

Few-Step Diffusion Sampling Through Instance-Aware Discretizations

Liangyu Yuan^{1,2*} Ruoyu Wang^{1*} Tong Zhao^{1,3*} Dingwen Fu¹
Mingkun Lei¹ Beier Zhu⁴ Chi Zhang^{1†}
¹Westlake University ²Tongji University
³Zhejiang University ⁴Nanyang Technological University

Abstract

Diffusion and flow matching models generate high-fidelity data by simulating paths defined by Ordinary or Stochastic Differential Equations (ODEs/SDEs), starting from a tractable prior distribution. The probability flow ODE formulation enables the use of advanced numerical solvers to accelerate sampling. Orthogonal yet vital to solver design is the discretization strategy. While early approaches employed handcrafted heuristics and recent methods adopt optimization-based techniques, most existing strategies enforce a globally shared timestep schedule across all samples. This uniform treatment fails to account for instance-specific complexity in the generative process, potentially limiting performance. Motivated by controlled experiments on synthetic data, which reveals the suboptimality of global schedules under instance-specific dynamics, we propose an instance-aware discretization framework. Our method learns to adapt timestep allocations based on input-dependent priors, extending gradient-based discretization search to the conditional generative setting. Empirical results across diverse settings, including synthetic data, pixel-space diffusion, latent-space images and video flow matching models, demonstrate that our method consistently improves generation quality with marginal tuning cost compared to training and negligible inference overhead.

1. Introduction

Diffusion Probabilistic Models (DPMs) [1, 2] and adjacent flow-matching models [3–5] generate high-fidelity data by simulating trajectories defined by ODEs/SDEs, starting from a simple prior distribution (typically isotropic Gaussian). This iterative refinement process underpins their strong generative capabilities across diverse modalities [6–10]. However, the generative power comes at a price: the tedious sampling time required for high-quality generation.

Acceleration methods for diffusion models can be divided into two main groups, model distillation [5, 11, 12] and training free acceleration [13–17]. Model distillation enables extreme few-step generation but often leads to distillation cost comparable to training. Conversely, training-free methods avoid the heavy tuning cost with a trade-off on more steps. Among them, solver-based methods stand out as an architecture-agnostic choice, leveraging numerical ODEs/SDEs techniques for higher order, multistep sampling, providing more portable acceleration for pre-trained models.

An essential aspect of solver design is the time discretization strategy. Initial approaches frequently relied on empirically derived heuristics, e.g. uniform [1] or logSNR [15]. However, these heuristics were identified as suboptimal for maximizing efficiency. Subsequently, research efforts have increasingly focused on optimization-based techniques to search for better discretization strategies [18–20]. Despite their improved efficiency, these optimization-driven approaches [18, 21–23] share a critical limitation: they enforce a single, globally optimized timestep schedule for all starting priors. This design may neglect the intrinsic variability in data complexity across samples. In practice, different inputs can give rise to distinct sampling trajectories [24, 25], each potentially benefiting from a different discretization.

To investigate this limitation, we first conduct a quantitative analysis using toy datasets (Fig. 2), revealing discernible performance gaps between globally optimized and instance-adaptive discretization schedules. These observations highlight a critical limitation of the current discretization strategy and motivate the development of adaptive discretization strategies that dynamically allocate timesteps based on the characteristics of each input. Building on this insight, we propose an effective method that generalizes previous gradient-based discretization search by taking the prior conditioning as input to produce instance-aware discretizations, as illustrated in Fig. 1 (right). To ensure scalability from synthetic analysis to high-dimensional image synthesis, we further introduce adaptation to handle conditional guidance and generalize the framework for alleviating mismatch issues commonly encountered in diffusion models literature.

* Equal contribution. † Corresponding author. This work was done during Liangyu Yuan’s visit at WestLake University in 2025.

