [\|\nabla_{\mathbf{x}} \log q_t(\mathbf{x}) - \mathbf{s}_\theta(\mathbf{x}, t)\|^2] \leq \eta^2 \epsilon_t^2, \end{aligned}$$

where $\eta > 0$ is an absolute constant.

More specifically, ϵ_t can be set to be 1 in the corresponding assumptions in [7, 8, 25], and ϵ_t can be set to be $\frac{1}{\sigma_t^2}$ in [26, Assumption 1]. Based on Assumption 1 and the transform between score network \mathbf{s}_θ and data prediction network \mathbf{x}_θ [22], we obtain the following lemma.

Lemma 1. For any $\mathbf{x}_0 \sim q_0$ and $P_0 \in (0, 1)$, with probability at least $1 - P_0$, the following event occurs: For all $t \in \{t_0, t_1, \dots, t_N\}$ and $\mathbf{x}_t \sim q_t$, we have

$$\|\mathbf{x}_\theta(\mathbf{x}_t, t) - \mathbf{x}_0\| \leq \tilde{\eta} \epsilon_t,$$

where $\tilde{\eta} := \sqrt{\frac{N+1}{P_0}} \eta$ and $\tilde{\varepsilon}_t := \frac{\varepsilon_t \sigma_t^2}{\alpha_t}$.

Proof. For any $u > 0$ and any $t \in \{t_0, t_1, \dots, t_N\}$, by the Markov inequality, we obtain

$$\begin{aligned} & \mathbb{P}(\|\nabla_{\mathbf{x}} \log q_t(\mathbf{x}_t) - \mathbf{s}_\theta(\mathbf{x}_t, t)\| > u) \\ &= \mathbb{P}(\|\nabla_{\mathbf{x}} \log q_t(\mathbf{x}_t) - \mathbf{s}_\theta(\mathbf{x}_t, t)\|^2 > u^2) \end{aligned} \quad (20)$$

$$\leq \frac{\mathbb{E}_{q_t} [\|\nabla_{\mathbf{x}} \log q_t(\mathbf{x}_t) - \mathbf{s}_\theta(\mathbf{x}_t, t)\|^2]}{u^2} \quad (21)$$

$$\leq \frac{\eta^2 \varepsilon_t^2}{u^2}, \quad (22)$$

where (22) follows from Assumption 1. Taking a union bound over all $t \in \{t_0, t_1, \dots, t_N\}$, we obtain that with probability at least $1 - \frac{(N+1)\eta^2 \varepsilon_t^2}{u^2}$, it holds for all $t \in \{t_0, t_1, \dots, t_N\}$ that

$$\|\nabla_{\mathbf{x}} \log q_t(\mathbf{x}_t) - \mathbf{s}_\theta(\mathbf{x}_t, t)\| \leq u. \quad (23)$$

Using the transforms $\nabla_{\mathbf{x}} \log q_t(\mathbf{x}_t) = \frac{\alpha_t \mathbf{x}_0 - \mathbf{x}_t}{\sigma_t^2}$ and $\mathbf{s}_\theta(\mathbf{x}_t, t) = \frac{\alpha_t \mathbf{x}_\theta(\mathbf{x}_t, t) - \mathbf{x}_t}{\sigma_t^2}$ [22], we derive that with probability at least $1 - \frac{(N+1)\eta^2 \varepsilon_t^2}{u^2}$, it holds for all $t \in \{t_0, t_1, \dots, t_N\}$ that

$$\|\mathbf{x}_0 - \mathbf{x}_\theta(\mathbf{x}_t, t)\| \leq \frac{\sigma_t^2}{\alpha_t} u. \quad (24)$$

Setting $u = \sqrt{\frac{N+1}{P_0}} \eta \varepsilon_t$, we obtain the desired result. \square

Note that we have $\tilde{\varepsilon}_t = \frac{\sigma_t^2}{\alpha_t}$ for the case that $\varepsilon_t = 1$, and $\tilde{\varepsilon}_t = \frac{1}{\alpha_t}$ for the case that $\varepsilon_t = \frac{1}{\sigma_t^2}$. Additionally, motivated by the training objective for the data prediction network that aims to minimize $\frac{\alpha_t}{\sigma_t} \|\mathbf{x}_0 - \mathbf{x}_\theta(\mathbf{x}_t, t)\|$ for $\mathbf{x}_0 \sim q_0$ and $\mathbf{x}_t \sim q_t$ [42], it is also natural to consider the case that $\tilde{\varepsilon}_t = \frac{\sigma_t}{\alpha_t}$. Therefore, in our experiments, we will set $\tilde{\varepsilon}_t$ as $\frac{\sigma_t^p}{\alpha_t}$ for some non-negative integer p .

