

# A-FloPS: Accelerating Diffusion Models via Adaptive Flow Path Sampler

Cheng Jin, Zhenyu Xiao, Yuantao Gu\*

Department of Electronic Engineering  
Tsinghua University  
Beijing 100084  
P. R. China

jinc21@mails.tsinghua.edu.cn, xzy23@mails.tsinghua.edu.cn, gyt@tsinghua.edu.cn

## Abstract

Diffusion models deliver state-of-the-art generative performance across diverse modalities but remain computationally expensive due to their inherently iterative sampling process. Existing training-free acceleration methods typically improve numerical solvers for the reverse-time ODE, yet their effectiveness is fundamentally constrained by the inefficiency of the underlying sampling trajectories. We propose A-FloPS (Adaptive Flow Path Sampler), a principled, training-free framework that reparameterizes the sampling trajectory of any pre-trained diffusion model into a flow-matching form and augments it with an adaptive velocity decomposition. The reparameterization analytically maps diffusion scores to flow-compatible velocities, yielding integration-friendly trajectories without retraining. The adaptive mechanism further factorizes the velocity field into a linear drift term and a residual component whose temporal variation is actively suppressed, restoring the accuracy benefits of high-order integration even in extremely low-NFE regimes. Extensive experiments on conditional image generation and text-to-image synthesis show that A-FloPS consistently outperforms state-of-the-art training-free samplers in both sample quality and efficiency. Notably, with as few as 5 function evaluations, A-FloPS achieves substantially lower FID and generates sharper, more coherent images. The adaptive mechanism also improves native flow-based generative models, underscoring its generality. These results position A-FloPS as a versatile and effective solution for high-quality, low-latency generative modeling.

## Introduction

Diffusion models have achieved remarkable generative performance across diverse modalities, producing high-fidelity samples that often rival real-world data (Rombach et al. 2022; Esser et al. 2024; Ramesh et al. 2022). However, their practical deployment is hindered by substantial sampling latency due to the inherently iterative nature of the generation process. Producing a single high-quality sample typically requires hundreds of network evaluations, making inference computationally expensive and limiting applicability in real-time or resource-constrained scenarios (Ho, Jain, and Abbeel 2020; Song et al. 2021).

\*Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

To address this bottleneck, a variety of acceleration strategies have been developed, which can be broadly categorized into two groups. The first group comprises *training-required* approaches, such as knowledge distillation (Salimans and Ho 2022; Schusterbauer et al. 2025)—which compress a teacher model’s multi-step generation process into a student model capable of producing comparable results in far fewer steps—and consistency models (Song et al. 2024; Sauer et al. 2023), which train networks to satisfy a self-consistency constraint across noise levels, enabling high-quality generation with an arbitrary (and often very small) number of inference steps. While highly effective, these methods require costly retraining and their performance is closely tied to the specific base model and training configuration.

The second group consists of *training-free* methods, which accelerate sampling by improving the numerical integration of the underlying generative ODEs without modifying model parameters. Early work such as DDIM (Song, Meng, and Ermon 2021)—interpretable as a first-order discretization of the deterministic ODE formulation of diffusion models—showed that this approach outperforms discretizing the stochastic SDE, especially in low-NFE regimes. More recent high-order solvers like DPM-Solver (Lu et al. 2022, 2023) further improve stability and accuracy, though their performance still degrades when the step count becomes extremely small, leaving room for improvement.

Meanwhile, the *Flow Matching* (FM) framework (Lipman et al. 2023) has emerged as a compelling paradigm for generative modeling. FM learns a continuous-time velocity field that deterministically transports a simple base distribution to the target data distribution, and often achieves competitive or superior results with fewer discretization steps (Liu, Gong, and Liu 2022; Esser et al. 2024). For Gaussian base distributions, FM can be interpreted as a special case of a diffusion model under a particular noise schedule and a deterministic sampling path (Patel, DeLoye, and Mathias 2024; Gao et al. 2025). However, this correspondence holds only for a specific schedule; a general diffusion model with an arbitrary noise schedule cannot be directly expressed in FM form.

Motivated by this connection, we introduce the *Flow Path Sampler* (FloPS), which reparameterizes the sampling trajectory of an arbitrary pre-trained diffusion model into a path

consistent with FM dynamics. This transformation retains the original score network while inheriting the temporal regularity of FM trajectories, enabling more stable and efficient ODE-based integration—entirely without retraining.

Building on this foundation, we further propose the *Adaptive Flow Path Sampler* (A-FloPS), which incorporates an adaptive decomposition of the velocity field to minimize the temporal variability of the residual component. This adaptivity allows high-order numerical solvers to operate more effectively, resulting in substantial improvements in both stability and sample quality, particularly in the few-step regime.