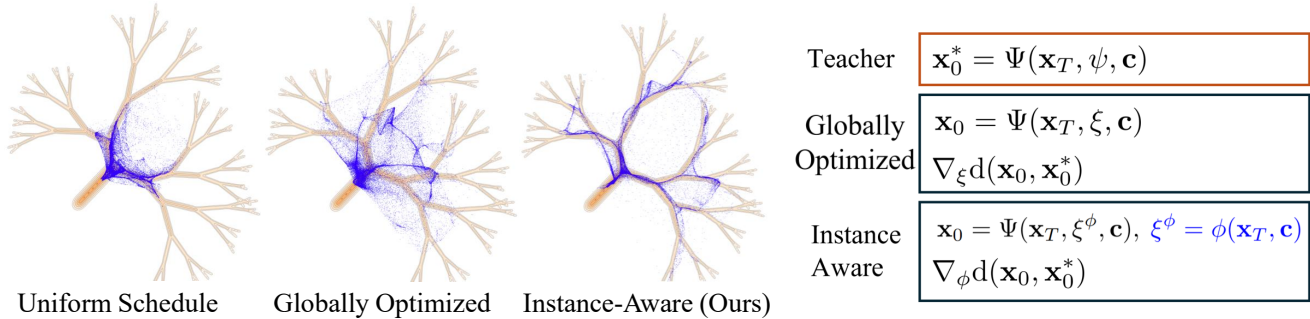


Figure 1. Our effective instance-aware discretization improves sampling quality, by generating a tailored discretization ξ^{ϕ} for each initial noise \mathbf{x}_T and condition \mathbf{c} , outperforming heuristic and globally optimized schedules. Orange contour represents the ground truth data distribution, blue dots represent the generated samples across different discretizations. ($\Psi(\cdot, \cdot, \cdot)$ represents the ODE path.)

We name our algorithm *INDIS* (*Instance-Specific Discretization*).

Empirical results across diverse settings, including synthetic datasets, pixel space [26], and latent space diffusion [27] and flow matching models on images [6] and videos [28], validate the effectiveness of our approach. Moreover, our approach is a lightweight solver acceleration method with negligible tuning overhead versus training or distilling the base model, and marginal additional sampling cost. Our contribution can be summarized as follows:

- We identify the limitations of global timestep discretization through synthetic experiments with quantitative analysis and propose an effective instance-aware discretization paradigm.
- We scale the paradigm of instance-aware discretization to high-dimensional image synthesis, by incorporating adaptations to manage conditional guidance and generalizing the framework to mitigate the exposure bias problem.
- Extensive experiments across diverse datasets and model types, including pixel-space diffusion, latent-space images and video flow matching models, validate the effectiveness of our approach.

2. Related work

Dedicated solvers for diffusion ODEs. Building upon the probability flow ODE formulation, significant research has aimed to accelerate diffusion sampling. DDIM [29] pioneered this by using a non-Markovian process to reduce DDPM [1] steps from thousands to fewer than a hundred. Subsequently, DPM-Solver [15] and DPM-Solver++ [30] realize the semi-linear form of diffusion ODE and use knowledge from exponential integrators to improve sampling quality. PNDM [31] and iPNDM [13] incorporated linear multistep methods for sampling efficiency. UniPC [14] further proposed a unified predictor-corrector framework to minimize local truncation errors. RX [32] uses Richardson Extrapolation to function as a plugin to improve sample quality

across multistep methods.

Diffusion solver fine-tuning. Recognizing the potential limitations of applying fixed solver parameters derived from general numerical analysis, more recent research [19, 33–38] has explored incorporating domain-specific information to enhance sampler performance. AMED [33] proposes to align the approximated mean value from low dimensionality across steps through tuning the intermediary time. Bespoke solver class (bespoke solver [34], bespoke non-stationary solver [35]) in flow matching models proposes to learn the general form of the solver parameters given global/local supervision.