Based on Lemma 1, we establish the following theorem, which provides an upper bound on the distance between the ground-truth solution \mathbf{x}_0 and the approximate solution $\tilde{\mathbf{x}}_\epsilon$ defined in (17).

Theorem 1. *Let the approximate solution $\tilde{\mathbf{x}}_\epsilon$ be defined as in (17) with $\mathbf{f}(\lambda) = \hat{\mathbf{x}}_\theta(\hat{\mathbf{x}}(\lambda), \lambda)$ and the weights $w_{n;k_n,j}$ satisfying (19). Suppose that Assumption 1 is satisfied for the score approximation. Then, for any $P_0 \in (0, 1)$, we have with probability at least $1 - P_0$ that*

$$\begin{aligned} \|\tilde{\mathbf{x}}_\epsilon - \mathbf{x}_0\| &\leq \left\| \frac{\sigma_\epsilon}{\sigma_T} \mathbf{x}_T + \sigma_\epsilon (e^{\lambda_\epsilon} - e^{\lambda_T}) \mathbf{x}_0 - \mathbf{x}_0 \right\| \\ &\quad + \sigma_\epsilon \tilde{\eta} \sum_{i=0}^{N-1} \tilde{\varepsilon}_{t_i} \cdot \left| \sum_{n-k_n+j=i} w_{n;k_n,j} \right|, \end{aligned} \quad (25)$$

where $\tilde{\eta} := \sqrt{\frac{N+1}{P_0}} \eta$ and $\tilde{\varepsilon}_t := \frac{\varepsilon_t \sigma_t^2}{\alpha_t}$.

Proof. From Lemma 1, we have with probability at least $1 - P_0$, it holds for all $n \in \{0, 1, \dots, N\}$ that $\mathbf{f}(\lambda_{t_n})$ can be written as

$$\mathbf{f}(\lambda_{t_n}) = \hat{\mathbf{x}}_\theta(\hat{\mathbf{x}}_{\lambda_{t_n}}, \lambda_{t_n}) = \mathbf{x}_\theta(\mathbf{x}_{t_n}, t_n) \quad (26)$$

$$= \mathbf{x}_0 + \boldsymbol{\xi}_{t_n}, \quad (27)$$

where $\boldsymbol{\xi}_{t_n}$ satisfies $\|\boldsymbol{\xi}_{t_n}\| \leq \tilde{\eta} \tilde{\varepsilon}_{t_n}$. Then, we obtain

$$\tilde{\mathbf{x}}_\epsilon = \frac{\sigma_\epsilon}{\sigma_T} \mathbf{x}_T + \sigma_\epsilon \sum_{n=1}^N \sum_{j=0}^{k_n-1} w_{n;k_n,j} \mathbf{f}(\lambda_{t_{n-k_n+j}}) \quad (28)$$

$$= \frac{\sigma_\epsilon}{\sigma_T} \mathbf{x}_T + \sigma_\epsilon \sum_{n=1}^N \sum_{j=0}^{k_n-1} w_{n;k_n,j} (\mathbf{x}_0 + \boldsymbol{\xi}_{t_{n-k_n+j}}) \quad (29)$$

$$\begin{aligned} &= \frac{\sigma_\epsilon}{\sigma_T} \mathbf{x}_T + \sigma_\epsilon \sum_{n=1}^N \sum_{j=0}^{k_n-1} w_{n;k_n,j} \mathbf{x}_0 \\ &\quad + \sigma_N \sum_{n=1}^N \sum_{j=0}^{k_n-1} w_{n;k_n,j} \boldsymbol{\xi}_{t_{n-k_n+j}} \end{aligned} \quad (30)$$

$$\begin{aligned} &= \frac{\sigma_\epsilon}{\sigma_T} \mathbf{x}_T + \sigma_\epsilon \sum_{n=1}^N (e^{\lambda_n} - e^{\lambda_{n-1}}) \mathbf{x}_0 \\ &\quad + \sigma_N \sum_{n=1}^N \sum_{j=0}^{k_n-1} w_{n;k_n,j} \boldsymbol{\xi}_{t_{n-k_n+j}} \end{aligned} \quad (31)$$

$$= \frac{\sigma_\epsilon}{\sigma_T} \mathbf{x}_T + \sigma_\epsilon (e^{\lambda_\epsilon} - e^{\lambda_T}) \mathbf{x}_0 + \sigma_N \sum_{i=0}^{N-1} \sum_{n-k_n+j=i} w_{n;k_n,j} \boldsymbol{\xi}_{t_i}, \quad (32)$$