Our work makes the following contributions:

- **Trajectory-level acceleration via diffusion-to-flow reparameterization.** We derive an analytical mapping that transforms the sampling trajectory of *any* pre-trained diffusion model into an integration-friendly flow-matching (FM) form—without retraining or altering network parameters. This enables plug-and-play acceleration for a wide range of generative models.
- **Adaptive velocity decomposition for few-step generation.** We introduce an inference-time adaptive mechanism that factorizes the FM velocity field into a dominant linear drift and a smooth residual, actively suppressing temporal variability. This restores the numerical advantages of high-order integration in the low-NFE regime.
- **Unified framework for diffusion and native FM models.** Beyond reparameterized diffusion models, our adaptive mechanism directly benefits native FM generators, demonstrating its generality across deterministic generative ODEs.
- **State-of-the-art few-step performance.** Extensive experiments on conditional image generation and text-to-image synthesis show that A-FloPS consistently outperforms strong training-free baselines, delivering sharper details, stronger semantic alignment, and substantially lower FID with as few as 5 steps.

## Background

### Diffusion Model

Modern diffusion models, originally introduced as discrete-time Markov chains, are now predominantly formulated from the perspective of stochastic differential equations (SDEs) (Song et al. 2021). This continuous-time view facilitates theoretical analysis and allows the use of advanced numerical solvers. In this framework, the forward diffusion process gradually perturbs data into noise according to:

$$d\mathbf{x}_\tau = -\alpha(\tau)\mathbf{x}_\tau d\tau + \sqrt{\beta(\tau)}d\mathbf{w}, \quad \mathbf{x}_0 \sim p_0, \quad (1)$$

where  $\mathbf{w}$  denotes the Wiener process,  $p_0$  is the target data distribution,  $\alpha(\tau)$  is the drift coefficient, and  $\beta(\tau)$  is the diffusion coefficient.

The marginal distribution of  $\mathbf{x}_\tau$  conditioned on  $\mathbf{x}_0$  is:

$$\mathbf{x}_\tau | \mathbf{x}_0 \sim \mathcal{N}(\bar{\alpha}_\tau \mathbf{x}_0, \sigma_\tau^2 \mathbf{I}), \quad (2)$$

where

$$\sigma_\tau^2 = \int_0^\tau \exp\left(-2 \int_s^\tau \alpha(r) dr\right) \beta(s) ds$$

and

$$\bar{\alpha}_\tau = \exp\left(-\int_0^\tau \alpha(s) ds\right).$$

At the terminal time  $\tau = T$ ,  $\sigma_T^2$  is large and  $\bar{\alpha}_T$  is small, meaning the state approaches an isotropic Gaussian dominated by noise.

The generative process is obtained by reversing the SDE:

$$d\tilde{\mathbf{x}}_\tau = [-\alpha(\tau)\tilde{\mathbf{x}}_\tau - \beta(\tau)\nabla_{\tilde{\mathbf{x}}_\tau} \log p_\tau(\tilde{\mathbf{x}}_\tau)] d\tau + \sqrt{\beta(\tau)}d\tilde{\mathbf{w}}, \quad (3)$$

where  $\nabla_{\tilde{\mathbf{x}}_\tau} \log p_\tau$  is the *score function* and  $\tilde{\mathbf{w}}$  is the reverse-time Wiener process. Alternatively, one can remove the stochastic term to obtain the reverse-time ODE:

$$d\bar{\mathbf{x}}_\tau = \left[-\alpha(\tau)\bar{\mathbf{x}}_\tau - \frac{\beta(\tau)}{2}\nabla_{\bar{\mathbf{x}}_\tau} \log p_\tau(\bar{\mathbf{x}}_\tau)\right] d\tau, \quad (4)$$

which is preferred in practice due to improved numerical stability and efficiency.

In practice, the exact score function is intractable. Since  $p_T$  is approximately Gaussian, a neural network  $s_\theta(\mathbf{x}_\tau, \tau)$  is trained via denoising score matching to approximate it, and the reverse process is initialized from  $\bar{\mathbf{x}}_T \sim \mathcal{N}(\mathbf{0}, \sigma_T^2 \mathbf{I})$ . Besides direct score prediction, alternative parameterizations—such as noise prediction and velocity prediction—are also widely used (Karras et al. 2022), and can be algebraically converted to the score function for sampling.

### Accelerated Sampling Methods

Due to their iterative nature, diffusion model samplers often require hundreds of neural evaluations, leading to high inference latency. Acceleration strategies fall into two broad categories: *training-required* and *training-free* approaches.

**Training-Required Acceleration.** These methods train specialized generative models to reproduce the denoising behavior of a high-fidelity model on a coarse time grid. Knowledge distillation transfers the generative process from a *teacher* to a *student* with fewer steps, as in Progressive Distillation (Salimans and Ho 2022), which halves the number of steps iteratively. Consistency models (Song et al. 2024; Sauer et al. 2023) enforce self-consistency across noise levels, enabling generation with very few steps. Hybrid designs, such as Truncated Diffusion Models (Bao et al. 2022), replace later denoising stages with fast feed-forward generators.

**Training-Free Acceleration.** These methods improve sampling efficiency without retraining by designing better numerical solvers for the reverse ODE. DDIM (Song, Meng, and Ermon 2021) reformulates sampling as a deterministic ODE and applies first-order integration. High-order solvers like DPM-Solver (Lu et al. 2022) and DPM-Solver++ (Lu et al. 2023) exploit the semi-linear structure for higher accuracy in few-step generation. UniPC (Zhao et al. 2023) further unifies predictor–corrector schemes for diverse noise schedules, reducing the quality gap to full-trajectory sampling.