Optimizing diffusion ODE timestep discretization. Apart from dedicated solver design, timesteps discretization fine-tuning has recently garnered significant research interest. Various approaches [18, 21–23] have been proposed, aiming to improve upon sub-optimal manually-crafted schedules. Both DMN [22] and AYS [23] formulate timestep selection as an optimization problem, solved using constrained trust-region methods and Monte Carlo sampling techniques, respectively. GITS [21] leverages assumptions about the geometric regularity of the sampling trajectory, modeling discretization as a shortest path problem that minimizes accumulated local truncation errors. LD3 [18] introduces a lightweight framework to explicitly learn the optimal discretization by minimizing the endpoint truncation error with respect to a teacher solver using gradient-based method. Optimization approaches for discretization also exist across other modalities [20].

3. Preliminaries

3.1. Diffusion ODEs for sampling

Diffusion Probabilistic Models (DPMs) [1, 2, 39] are generative models that learn to reverse a noising process. Data generation involves a learned score network, $s_{\theta}(\mathbf{x}_t, t)$, which approximates the score function $\nabla_{\mathbf{x}} \log q_t(\mathbf{x}_t)$ of the perturbed data density at time t . This score function can be equivalently parameterized via noise prediction $\epsilon_{\theta}(\mathbf{x}_t, t)$ or

data prediction $\mathbf{x}_\theta(\mathbf{x}_t, t)$.

For sampling, a widely adopted approach is to use the deterministic probability flow ordinary differential equation (PF-ODE), whose trajectories share the same marginal densities as the reverse-time SDE:

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

$$\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I}).$$

Here, the drift and diffusion-related coefficients $\mathbf{f}(t) = \frac{\dot{\alpha}_t}{\alpha_t}$ and $\mathbf{g}^2(t) = 2\dot{\sigma}_t\sigma_t - 2\frac{\dot{\alpha}_t}{\alpha_t}\sigma_t^2$ are determined by the forward process noise schedule α_t and σ_t . Common schedules are EDM-VE [26] ($\alpha_t = 1, \sigma_t = t$), Flow-Matching Optimal Transport [3–5] ($\alpha_t = 1 - t, \sigma_t = t$) Starting from a prior sample $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ (and incorporating conditional information \mathbf{c} if available), integrating Equation (1) from $t = T$ down to $t \approx 0$ using numerical ODE solvers forms the basis for efficient sample generation.

Building upon the ODE integration path in Equation (1) with noise prediction $\epsilon_\theta(\cdot, \cdot)$, various advanced solvers were designed to boost sampling [13–15]. Given a predefined discretization schedule $\xi = \{\tau_k\}_{k=0}^N$ where $\tau_0 = t_0 \rightarrow 0$ and $\tau_N = T$, the higher order multistep methods can be seen as approximations to the integrating ODE path. A generalized form can be expressed as:

$$\mathbf{x}_{k-1} = \mathcal{F}^k(\mathbf{x}_k, \xi) := u_k \cdot \mathbf{x}_k + \sum_{j=1}^M w_{k,j} \cdot \epsilon_{k+j}, \quad (2)$$

$$\epsilon_k := \epsilon_\theta(\mathbf{x}_k, \tau_k).$$

M is the order of steps, $u_k, w_{k,j}$ are the multistep coefficients dependent on the subset of discretization schedule ξ . We follow the notation of [18] and define the sampling endpoint as:

$$\begin{aligned} \mathbf{x}_0 &= \Psi(\mathbf{x}_T, \xi) = \mathcal{F}_1 \circ \mathcal{F}_2 \cdots \mathcal{F}_N(\mathbf{x}_T, \xi) \\ &= \bar{u}_1 \mathbf{x}_T + \sum_{j=1}^N \bar{w}_j \epsilon_j. \end{aligned} \quad (3)$$

Where $\bar{u}_1 = \prod_{k=1}^N u_k$, \bar{w}_j is a linear combination of $u_k, w_{k,j}$. Given a predefined network ϵ_θ and solver parameterization choice (e.g. iPNDM [15], DPM-Solver++ [30]), then the ODE path can be considered as a function of the initial value \mathbf{x}_T and timesteps ξ .