where (29) follows from (27), (31) follows from (19), and (32) follows from the condition that $j+1 \leq k_n \leq n$. Therefore, from the triangle inequality, we have

$$\begin{aligned} \|\tilde{\mathbf{x}}_\epsilon - \mathbf{x}_0\| &\leq \left\| \frac{\sigma_\epsilon}{\sigma_T} \mathbf{x}_T + \sigma_\epsilon (e^{\lambda_\epsilon} - e^{\lambda_T}) \mathbf{x}_0 - \mathbf{x}_0 \right\| \\ &\quad + \sigma_\epsilon \sum_{i=0}^{N-1} \left| \sum_{n-k_n+j=i} w_{n;k_n,j} \right| \|\boldsymbol{\xi}_{t_i}\| \end{aligned} \quad (33)$$

$$\begin{aligned} &\leq \left\| \frac{\sigma_\epsilon}{\sigma_T} \mathbf{x}_T + \sigma_\epsilon (e^{\lambda_\epsilon} - e^{\lambda_T}) \mathbf{x}_0 - \mathbf{x}_0 \right\| \\ &\quad + \sigma_\epsilon \tilde{\eta} \sum_{i=0}^{N-1} \tilde{\varepsilon}_{t_i} \cdot \left| \sum_{n-k_n+j=i} w_{n;k_n,j} \right|, \end{aligned} \quad (34)$$

which gives the desired upper bound. \square

The starting time of sampling T and the end time of sampling ϵ are fixed, and we typically have $\sigma_\epsilon \approx 0$ and $\alpha_T \approx 0$. Then, the first term in the upper bound of (25) will be fixed and small. Note that $\{\lambda_{t_n}\}_{n=0}^N$ needs to be a monotonically increasing sequence. Additionally, since $\tilde{\eta}$ is a fixed positive constant, we observe from (25) that

Algorithm 1 Finding the time steps via (35)

Input: Number of time steps N , initial time of sampling T , end time of sampling ϵ , any sampling algorithm that is characterized by local polynomials $\mathcal{P}_{n;k_n-1}(\lambda)$ of the form (10), the function $\tilde{\epsilon}_t := \frac{\sigma_t^p}{\alpha_t}$ with a fixed positive integer p for score approximation, the initial values of $\lambda_{t_1}, \dots, \lambda_{t_{N-1}}$

1: Set $\lambda_{t_0} = \lambda_T$ and $\lambda_{t_N} = \lambda_\epsilon$ and calculate $\tilde{\epsilon}_{t_i}$ for $i = 0, 1, \dots, N - 1$

2: Calculate the weights $w_{n;k_n,j}$ from (18)

3: Solve the optimization problem in (35) via the constrained trust region method

Output: Optimized λ (or equivalently, time) steps $\hat{\lambda}_{t_1}, \dots, \hat{\lambda}_{t_{N-1}}$

we should choose appropriate $\lambda_{t_1}, \lambda_{t_2}, \dots, \lambda_{t_{N-1}}$ (which lead to the corresponding time steps t_1, t_2, \dots, t_{N-1}) such that $\sum_{i=0}^{N-1} \tilde{\epsilon}_{t_i} \cdot \left| \sum_{n-k_n+j=i} w_{n;k_n,j} \right|$ is minimized, which gives the following optimization problem:

$$\begin{aligned} \min_{\lambda_{t_1}, \dots, \lambda_{t_{N-1}}} \quad & \sum_{i=0}^{N-1} \tilde{\epsilon}_{t_i} \cdot \left| \sum_{n-k_n+j=i} w_{n;k_n,j} \right| \\ \text{s.t.} \quad & \lambda_{t_{n+1}} > \lambda_{t_n}, \text{ for } n = 0, 1, \dots, N - 1, \end{aligned} \quad (35)$$

where $\lambda_{t_0} = \lambda_T$ and $\lambda_{t_N} = \lambda_\epsilon$ are fixed. Note that both $\tilde{\epsilon}_{t_i}$ and $w_{n;k_n,j}$ are functions dependent on $\lambda_{t_1}, \dots, \lambda_{t_{N-1}}$. More specifically, for any sampling algorithm that uses the local approximation polynomials $\mathcal{P}_{n;k_n-1}(\lambda)$ of the form (10), the weights $w_{n;k_n,j}$ can be easily calculated from (18). In addition, as mentioned above, we can set $\tilde{\epsilon}_{t_i}$ to be $\frac{\sigma_{t_i}^p}{\alpha_{t_i}}$ for some non-negative integer p ,⁵ which can be represented as a functions of λ_{t_i} . Furthermore, the optimization problem in (35) can be approximately solved highly efficiently using the constrained trust region method.⁶ Our method is summarized in Algorithm 1.