While these approaches improve integration accuracy, they remain constrained by the inherent structure of the sampling trajectory. This motivates exploring *trajectory reparameterization*, which can potentially make the path itself more integration-friendly.

---

**Algorithm 1: Flow Path Sampler(FloPS)**

---

**Input:** the estimation of score function  $s_\theta(\cdot, \cdot)$ , the number of sampling steps  $N$ , and the noise scheduler of the diffusion model  $\{\sigma_\tau, \bar{\alpha}_t | \tau = 0, 1, 2, \dots, T\}$ .

**Output:** Your algorithm’s output

```
1: Sample  $\mathbf{x}_0 \sim \mathcal{N}(0, \mathbf{I})$ .
2:  $\Delta t \leftarrow 1/N$ 
3:  $t_n \leftarrow n\Delta t$ 
4: for  $n = 0$  to  $N - 1$  do
5:   if  $t_n < \frac{1}{1+\sigma_T/\bar{\alpha}_T}$  then
6:      $\mathbf{v}_{t_n} \leftarrow \frac{\sigma_\tau[\mathbf{x}_t + \sigma_T s_\theta((\bar{\alpha}_T + \sigma_T)\mathbf{x}_T, T)]}{\bar{\alpha}_\tau}$ .
7:   else
8:      $\tau \leftarrow \arg \min_s \left| t_n - \frac{1}{1+\sigma_s/\bar{\alpha}_s} \right|$ .
9:      $\mathbf{v}_{t_n} \leftarrow \frac{\sigma_\tau[\mathbf{x}_t + \sigma_\tau s_\theta((\bar{\alpha}_\tau + \sigma_\tau)\mathbf{x}_\tau, \tau)]}{\bar{\alpha}_\tau(1-t_n)}$ .
10:  end if
11:   $\mathbf{x}_{t_{n+1}} \leftarrow \mathbf{x}_{t_n} + \mathbf{v}_{t_n} \Delta t$ 
12: end for
13: return  $\mathbf{x}_1$ 
```

---

### Flow Matching Model

Flow matching (Lipman et al. 2023) formulates generative modeling as learning a *deterministic* continuous-time flow that transports a simple reference distribution  $q_0$  to a complex target distribution  $q_1$ . It defines a time-dependent velocity field  $\mathbf{v}(\mathbf{x}_t, t)$  evolving samples via:

$$\frac{d\mathbf{x}_t}{dt} = \mathbf{v}^*(\mathbf{x}_t, t), \quad (5)$$

with  $\mathbf{x}_0 \sim q_0$  and the terminal state  $\mathbf{x}_1$  following  $q_1$ .

A common example is Gaussian Optimal Transport, where  $q_0 = \mathcal{N}(\mathbf{0}, \mathbf{I})$  and  $q_1$  is the data distribution. The forward interpolation is:

$$\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\mathbf{x}_1,$$

where  $\mathbf{x}_0 \sim q_0$  and  $\mathbf{x}_1 \sim q_1$ . The optimal velocity field minimizes:

$$\min_{\mathbf{v}} \mathbb{E}_{t, \mathbf{x}_0, \mathbf{x}_1} \|\mathbf{v}(\mathbf{x}_t, t) - (\mathbf{x}_1 - \mathbf{x}_0)\|^2, \quad (6)$$

whose solution is the conditional expectation:

$$\mathbf{v}^*(\mathbf{x}_t, t) = \mathbb{E}[\mathbf{x}_1 - \mathbf{x}_0 | \mathbf{x}_t]. \quad (7)$$

This guarantees a globally consistent transport from  $q_0$  to  $q_1$ .

Notably, flow matching can be interpreted as a deterministic sampler variant of diffusion models under a specific noise schedule. This connection enables reparameterizing diffusion trajectories into FM-like paths for potentially more efficient sampling.

### Method

Our approach introduces a training-free acceleration framework for diffusion model sampling, composed of two key components: **Flow Path Sampler (FloPS)** and its adaptive variant **A-FloPS**. FloPS reparameterizes the sampling

trajectory of any pre-trained diffusion model into a flow-matching (FM) trajectory, enabling more efficient ODE integration without retraining. Building on this foundation, A-FloPS further incorporates an *adaptive decomposition* of the velocity field, which suppresses temporal variability in the residual term and thus restores the effectiveness of high-order integration under the FM framework. The following subsections detail each component in turn, from the motivation and formulation of FloPS to the design and implementation of our adaptive mechanism.

### Flow Path Sampler (FloPS)

Flow Matching (FM) models have been shown to consistently outperform standard diffusion models under the same number of discretization steps, primarily due to their smoother and more regular sampling trajectories induced by the noise schedule (Lipman et al. 2023). This observation naturally raises the question: *can we reconstruct the sampling trajectory of a diffusion model in an analogous manner, thereby improving both sampling efficiency and generation quality?*

We address this question by deriving an explicit mapping from the score function of a diffusion model to the velocity field of an FM model targeting the same data distribution. This mapping is formalized in Theorem 1, with a detailed proof provided in the Appendix. Intuitively, it enables the conversion of any pre-trained diffusion model into a valid FM trajectory *without* retraining or architectural modifications. From a practical standpoint, this transformation bridges the algorithmic gap between two previously distinct paradigms: score-based stochastic generative modeling and deterministic flow-based modeling. It reveals that the efficiency of FM sampling is not an inherent property of its architecture, but rather a consequence of a more integration-friendly trajectory structure, which can in principle be induced from any diffusion model.