3.2. Gradient based discretization search

Traditional solver designs often rely on heuristic discretization schedules, such as Uniform $\xi = \{\tau_i = \frac{i}{N}(T - t_0) + t_0\}$ or LogSNR $\xi = \{\tau_i = \frac{i}{N}(\lambda_T - \lambda_{t_0}) + \lambda_{t_0}\}$, where $\lambda_t = \log(\alpha_t/\sigma_t)$. Because these manually crafted heuristics are typically suboptimal, recent research has increasingly focused on optimizing discretizations [18, 21, 22].

Among these efforts, gradient-based search with endpoint error supervision has proven highly competitive, as it fundamentally accounts for both approximation and accumulated truncation errors. Specifically, let ξ and ψ denote the student and teacher discretization strategies, respectively, where $|\psi| > |\xi|$. The optimization objective is formulated as:

$$\arg \min_{\xi} \mathbb{E}_{\mathbf{x}_T \sim \mathcal{N}(0, \sigma_T^2 \mathbf{I})} [d(\Psi(\mathbf{x}_T, \psi), \Psi(\mathbf{x}_T, \xi))], \quad (4)$$

where $d(\cdot, \cdot) : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ represents a distance metric (e.g., MSE or LPIPS). Minimizing this objective is equivalent to optimizing the KL divergence between student and teacher samples in the data domain.

In LD3 [18], this is termed the ‘‘hard optimization’’ objective. To ease this, LD3 introduces a soft bound by treating the initial noise \mathbf{x}_T as a learnable parameter, dynamically modifying noise-data pairs during training to reduce loss. We offer a new interpretation of this mechanism and compare it with our proposed method in Section 4.3.

4. Method

4.1. Observations on toy examples

Formally, recent learning-based methods [18, 19] aim to optimize a single global set of timestep parameters, denoted as ξ^* , which is uniformly applied across all initial samples \mathbf{x}_T . While this global optimization strategy can outperform fixed heuristic schedules, it inherently yields an optimal compromise only on average across all instances, rather than for each instance. In other words, if each sample were assigned its own optimal schedule, the global performance would naturally be optimal as well. But the reverse does not hold: a globally optimal schedule does not guarantee optimality at instance level.

This implies that the expected loss under a globally shared schedule, ϵ_g , serves as an upper bound on the expected loss achievable by an instance-specific approach, ϵ_i ; ($\epsilon_g \geq \epsilon_i$). This asymmetry motivates the following research questions: **(RQ1)** *To what extent can instance-specific discretization improve sampling performance compared to globally optimized schedule?* **(RQ2)** *How can we effectively design a conditioning mechanism that produces a tailored timestep schedule for each instance?* To answer these questions, we conduct a set of controlled experiments on synthetic example designed to isolate and quantify the benefits of instance-level scheduling.

Controlled experiment setup. We consider the recursive tree branch data distribution from [40] for better resemblance of the actual high-dimensional image data distribution, and change the noise schedule from VE to flow matching OT ($[T = 80, t_0 = 0.002] \rightarrow [T = 0.988, t_0 = 0.002]$) for better trajectory and prior analysis. (Details of the toy are in the appendix).