6. Experiments

In this section, we demonstrate the effectiveness of our method combined with the currently most popular ODE solvers DPM-Solver++ [32] and UniPC [55]⁷ for the case

⁵We find in the experiments that setting $p = 1$ is a good choice for pixel-space generation, and setting $p = 2$ is appropriate for latent-space generation.

⁶Such an algorithm has been implemented in standard Python libraries such as `scipy`. See <https://docs.scipy.org/doc/scipy/tutorial/optimize.html#defining-linear-constraints>.

⁷We observe from the experiments that our optimized time steps work effectively for UniPC, although the optimization problem does not take the corrector of UniPC into account, which corresponds to implicit numerical ODE solvers.

of using 5 \sim 15 neural function evaluations (NFEs) with extensive experiments on various datasets. The order is 3 for both DPM-Solver++ and UniPC. We use Fenchel Inception Distance (FID) [14] as the evaluation metric to show the effectiveness of our method. Unless otherwise specified, 50K images are sampled for evaluation. The experiments are conducted on a wide range of datasets, with image sizes ranging from 32x32 to 512x512. We also evaluate the performance of various previous state-of-the-art pre-trained diffusion models, including Score-SDE [46], ADM [11], EDM [19], and DiT [38]. which cover pixel-space and latent-space diffusion models, with unconditional generation, conditional generation and generation with classifier-free guidance settings.

6.1. Pixel Diffusion Model Generation

Results on CIFAR-10 32x32 For the CIFAR-10 32x32 experiment, We use the *ddpmpp_deep_continuous* Score-SDE model from [46], which is an unconditional VP-schedule model. The experiment results for UniPC are shown in Fig. 1. Our proposed optimized sampling schedule consistently outperforms the other three baseline sampling schedules and achieves state-of-the-art FID results. For example, we improve the FID on CIFAR-10 with only 5 NFEs from 23.22 to 12.11. We also conducted the experiment for DPM-Solver++ on CIFAR-10, see Table 1. We found that although our optimized sampling schedule also consistently improves DPM-Solver++, UniPC still performs better no matter which sampling schedule is used. Therefore, in the remaining part of this section, we only report the results for UniPC. Due to the page limit, some additional experimental results have been deferred to the supplementary material.

Results on ImageNet 64x64 For the ImageNet 64x64 experiment, we use the *64x64 diffusion* model from ADM [11], which is a conditional VP-schedule model. The experimental results are shown in Fig. 1. It can be noticed that among the baseline methods, the uniform- t performs the best for NFEs 5 \sim 6, and uniform- λ performs best for NFEs > 7 . This indicates that the previous baseline sampling schedules may be far from optimal. Our proposed optimized sampling schedule consistently outperforms the other three baseline sampling schedules and achieves state-of-the-art FID results. For example, we achieved the FID of 10.47 on ImageNet 64x64 with only 5 NFEs.

Results on FFHQ 64x64 and AFHQv2 64x64 For the FFHQ 64x64 and AFHQv2 64x64 experiments, we use the EDM [19] unconditional model. The experiment results are shown in Fig. 1. For the EDM model, the time ranges from 0.002 to 80 during sampling rather than 0 to 1, which makes the result of uniform- t significantly worse than the other two baseline discretization schemes. Thus we do

Methods \ NFEs	5	6	7	8	9	10	12	15
DPM-Solver++ with uniform- λ	29.22	13.28	7.18	5.12	4.40	4.03	3.45	3.17
DPM-Solver++ with uniform- t	28.16	19.63	15.29	12.58	11.18	10.15	8.50	7.10
DPM-Solver++ with EDM	40.48	25.10	15.68	10.22	7.42	6.18	4.85	3.49
DPM-Solver++ with optimized step (Ours)	12.91	8.35	5.44	4.74	3.81	3.51	3.24	3.15
UniPC with uniform- λ	23.22	10.33	6.18	4.80	4.19	3.87	3.34	3.17
UniPC with uniform- t	25.11	17.40	13.54	11.33	9.83	8.89	7.38	6.18
UniPC with EDM	38.24	23.79	14.62	8.95	6.60	5.59	4.18	3.16
UniPC with optimized step (Ours)	12.11	7.23	4.96	4.46	3.75	3.50	3.19	3.13

Table 1. Sampling quality measured by FID (\downarrow) of different discretization schemes of time steps for DPM-Solver++ [32] and UniPC [55] with varying NFEs on CIFAR-10.