**Theorem 1** (Diffusion-to-Flow Transformation). *If a diffusion model and a flow-matching model target the same distribution, i.e.,  $q_1$  in Eq. (6) equals  $p_0$  in Eq. (1), their velocity field and score function satisfy:*

$$(1-t)\mathbf{v}^*(\mathbf{x}_t, t) = \frac{\sigma_\tau}{\bar{\alpha}_\tau} [\mathbf{x}_t + \sigma_\tau \nabla \log p_\tau((\bar{\alpha}_\tau + \sigma_\tau)\mathbf{x}_t)],$$

where  $t = \frac{1}{1+\sigma_\tau/\bar{\alpha}_\tau}$  and this transformation is bijective since  $\sigma_\tau/\bar{\alpha}_\tau$  increases strictly with  $\tau$ , ensuring a one-to-one correspondence between  $\tau$  and  $t$ .

Building on this theoretical insight, we propose the Flow Path Sampler (FloPS), a training-free, plug-and-play framework that reuses the pre-trained score network of a diffusion model to construct its FM trajectory. Unlike conventional FM training, FloPS does not require a separately learned velocity field and is applicable to models with *arbitrary* noise schedules or prediction parameterizations. This universality is crucial for real-world adoption: FloPS can be seamlessly integrated into existing pipelines, including large text-to-image models and high-resolution conditional generators, without access to their training data or checkpoints beyond the original network weights.

---

**Algorithm 2: Adaptive Flow Path Sampler(A-FloPS)**


---

**Input:** the estimation of score function  $s_\theta(\cdot, \cdot)$ , the number of sampling steps  $N$ , and the noise scheduler of the diffusion model  $\{\sigma_\tau, \bar{\alpha}_t | \tau = 0, 1, 2, \dots, T\}$ .

**Output:** Your algorithm’s output

```

1: Sample  $\mathbf{x}_0 \sim \mathcal{N}(0, \mathbf{I})$ .
2:  $\Delta t \leftarrow 1/N$ 
3:  $t_n = n\Delta t$ 
4: for  $n = 0$  to  $N - 1$  do
5:   if  $t_n < \frac{1}{1 + \sigma_T / \bar{\alpha}_T}$  then
6:      $\mathbf{v}_{t_n} \leftarrow \frac{\sigma_\tau [\mathbf{x}_t + \sigma_T s_\theta((\bar{\alpha}_T + \sigma_T)\mathbf{x}_T, T)]}{\bar{\alpha}_\tau}$ .
7:   else
8:      $\tau \leftarrow \arg \min_s \left| t_n - \frac{1}{1 + \sigma_s / \bar{\alpha}_s} \right|$ .
9:      $\mathbf{v}_{t_n} \leftarrow \frac{\sigma_\tau [\mathbf{x}_t + \sigma_\tau s_\theta((\bar{\alpha}_\tau + \sigma_\tau)\mathbf{x}_\tau, \tau)]}{\bar{\alpha}_\tau (1 - t_n)}$ .
10:  end if
11:  if  $n=0$  then
12:     $\mathbf{x}_{t_{n+1}} \leftarrow \mathbf{x}_{t_n} + \mathbf{v}_{t_n} \Delta t$ 
13:  else
14:    Calculate  $\lambda^{(n)}$  as (12)
15:    Calculate  $a, b$  as in (15)
16:    Calculate  $\mathbf{x}_{t_{n+1}}$  as in (14) using the finite-difference form of the derivative in (16)
17:  end if
18: end for
19: return  $\mathbf{x}_1$ 

```

---

In practice, many diffusion schedulers maintain a non-zero signal strength at the start of sampling ( $\bar{\alpha}_T \neq 0$ ). Consequently, the mapping in Theorem 1 is exact only after an initial *start time*  $t_{\min}$ , given by

$$t_{\min} = \frac{1}{1 + \sigma_T / \bar{\alpha}_T}.$$

For earlier times  $t < t_{\min}$ , we approximate the velocity using the value at  $t_{\min}$ , i.e.,  $\mathbf{v}(\mathbf{x}_t, t_{\min})$ . The complete FloPS algorithm is summarized in Algorithm 1, where we present the method in its score-prediction form; equivalent formulations for noise- or data-prediction are straightforward.

### Adaptive Mechanism

High-order samplers such as DPM-Solver++ (Lu et al. 2022) achieve acceleration primarily through two mechanisms: (i) employing a high-order expansion of the estimator over the integration interval, thereby improving numerical accuracy; and (ii) leveraging the approximate *temporal invariance* of the estimator  $\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_\tau]$ , which enables accurate extrapolation from future time steps. The second factor is particularly important, as it allows the solver to exploit smooth temporal behavior for better high-order integration.