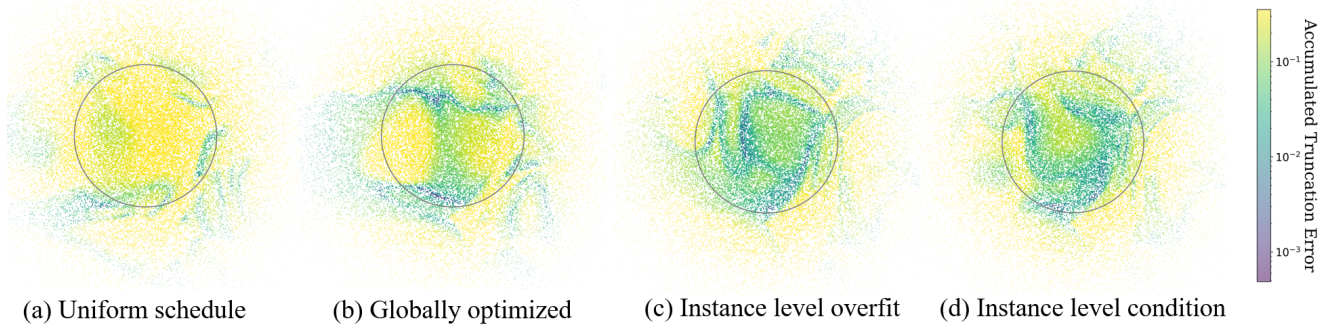


Figure 2. Comparison of endpoint(accumulated) errors for different timestep strategies (NFE=3). Each point is an initial noise sample $\mathbf{x}_T \sim \mathcal{N}(0, \sigma_T^2 \mathbf{I})$, colored by L2 error relative to 100-step euler as the ground truth. Methods: (a) Uniform timesteps. (b) Globally optimized timesteps. (c) Instance-specific timesteps (overfitted). (d) Instance-specific timesteps (learned through network ϕ).

We compare the following timestep strategies for sampling via the Euler method using 3 steps: **(a)** As a baseline, we employ uniform discretization, where timesteps are evenly spaced as $\tau_i = \frac{i}{N}(T - t_0) + t_0$. **(b)** Representing global optimization methods, we learn a single shared set of 3 timesteps by minimizing Equation (4), averaged over a training set of 20,000 prior samples $\{\mathbf{x}_T^i\}$. **(c)** As an oracle-style upper bound, we overfit a dedicated set of 3 timesteps for each individual prior sample \mathbf{x}_T^i , optimizing the error specific to that trajectory. For (b) and (c), we adopt MSE (Mean Square Error) as the distance metric $d(\cdot, \cdot)$.

Performance gap. To ensure a fair comparison, (b) and (c) are trained and evaluated using the same set of prior samples $\{\mathbf{x}_T^i\}$, which also serves as the sampling set in (a). The endpoint error is measured by comparing the final state of the sampled trajectory (using each strategy) to the ground truth obtained via 100-step Euler method from the same prior \mathbf{x}_T^i . As illustrated in Fig. 2, strategy (c), which fits timesteps individually per sample, achieves significantly lower endpoint errors than both the uniform and globally optimized schedules, particularly in the high-density region near the center of the Gaussian prior distribution. Quantitatively, the instance-specific (overfitted) achieves an average MSE of $\varepsilon_o = 0.0122$, representing a **50.2%** reduction compared to the globally optimized schedule ($\varepsilon_g = 0.0245$). This discernible performance gap forms our **core motivation**: To realize effective instance-level discretizations, we propose to directly condition the timestep strategy on the starting point \mathbf{x}_T of the deterministic ODE paths.

Conditioning on prior. Based on this motivation, we then train a lightweight network $\phi(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^N$ (where d is the noise/data dimension, N is the number of steps), taken the noise prior as input, and output instance level timesteps. We mark it as **(d)**. As illustrated in Fig. 2 (d), the conditioning is effective with regard to the global error estimation.

Quantitatively, we compare the instance-specific timestep generation (ϕ) against globally optimized timestep sched-

ules across various NFEs. The evaluation employs metrics assessing both average per-instance accuracy as MSE, and overall distributional similarity, using KL divergence and Wasserstein distance compared to ground truth distribution, to see how this instance-level correction contributes to distributional optimization objective. As illustrated in Figure 3, the results consistently demonstrate that incorporating instance-level information translates into substantial global performance gains. Statistically, this contributes to 25.86%, 21.85% on KL divergence and Wasserstein distance. The effect of instance-level condition becomes more apparent in the low NFE regimes (3-6), increasing these contributions to an average of 45.67% and 29.54%.