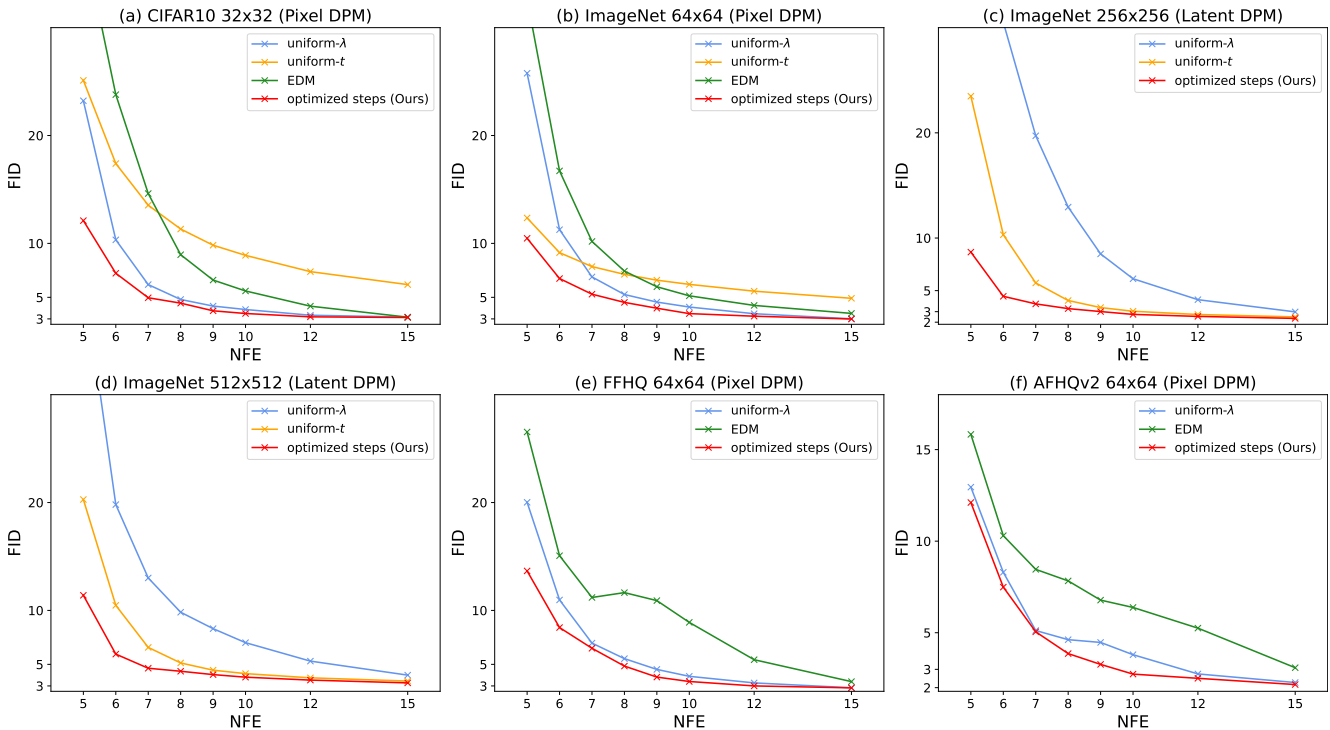


Figure 1. Sampling quality measured by FID (\downarrow) of different discretization schemes of time steps for UniPC [55] with varying NFEs on various DPMs and datasets.

not include the results of the uniform- t scheme. In comparison, our proposed optimized sampling schedule consistently outperforms the other two baseline sampling schedules and achieves state-of-the-art FID results. For example, we achieved an FID of 13.66 on FFHQ 64x64 and an FID of 12.11 on AFHQv2 64x64 with only 5 NFEs.

6.2. Latent Diffusion Model Generation

Results on ImageNet 256x256 and ImageNet 512x512

We evaluate our method on the ImageNet dataset using

the DiT-XL-2 [38] model, which is a Vision-Transformer-based model in the latent space of KL-8 encoder-decoder. The corresponding classifier-free guidance scale, namely $s = 1.5$, is adopted for evaluation. The experiment results are shown in Fig. 1. Our proposed optimized sampling schedule consistently outperforms the other two baseline sampling schedules and achieves state-of-the-art FID results. For example, we improve the FID from 23.48 to 8.66 on ImageNet 256x256 and from 20.28 to 11.40 on ImageNet 512x512 with only 5 NFEs. Figure 2 visu-