However, in flow matching, the velocity field  $\frac{d\mathbf{x}_t}{dt}$  exhibits only *weak temporal variation* along the sampling trajectory (e.g., exactly constant for a Dirac delta target), which renders factor (ii) largely ineffective. As a result, even a simple Euler integrator can already achieve excellent performance, and the benefits of applying high-order samplers

such as DPMSolver diminish considerably. This means that naively applying high-order methods to FM-like trajectories can even waste computation, since the additional polynomial fitting captures almost no extra curvature information, which has been found in practice (Tan et al. 2025).

To overcome this limitation, we propose an *adaptive decomposition* that factorizes the velocity field into a linear term and a higher-order residual term in an adaptive manner. By suppressing the temporal variation of the residual, our method restores the effectiveness of high-order integration within the flow-matching framework.

**Adaptive coefficient estimation.** We rewrite the flow-matching ODE (5) as:

$$\frac{d\mathbf{x}_t}{dt} = \lambda_t \mathbf{x}_t + [\mathbf{v}(\mathbf{x}_t, t) - \lambda_t \mathbf{x}_t], \quad (8)$$

where  $\lambda_t$  is a time-dependent coefficient chosen to make the residual  $\mathbf{v}(\mathbf{x}_t, t) - \lambda_t \mathbf{x}_t$  as temporally smooth as possible. Unlike the fixed or pre-trained decompositions in the DPM-Solver family,  $\lambda_t$  is *adaptive*—it is re-estimated at every integration step to track local variations in the velocity field. The intuition is simple: if the remaining term varies slowly, high-order scheme is allowed to exploit temporal smoothness effectively.

For stability and efficiency, we treat  $\lambda_t$  as piecewise constant within each integration interval:

$$\lambda_t = \lambda^{(n)}, \quad t \in [t_n, t_{n+1}), \quad (9)$$

where  $0 = t_0 < t_1 < \dots < t_N = 1$ . Ideally,  $\lambda^{(n)}$  minimizes the instantaneous temporal change of the residual:

$$\lambda_{\text{ideal}}^{(n)} = \arg \min_\lambda \left\| \left. \frac{d}{dt} \mathbf{h}(\mathbf{x}_t, t; \lambda) \right| \right\|_{t=t_n}^2, \quad (10)$$

with  $\mathbf{h}(\mathbf{x}_t, t; \lambda) := \mathbf{v}(\mathbf{x}_t, t) - \lambda \mathbf{x}_t$ .

Direct derivative computation is costly, so we approximate it by a finite difference between consecutive steps:

$$\lambda^{(n)} = \arg \min_\lambda \left\| \mathbf{h}(\mathbf{x}_{t_n}, t_n; \lambda) - \mathbf{h}(\mathbf{x}_{t_{n-1}}, t_{n-1}; \lambda) \right\|_2^2. \quad (11)$$

This quadratic problem has the closed-form solution:

$$\lambda^{(n)} = \frac{\langle \Delta \mathbf{v}, \Delta \mathbf{x} \rangle}{\|\Delta \mathbf{x}\|_2^2}, \quad (12)$$

where  $\Delta \mathbf{v} := \mathbf{v}(\mathbf{x}_{t_n}, t_n) - \mathbf{v}(\mathbf{x}_{t_{n-1}}, t_{n-1})$  and  $\Delta \mathbf{x} := \mathbf{x}_{t_n} - \mathbf{x}_{t_{n-1}}$ . This choice minimizes local temporal variation in the residual, improving high-order integrator accuracy.

**High-order integration.** To achieve a better trade-off between discretization error and the number of function evaluations (NFEs), we employ a high-order integration scheme for the flow ODE. With (9), the flow ODE can be written as:

$$\begin{aligned} \mathbf{x}_{t_{n+1}} &= e^{\lambda^{(n)} \Delta t} \mathbf{x}_{t_n} \\ &+ \int_{t_n}^{t_{n+1}} e^{\lambda^{(n)}(t_{n+1}-\tau)} \mathbf{h}(\mathbf{x}_\tau, \tau; \lambda^{(n)}) d\tau, \end{aligned} \quad (13)$$




Figure 1: Qualitative comparison of generated images for the “golden retriever” class using different diffusion samplers across varying inference steps. Each row corresponds to a specific sampler (from top to bottom: DDIM, DPM-Solver++, UniPC, FloPS, and A-FloPS), while each column shows results at a different number of function evaluations (NFE = 5, 6, 7, 8, 9, 25). All methods are initialized with the same random seed to ensure a fair comparison. Our proposed A-FloPS generates sharper textures and more coherent object structures, achieving superior visual fidelity with substantially fewer sampling steps.

where  $\Delta t = t_{n+1} - t_n$ . Approximating the integral by a second-order Taylor expansion at  $t_n$  yields:

$$\mathbf{x}_{t_{n+1}} \approx e^{\lambda^{(n)}\Delta t} \mathbf{x}_{t_n} + a \mathbf{h}(\mathbf{x}_{t_n}, t_n; \lambda^{(n)}) + b \left. \frac{d\mathbf{h}(\mathbf{x}_t, t; \lambda^{(n)})}{dt} \right|_{t=t_n}, \quad (14)$$

where

$$a = \frac{1 - e^{-\lambda^{(n)}\Delta t}}{\lambda^{(n)}}, \quad b = \frac{1 - (1 + \lambda^{(n)}\Delta t)e^{-\lambda^{(n)}\Delta t}}{(\lambda^{(n)})^2}. \quad (15)$$

In practice, instead of computing the time derivative in (14) explicitly, we use its backward finite-difference approximation:

$$\left. \frac{d\mathbf{h}}{dt} \right|_{t=t_n} \approx \frac{\mathbf{h}(\mathbf{x}_{t_n}, t_n; \lambda^{(n)}) - \mathbf{h}(\mathbf{x}_{t_{n-1}}, t_{n-1}; \lambda^{(n)})}{\Delta t}. \quad (16)$$

Hence, the update in (14) is implemented using the finite-difference estimate (16) in place of the time derivative.