4.2. Scaling to discretization search for image synthesis

We now develop our practical method for high-dimensional image synthesis. We first detail two effective adaptations, then proceed to present the general framework.

Incorporating conditional guidance. Many contemporary applications of diffusion models involve conditional generation, where the sampling process is guided by auxiliary condition \mathbf{c} , such as class labels [26, 41] or text prompts [6, 27]. This guidance mechanism actively influences the evolution of the state \mathbf{x}_t along the ODE trajectory, alongside the initial noise sample \mathbf{x}_T . Our framework accommodates this by incorporating the conditional information \mathbf{c} as an additional input to the network ϕ that predicts the timestep parameters $\xi = \phi(\mathbf{x}_T, \mathbf{c})$. Specifically, we consider two types of conditional guidance \mathbf{c} , *i.e.*, class labels and prompt embedding. Details of the implemented network architecture can be referred to Fig. 4 and in the appendix.

Time and scale shift factors. The noise end state of training diffusion model often contains a minority part of data information (Take VE EDM for example, $\mathbf{x}_T = \mathbf{x}_0 + \sigma$, $\sigma \sim \mathcal{N}(0, \sigma_T \mathbf{I})$), while during sampling we always start from isotropic Gaussian $\mathbf{x}_T \sim \mathcal{N}(0, \sigma_T \mathbf{I})$) This is often noted




Figure 3. Quantitative comparison on synthetic experiments, evaluating MSE to teacher samples, KL divergence, and Wasserstein distance across various NFEs(log-scale). Methods include: uniform heuristics, globally optimized timesteps, and our proposed instance-level optimized timesteps conditioned on prior sample.

as the **exposure bias/mismatch** problem in diffusion models [42–44]. Taking this into consideration, we integrate potential correction factors (time input shifts and output scaling) as learnable parameters within our instance-aware optimization framework, rather than employing shared heuristics or statistically dependent parameters [44] used in prior work. Therefore, we adhere the setting in [18, 33] to generalize the framework for alleviating the bias problem. Specifically, we reframe the transformed function evaluation $\hat{\epsilon}_\theta$ as follows:

$$\begin{aligned} \hat{\epsilon}_\theta(\mathbf{x}_n, \tau_n, \Delta\tau_n, \gamma_n) &:= \gamma_n \cdot \epsilon_\theta(\mathbf{x}_n, \tau_n + \Delta\tau_n), \\ \xi^\phi &= \{\tau_n, \Delta\tau_n, \gamma_n\}_{n=1}^N = \phi(\mathbf{x}_T, \mathbf{c}). \end{aligned} \quad (5)$$

Where we incorporate τ and $\Delta\tau$ to make temporal and spatial shift, in an attempt to alleviate exposure bias. We define $\{\tau_n, \Delta\tau_n, \gamma_n\}_{n=1}^N$ as the general discretization set, deciding the timesteps along with the function calls, applicable to be combined with existing solver parameterization. Here we present practical implementation of the network output ϕ , given N step, we first output the instance factors:

$$\begin{aligned} O &= \phi(\mathbf{x}_T, \mathbf{c}), O = [o_\tau, o_{\Delta\tau}, o_\gamma]^T \in \mathbb{R}^{3 \times N} \\ \Delta\tau &= b_{\Delta\tau} \cdot \tanh(o_{\Delta\tau}/2), \\ \gamma &= b_\gamma \cdot \tanh(o_\gamma/2) + 1, \end{aligned} \quad (6)$$

Then we apply the softmax parameterization along with Equation (5) to obtain the final discretizations.

$$\begin{aligned} \tau &= \frac{f(\cdot) - f(0)}{f(N) - f(0)} \cdot (T - t_0) + t_0 \\ \text{where } .f(i) &= \sum_{n=i}^N \text{softmax}(o_\tau)[n]. \end{aligned} \quad (7)$$

Where $b_{\Delta\tau}, b_\gamma$ is the predefined bounding parameter to ensure stability. For parameterization of the main timestep $\{\tau_i\}_{i=1}^N$, we adhere to the setting in LD3 to ensure monotonicity. Further details are provided in the appendix.