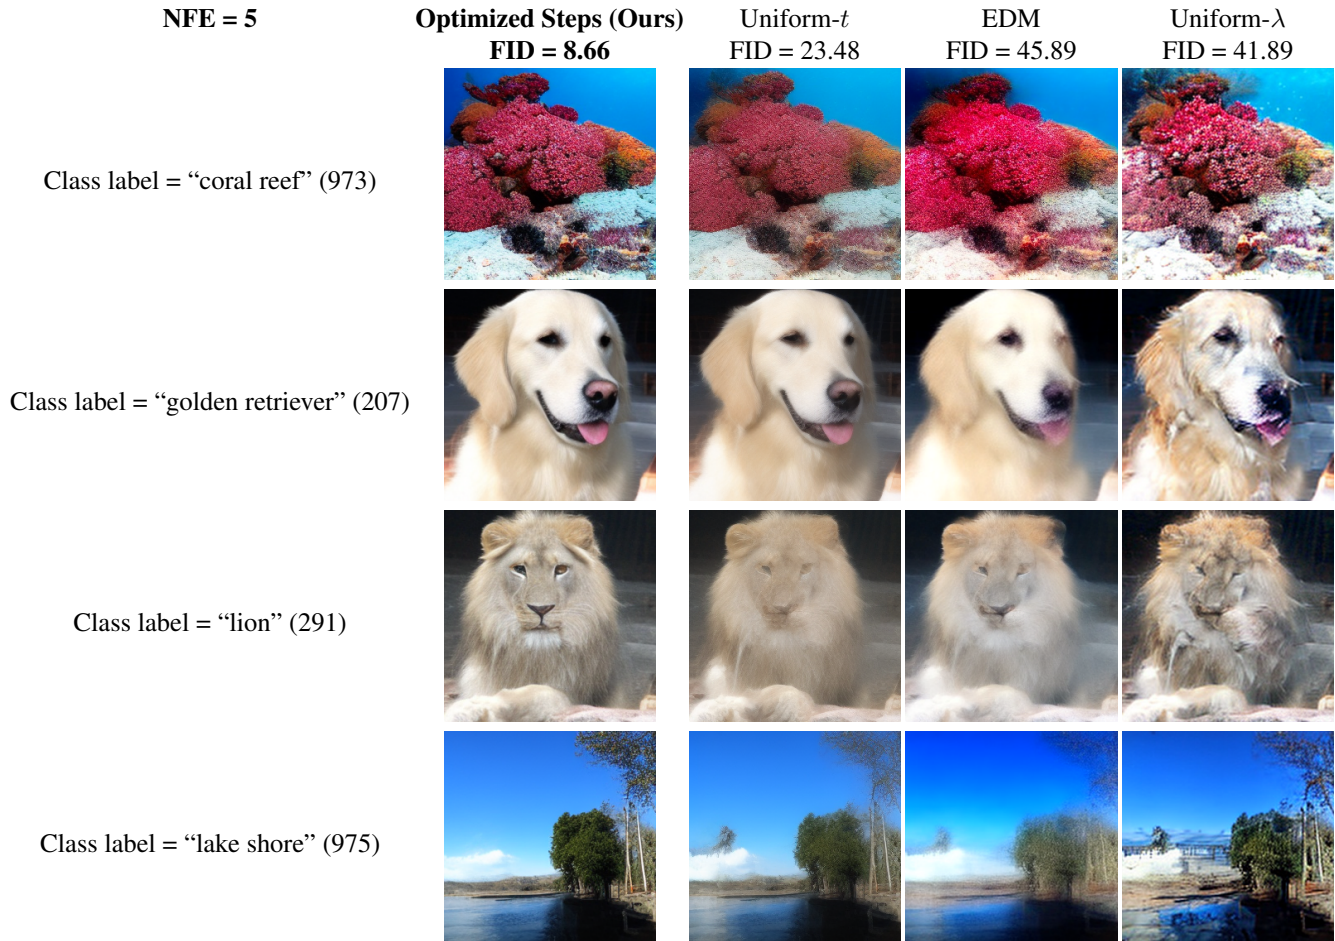


Figure 2. Generated images by UniPC [55] with only 5 NFEs for various discretization schemes of time steps from DiT-XL-2 ImageNet 256x256 model [38] (with cfg scale $s = 1.5$ and the same random seed).

ally demonstrates the effectiveness of our proposed method. Since these experiments are a bit time-consuming, and the EDM discretization scheme for time steps has been shown to not work well for small NFEs, we mainly use uniform- t and uniform- λ schemes as baselines.

6.3. Running Time Analysis

We test the running time of our Algorithm 1 on *Intel(R) Xeon(R) Gold 6278C CPU @ 2.60GHz*. We report the longest running time observed for performing Algorithm 1 for 5, 6, 7, 8, 9, 10, 12, 15 NFEs. The experiment results are shown in Table 2. Our algorithm can be pre-computed before sampling and the optimization result can be reutilized. The result shows that our optimization problem can be solved within 15 seconds for NFEs less than or equal to 15, which is negligible. In comparison, the learning to schedule methods [27, 30, 48, 50, 51] usually takes several GPU hours for optimization and the overall performance is comparable to ours.