This adaptive decomposition is applied entirely at inference time, requiring *no* retraining or architectural changes. By continually re-estimating  $\lambda_t$ , we suppress residual temporal variation and restore the accuracy gains of high-order integration—particularly in the low-step regime where fixed-coefficient methods often underperform. We emphasize that this adaptivity is local in time and model-agnostic: it can be applied not only to FloPS but also to any FM-style sampler, and even to certain deterministic ODE solvers outside the generative modeling domain.

## Experiments

We evaluate A-FloPS across diverse generative settings to assess both sampling efficiency and generation quality. Our study addresses three main questions: (1) How does A-FloPS compare to existing training-free samplers under varying NFEs? (2) What is the individual contribution of

the proposed flow-based reparameterization? (3) How effective is the adaptive mechanism in both diffusion and flow-matching models?

## Experimental Setup

We consider two representative generative modeling scenarios:

- Conditional image generation: 50K ImageNet (Deng et al. 2009)-256×256 samples generated by DiT (Peebles and Xie 2023).
- Text-to-image generation: 5K COCO-prompt (Lin et al. 2014) samples generated by Stable Diffusion v3.5 (d=38).

**Metrics.** For conditional generation, we report Fréchet Inception Distance (FID), Inception Score (IS), and improved Precision/Recall (Kynkäänniemi et al. 2019). For text-to-image generation, we use CLIP score (Radford et al. 2021), ImageReward (IR) (Xu et al. 2023), and an LLM-based evaluation for text–image alignment and perceptual quality.

**Protocol.** All experiments vary the number of function evaluations (NFE=5–10) with fixed random seeds for fair comparison, focusing on the low-NFE regime most relevant to real-time deployment. Experiments are conducted on a workstation equipped with 8×NVIDIA A100 GPUs (40GB each). The code is available at <https://github.com/jinc7461/AFloPS>.

## Main Results on DiT Backbone

We compare A-FloPS with representative training-free samplers: DDIM (Song, Meng, and Ermon 2021), DPM-Solver++ (Lu et al. 2022), and UniPC (Zhao et al. 2023). Following (Peebles and Xie 2023), we use a CFG scale of 1.5. For stability,  $\lambda^{(n)}$  is constrained to  $[-1, 1]$ , and  $a, b$  in (15) retain terms only up to second order in  $\Delta t$ .

Table 1 shows that A-FloPS consistently outperforms all baselines. At NFE = 5, it achieves a notably lower FID than the next-best method, confirming its efficiency in the few-step regime. The advantage persists as NFE increases, demonstrating the robustness of our trajectory reparameterization and adaptive integration.

Qualitative results in Figure 1 further support these findings: competing methods tend to yield overly smooth or blurry edges at NFE = 5, while A-FloPS produces sharper textures and more coherent structures. These results align with our objective of improving temporal regularity in the sampling trajectory—particularly critical for few-step generation. Additional examples are provided in the appendix.

## Ablation Study

**Flow Path Reparameterization without Adaptive Mechanism.** To isolate the contribution of flow-path reparameterization, we remove the adaptive mechanism while retaining the flow-transformed velocity field. As shown in Table 2, this non-adaptive variant (FloPS) almost consistently surpasses DDIM across all NFEs, demonstrating that reparameterizing the sampling trajectory alone yields a more sample-efficient process and achieves better alignment with



(a) A-Euler (Ours)



(b) Euler (Baseline)

Figure 2: Qualitative comparison between A-Euler and baseline Euler on SDv3.5 (d=38) at NFE = 5. Prompt: “A cat is sitting on top of a bed in front of a living room”. A-Euler produces sharper textures and more coherent structures.

the target distribution. While it falls short of the perceptual quality attained by the full A-FloPS, it still delivers markedly sharper and more coherent results than DDIM, particularly in the challenging low-NFE regime. This finding suggests that re-normalizing the sampling path of existing diffusion models offers a practical and virtually cost-free improvement.