4.3. The proposed INDIS method

Building upon synthetic analysis and tailored adaptation, here we present the implementation of instance-aware discretization INDIS in a formulated way. Given the initial sampling point $\mathbf{x}_T \in \mathbb{R}^d$ with conditional information $\mathbf{c} \in \mathbb{R}^e$

available, we design a network $\phi(\cdot, \cdot) : \mathbb{R}^d \times \mathbb{R}^e \rightarrow \mathbb{R}^{3 \times N}$, taking prior conditioning as input, then output the tailored discretizations. Given the teacher discretization strategy ψ , the optimization objective can be defined as:

$$\arg \min_{\phi} \mathbb{E}_{\mathbf{c} \sim \mathcal{C}, \mathbf{x}_T \sim \mathcal{N}(0, \sigma_T^2 \mathbf{I})} [\text{d}(\Psi(\mathbf{x}_T, \psi, \mathbf{c}), \Psi(\mathbf{x}_T, \xi^\phi, \mathbf{c}))]. \quad (8)$$

Where \mathcal{C} is the set of conditions (*e.g.* class labels, text prompts). We then present our discretization training pipeline. INDIS can be integrated into various differentiable ODE solvers for improved discretization search and with negligible additional computational overhead.

Algorithm 1 Tuning INDIS

- 1: Solver parameterization $\Psi(\cdot, \cdot)$. Condition Set \mathcal{C} if available, teacher discretization ψ , prior conditioning net ϕ
 - 2: Dataset: $\mathcal{D} \leftarrow \{\mathbf{c} \sim \mathcal{C}, \mathbf{x}_T \sim \mathcal{N}(0, \sigma_T \mathbf{I}), \mathbf{x}_0^* = \Psi(\mathbf{x}_T, \psi, \mathbf{c})\}$. ▷ Data Preparation
 - 3: **repeat**
 - 4: Sample $\mathbf{x}_T, \mathbf{c}, \mathbf{x}_0^*$ from \mathcal{D}
 - 5: $\xi^\phi = \{\tau_n, \Delta\tau_n, \gamma_n\}_{n=1}^N = \phi(\mathbf{x}_T, \mathbf{c})$
 - 6: ▷ Forward pass of prior conditioning network
 - 7: $\mathbf{x}_0 = \Psi(\mathbf{x}_T, \xi^\phi, \mathbf{c})$ ▷ $\hat{\epsilon}_n = \gamma_n \cdot \epsilon_\theta(\mathbf{x}_n, \tau_n + \Delta\tau_n)$
 - 8: Take gradient step: $\nabla_{\phi} \text{d}(\mathbf{x}_0, \mathbf{x}_0^*)$
 - 9: **until** convergence
-

Implementation details. As illustrated in Algorithm 1, our training process commences with the generation of a dataset comprising tuples of $(\mathbf{c}, \mathbf{x}_T, \mathbf{x}_0^*)$, where $\mathbf{x}_T \sim \mathcal{N}(0, \sigma_T^2 \mathbf{I})$ is the initial noise, \mathbf{c} represents optional conditional information, and \mathbf{x}_0^* is the corresponding target endpoint pre-computed using a higher-NFE teacher solver with heuristic discretization. Various multistep methods (*e.g.* DPM-Solver, UniPC, iPNDM) were considered for both teacher and student roles, we empirically found iPNDM to yield superior performance and thus selected it as the base solver structure for both. We store the random generator state instead of raw noise, thus the memory cost of this part is negligible.