NFEs	5	6	7	8	9	10	12	15
Time(s)	1.9	2.3	5.3	5.9	7.8	8.8	11.0	14.1

Table 2. Running time of our optimization algorithm.

7. Conclusion

In this paper, we propose an optimization-based method to find appropriate time steps to accelerate the sampling of diffusion models, and thus generating high-quality images in a small number of sampling steps. We formulate the problem as a surrogate optimization method with respect to the time steps, which can be efficiently solved via the constrained trust region method. Experimental results on popular image datasets demonstrate that our method can be employed in a plug-and-play manner and achieves state-of-the-art sampling performance based on various pre-trained diffusion models. Our work solves an approximated optimization object which can be further improved if a more accurate formulation can be found.

References

- [1] Brian D.O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982. 2
- [2] 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. 1
- [3] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1
- [4] Hongrui Chen, Holden Lee, and Jianfeng Lu. Improved analysis of score-based generative modeling: User-friendly bounds under minimal smoothness assumptions. In *International Conference on Machine Learning*, pages 4735–4763. PMLR, 2023. 4
- [5] Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Yue Wu, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart- α : Fast training of diffusion transformer for photorealistic text-to-image synthesis, 2023. 1, 2, 4, 5, 6, 7
- [6] Minshuo Chen, Kaixuan Huang, Tuo Zhao, and Mengdi Wang. Score approximation, estimation and distribution recovery of diffusion models on low-dimensional data. *arXiv preprint arXiv:2302.07194*, 2023. 4
- [7] Sitan Chen, Sinho Chewi, Holden Lee, Yuanzhi Li, Jianfeng Lu, and Adil Salim. The probability flow ode is provably fast. *arXiv preprint arXiv:2305.11798*, 2023. 4
- [8] Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. In *International Conference on Learning Representations*, 2023. 4
- [9] Valentin De Bortoli. Convergence of denoising diffusion models under the manifold hypothesis. *Transactions on Machine Learning Research*, 2022. 4
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255. IEEE, 2009. 2
- [11] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems*, pages 8780–8794, 2021. 1, 6
- [12] Yansong Gao, Zhihong Pan, Xin Zhou, Le Kang, and Pratik Chaudhari. Fast diffusion probabilistic model sampling through the lens of backward error analysis. *arXiv preprint arXiv:2304.11446*, 2023. 2
- [13] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014. 1
- [14] 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*, pages 6626–6637, 2017. 6
- [15] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, pages 6840–6851, 2020. 1, 2
- [16] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 23(47):1–33, 2022. 1
- [17] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. In *Advances in Neural Information Processing Systems*, 2022. 1
- [18] Alexia Jolicoeur-Martineau, Ke Li, Rémi Piché-Taillefer, Tal Kachman, and Ioannis Mitliagkas. Gotta go fast when generating data with score-based models. *arXiv preprint arXiv:2105.14080*, 2021. 1, 2
- [19] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Proc. NeurIPS*, 2022. 1, 3, 6
- [20] Dongjun Kim, Seungjae Shin, Kyungwoo Song, Wanmo Kang, and Il-Chul Moon. Soft truncation: A universal training technique of score-based diffusion model for high precision score estimation. In *ICML*, pages 11201–11228. PMLR, 2022. 3
- [21] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014. 1
- [22] Diederik P Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. In *Advances in Neural Information Processing Systems*, 2021. 2, 3, 4, 5
- [23] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. The CIFAR-10 Dataset. *online: <http://www.cs.toronto.edu/kriz/cifar.html>*, 55, 2014. 2
- [24] Max WY Lam, Jun Wang, Dan Su, and Dong Yu. Bddm: Bilateral denoising diffusion models for fast and high-quality speech synthesis. In *International Conference on Learning Representations*, 2022. 1
- [25] Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence for score-based generative modeling with polynomial complexity. *Advances in Neural Information Processing Systems*, 35: 22870–22882, 2022. 4
- [26] Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence of score-based generative modeling for general data distributions. In *International Conference on Algorithmic Learning Theory*, pages 946–985. PMLR, 2023. 4
- [27] Lijiang Li, Huixia Li, Xiawu Zheng, Jie Wu, Xuefeng Xiao, Rui Wang, Min Zheng, Xin Pan, Fei Chao, and Rongrong Ji. Autodiffusion: Training-free optimization of time steps and architectures for automated diffusion model acceleration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7105–7114, 2023. 2, 8
- [28] Shigui Li, Wei Chen, and Delu Zeng. Scire-solver: Accelerating diffusion models sampling by score-integrand solver with recursive difference. 2023. 2