**Effectiveness of Adaptivity in Flow-Based Models.** Building on the ablation results above, where A-FloPS consistently outperforms its non-adaptive counterpart (FloPS), it is natural to ask: while FloPS reparameterizes a diffusion model into a flow-matching (FM) model and the adaptive mechanism is tailored to FM dynamics, can this mechanism also benefit *native* FM models that are not obtained via repara-

Step	DDIM				UniPC				DPM-Solver				A-FloPS (proposed)			
	FID	IS	Prec.	Rec.	FID	IS	Prec.	Rec.	FID	IS	Prec.	Rec.	FID	IS	Prec.	Rec.
5	38.124	78.897	0.447	0.438	19.165	124.850	0.598	0.431	22.267	115.066	0.571	0.416	<b>6.977</b>	<b>195.176</b>	<b>0.751</b>	<b>0.465</b>
6	22.449	119.865	0.570	0.456	10.165	172.094	0.699	0.472	11.852	162.066	0.681	0.454	<b>4.947</b>	<b>224.613</b>	<b>0.790</b>	<b>0.486</b>
7	14.650	151.406	0.646	0.463	6.698	206.204	0.752	<b>0.493</b>	7.555	196.617	0.739	0.486	<b>4.196</b>	<b>239.762</b>	<b>0.809</b>	0.492
8	10.325	175.733	0.700	0.481	5.320	225.468	0.780	<b>0.504</b>	5.711	216.763	0.772	0.502	<b>3.753</b>	<b>248.932</b>	<b>0.817</b>	0.500
9	7.669	196.445	0.735	0.489	4.828	237.609	0.794	0.514	4.947	231.484	0.787	0.513	<b>3.647</b>	<b>253.016</b>	<b>0.821</b>	<b>0.516</b>
10	6.234	209.926	0.759	0.506	4.638	243.356	0.803	0.516	4.615	239.502	0.799	0.517	<b>3.440</b>	<b>258.486</b>	<b>0.826</b>	<b>0.521</b>

Table 1: Quantitative comparison of different samplers under varying numbers of steps. FID ( $\downarrow$ ), IS ( $\uparrow$ ), Precision ( $\uparrow$ ), and Recall ( $\uparrow$ ) are reported.

Step	DDIM				FloPS (proposed)			
	FID	IS	Prec.	Rec.	FID	IS	Prec.	Rec.
5	38.12	78.90	0.4465	0.4381	14.9 (−23)	151.44 (+72.5)	0.6425 (+0.2)	0.458 (+0.02)
6	22.45	119.87	0.5702	0.456	9.721 (−13)	181.12 (+61.3)	0.7057 (+0.14)	0.4814 (+0.025)
7	14.65	151.41	0.6465	0.4631	7.005 (−7.6)	204.47 (+53.1)	0.7469 (+0.1)	0.4869 (+0.024)
8	10.33	175.73	0.6996	0.4813	5.587 (−4.7)	219.57 (+43.8)	0.7713 (+0.072)	0.4964 (+0.015)
9	7.669	196.45	0.7351	0.4889	4.623 (−3)	231.76 (+35.3)	0.7904 (+0.055)	0.5042 (+0.015)
10	6.234	209.93	0.7591	0.5062	4.11 (−2.1)	241.02 (+31.1)	0.8026 (+0.044)	0.5052 (−0.001)

Table 2: Comparison of DDIM and FloPS across steps. FloPS values are annotated with color-coded deltas relative to DDIM.

NFE	Euler			A-Euler (Ours)		
	CLIP $\uparrow$	IR $\uparrow$	GPT $\uparrow$	CLIP $\uparrow$	IR $\uparrow$	GPT $\uparrow$
5	0.316	0.300	0.172	<b>0.319</b>	<b>0.579</b>	<b>0.828</b>
7	<b>0.319</b>	0.629	0.234	0.318	<b>0.785</b>	<b>0.766</b>
9	<b>0.319</b>	0.755	0.315	0.317	<b>0.843</b>	<b>0.685</b>

Table 3: Comparison of Euler and A-Euler on SDv3.5 under different NFEs (CFG = 2.0). Higher CLIP, IR, and GPT indicate better performance.

reparameterization?

To answer this, we integrate the adaptive mechanism into a representative FM-based model, Stable Diffusion v3.5 (d=38). As shown in Table 3, it yields consistent gains across most metrics, with the largest improvements at low NFEs. Although CLIP scores for A-Euler are slightly lower than Euler at  $N = 7, 9$ , this is likely due to CLIP’s known saturation and stylistic bias (Yarom et al. 2023). In contrast, IR and GPT-based metrics—better aligned with human judgment—show substantial improvements. The GPT-based evaluation compares two images from the same prompt, selecting the one with better text alignment and perceptual quality; details are in the appendix. Figure 2 further shows that A-Euler produces sharper textures and more coherent structures than baseline Euler at NFE = 5.

These improvements on a native FM model suggest that the adaptive mechanism is not tied to the specifics of our diffusion-to-flow reparameterization. In fact, the mechanism is agnostic to how the FM velocity field is obtained—whether learned directly, distilled from a diffusion model, or analytically derived. It can thus be applied to a broad range of deterministic generative ODEs, such as rec-

tified flows (Liu, Gong, and Liu 2022) and hybrid diffusion–flow systems. While our experiments focus on image generation, the method naturally extends to other modalities such as audio or video synthesis.

Our experiments lead to three main observations. First, flow-based reparameterization substantially improves sampling efficiency by producing trajectories that are inherently more amenable to numerical integration. Second, the adaptive mechanism provides additional accuracy gains, particularly in the low-NFE regime. Finally, combining these two components allows A-FloPS to achieve both high sampling speed and strong generation quality, all without requiring model retraining.

## Conclusion

We introduced **A-FloPS**, a training-free framework that accelerates diffusion sampling by reparameterizing arbitrary sampling paths into flow-matching (FM) trajectories and augmenting them with an adaptive velocity decomposition. This addresses two core limitations of prior training-free methods: (i) inefficient integration of conventional diffusion trajectories, and (ii) reduced benefits of high-order solvers under weakly time-varying velocity fields.

The reparameterization alone yields integration-friendly paths, while the adaptive mechanism restores the accuracy advantages of high-order integration in the few-step regime. Together, they deliver superior generation quality and semantic fidelity with far fewer function evaluations, across both diffusion and native FM models.

Experiments on ImageNet and COCO show consistent state-of-the-art performance among training-free samplers, especially at extremely low NFEs. Future work will extend A-FloPS to multi-modal synthesis, real-time generation, and richer temporal models for the adaptive coefficient.

## Ethics Statement

This work focuses on improving the sampling efficiency of pretrained diffusion models. All experiments are conducted on publicly available datasets with appropriate licenses (ImageNet, COCO). No personally identifiable information or sensitive data is used. The proposed method does not introduce new ethical risks beyond those already associated with generative models, such as potential misuse for generating misleading content. We emphasize that our method should be applied responsibly and strictly for research and non-harmful purposes. No human subjects or animals were involved, and no IRB approval was required.

## Acknowledgements

This work was supported by NSAF (Grant No. U2230201). All authors are affiliated with the Department of Electronic Engineering, Tsinghua University. In addition, Y. Gu is also affiliated with the Tsinghua Shenzhen International Graduate School and the Beijing National Research Center for Information Science and Technology at Tsinghua University.

## References

- Bao, Y.; Chen, Y.; Li, C.; and Zhu, J. 2022. Truncated Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*.
- Gao, R.; Hoogeboom, E.; Heek, J.; De Bortoli, V.; Murphy, K. P.; and Salimans, T. 2025. Diffusion Meets Flow Matching: Two Sides of the Same Coin. <https://diffusionflow.github.io/>.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *NeurIPS*.
- Karras, T.; Aittala, M.; Aila, T.; and Laine, S. 2022. Elucidating the Design Space of Diffusion-Based Generative Models. In *Proc. NeurIPS*.
- Kynkäänniemi, T.; Karras, T.; Laine, S.; Lehtinen, J.; and Aila, T. 2019. Improved precision and recall metric for assessing generative models. In *Advances in Neural Information Processing Systems*, volume 32.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Lipman, Y.; Chen, R. T. Q.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2023. Flow Matching for Generative Modeling. In *International Conference on Learning Representations (ICLR)*.
- Liu, X.; Gong, C.; and Liu, Q. 2022. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2022. DPM-Solver: A Fast ODE Solver for Diffusion Probabilistic Model Sampling in Around 10 Steps. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, 5775–5787.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2023. DPM-Solver++: Fast Solver for Guided Sampling of Diffusion Models. In *International Conference on Learning Representations (ICLR)*.
- Patel, Z.; DeLoye, J.; and Mathias, L. 2024. Exploring Diffusion and Flow Matching Under Generator Matching. *arXiv preprint arXiv:2412.11024*.
- Peebles, W.; and Xie, S. 2023. Scalable Diffusion Models with Transformers. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 16709–16721.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning*, 8748–8763.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2): 3.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *CVPR*.
- Salimans, T.; and Ho, J. 2022. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations (ICLR)*.
- Sauer, A.; Karras, T.; Laine, S.; and Geiger, A. 2023. Latent Consistency Models: Synthesizing High-Resolution Images with Few-Step Inference. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Schusterbauer, J.; Gui, M.; Fundel, F.; and Ommer, B. 2025. Diff2Flow: Training Flow Matching Models via Diffusion Model Alignment. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 28347–28357.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising Diffusion Implicit Models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, 6115–6126.
- Song, Y.; Dhariwal, P.; Chen, M.; and Sutskever, I. 2024. Consistency Models. In *International Conference on Learning Representations (ICLR) 2024*.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-based generative modeling through stochastic differential equations. In *ICLR*.
- Tan, Z.; Wang, W.; Bertozzi, A. L.; and Ryu, E. K. 2025. STORK: Improving the Fidelity of Mid-NFE Sampling for Diffusion and Flow Matching Models. *arXiv preprint arXiv:2505.24210*. May 30 2025.

Xu, J.; Liu, X.; Wu, Y.; Tong, Y.; Li, Q.; Ding, M.; Tang, J.; and Dong, Y. 2023. ImageReward: Learning and Evaluating Human Preferences for Text-to-Image Generation. In *Advances in Neural Information Processing Systems*.

Yarom, M.; Bitton, Y.; Changpinyo, S.; Aharoni, R.; Herzig, J.; Lang, O.; Ofek, E.; and Szpektor, I. 2023. What You See Is What You Read? Improving Text-Image Alignment Evaluation. In *Advances in Neural Information Processing Systems*. Poster, NeurIPS 2023.

Zhao, W.; Bai, L.; Rao, Y.; Zhou, J.; and Lu, J. 2023. UniPC: A Unified Predictor–Corrector Framework for Fast Sampling of Diffusion Models. In *NeurIPS 2023*.