- [29] 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 *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. [1](#)
- [30] Enshu Liu, Xuefei Ning, Zinan Lin, Huazhong Yang, and Yu Wang. Oms-dpm: Optimizing the model schedule for diffusion probabilistic models. *arXiv preprint arXiv:2306.08860*, 2023. [2](#), [8](#)
- [31] 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. In *Advances in Neural Information Processing Systems*, pages 5775–5787, 2022. [1](#), [2](#), [3](#), [4](#)
- [32] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models, 2023. [2](#), [3](#), [4](#), [6](#), [7](#)
- [33] Eric Luhman and Troy Luhman. Knowledge distillation in iterative generative models for improved sampling speed. *arXiv preprint arXiv:2101.02388*, 2021. [1](#), [2](#)
- [34] Weijian Luo, Tianyang Hu, Shifeng Zhang, Jiacheng Sun, Zhenguo Li, and Zhihua Zhang. Diff-instruct: A universal approach for transferring knowledge from pre-trained diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024.
- [35] Chenlin Meng, Ruiqi Gao, Diederik P Kingma, Stefano Ermon, Jonathan Ho, and Tim Salimans. On distillation of guided diffusion models. In *NeurIPS 2022 Workshop on Score-Based Methods*, 2022. [1](#), [2](#)
- [36] Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. GLIDE: towards photorealistic image generation and editing with text-guided diffusion models. In *International Conference on Machine Learning (ICML)*, 2022. [1](#)
- [37] Francesco Pedrotti, Jan Maas, and Marco Mondelli. Improved convergence of score-based diffusion models via prediction-correction. *arXiv preprint arXiv:2305.14164*, 2023. [4](#)
- [38] William Peebles and Saining Xie. Scalable diffusion models with transformers. *arXiv preprint arXiv:2212.09748*, 2022. [6](#), [7](#), [8](#), [2](#), [4](#)
- [39] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents, 2022. [1](#)
- [40] 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 (CVPR)*, pages 10684–10695, 2022. [1](#)
- [41] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems*, 2022. [1](#)
- [42] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2022. [1](#), [2](#), [5](#)
- [43] 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. [1](#), [2](#)
- [44] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. [1](#)
- [45] Kaitao Song, Yichong Leng, Xu Tan, Yicheng Zou, Tao Qin, and Dongsheng Li. Transcormer: Transformer for sentence scoring with sliding language modeling. In *Advances in Neural Information Processing Systems*, 2022. [1](#)
- [46] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. [1](#), [2](#), [3](#), [6](#)
- [47] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *arXiv preprint arXiv:2303.01469*, 2023. [1](#), [2](#)
- [48] Yunke Wang, Xiyu Wang, Anh-Dung Dinh, Bo Du, and Charles Xu. Learning to schedule in diffusion probabilistic models. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2478–2488, 2023. [2](#), [8](#)
- [49] Zhendong Wang, Huangjie Zheng, Pengcheng He, Weizhu Chen, and Mingyuan Zhou. Diffusion-GAN: Training GANs with diffusion. In *The Eleventh International Conference on Learning Representations*, 2023. [1](#), [2](#)
- [50] Daniel Watson, William Chan, Jonathan Ho, and Mohammad Norouzi. Learning fast samplers for diffusion models by differentiating through sample quality. In *International Conference on Learning Representations*, 2022. [1](#), [2](#), [8](#)
- [51] Mengfei Xia, Yujun Shen, Changsong Lei, Yu Zhou, Ran Yi, Deli Zhao, Wenping Wang, and Yong-jin Liu. Towards more accurate diffusion model acceleration with a timestep aligner. *arXiv preprint arXiv:2310.09469*, 2023. [2](#), [8](#)
- [52] Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with denoising diffusion GANs. In *International Conference on Learning Representations*, 2022. [1](#), [2](#)
- [53] Shuchen Xue, Mingyang Yi, Weijian Luo, Shifeng Zhang, Jiacheng Sun, Zhenguo Li, and Zhi-Ming Ma. Sa-solver: Stochastic adams solver for fast sampling of diffusion models, 2023. [1](#)
- [54] Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator. In *The Eleventh International Conference on Learning Representations*, 2023. [1](#), [4](#)
- [55] Wenliang Zhao, Lujia Bai, Yongming Rao, Jie Zhou, and Jiwen Lu. Unipc: A unified predictor-corrector framework for fast sampling of diffusion models. *arXiv preprint arXiv:2302.04867*, 2023. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#), [5](#)