During each training iteration, given an initial noise \mathbf{x}_T and \mathbf{c} from a batch, our parameter prediction network $\phi(\cdot, \cdot)$

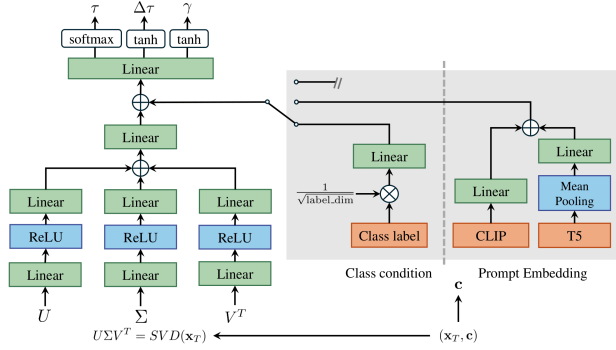


Figure 4. Architectural design of the proposed lightweight prior conditioning network. When conditional information is available, class indices are first scaled by a factor of $\frac{1}{\sqrt{\text{label_dim}}}$ and then processed through a linear layer. For prompt embeddings (FLUX.1-dev), T5 embeddings undergo mean pooling to reduce dimensionality before being concatenated with CLIP embeddings.

first calculates the instance-specific discretization parameters ξ . Subsequently, ODE sampling is performed using these tailored parameters to generate a student sample \mathbf{x}_0 . Consistent with practices in [12, 18], we employ the LPIPS [45] as the distance metric $d(\cdot, \cdot)$ in pixel domain between \mathbf{x}_0 and the cached teacher target \mathbf{x}_0^* . The parameters of the network ϕ are then updated via gradient-based optimization using Adam with cosine learning rate schedule. Further details on training implementation are provided in appendix.

At inference, generating the instance-specific parameters requires a single forward pass of $\phi(\cdot, \cdot) : \mathbb{R}^d \times \mathbb{R}^e \rightarrow \mathbb{R}^{3 \times N}$. This introduces minimal computational overhead compared to the total N evaluations of the main diffusion model ϵ_θ . Formally, the sampling requires an extra forward pass of the lightweight network ϕ to get instance-level discretizations:

$$\mathbf{x}_0 = \Psi(\mathbf{x}_T, \xi^\phi, \mathbf{c}), \quad \xi^\phi = \phi(\mathbf{x}_T, \mathbf{c}) \quad (9)$$

$$\mathbf{x}_T \sim \mathcal{N}(0, \sigma_T \mathbf{I}), \quad \mathbf{c} \sim \mathcal{C}$$

Efficiency analysis. The forward pass of our prior condition network is negligible to the function calls of diffusion model. For example, on 5-NFE setting, this overhead constitutes 2.5% of the total sampling time for CIFAR10 and 2.3% on FLUX.1-dev. Details could be referred in the appendix.

Discussion. The impact of initial noise of diffusion sampling has been extensively studied in recent literature [24, 25, 46], indicating that some noise are better than others. LD3 takes noise \mathbf{x}_T as a learnable parameter, we hypothesize that the noise is updated to get closer to better noise that result more closely to the data manifold. Then LD3 use a single set of hyperparameter timesteps to better align with those more important noise. This will help improve convergence speed and improve quality, but might quickly reach performance plateau when scaling up the dataset (e.g. > 100), due to the error between the original and updated noise given fixed image. Our work improves by directly assigning each prior \mathbf{x}_T a tailored discretization, giving the discretization more expressive power, thus shows better empirical performance when scaling up the dataset to thousands.

5. Experiments

5.1. Setup

Pretrained models. We use established pretrained diffusion and flow matching models for pixel-space and latent-space generation tasks. For pixel-space diffusion models, we adopt the official EDM [26] pretrained checkpoints for CIFAR-10 (32×32) [47], ImageNet (64×64) [48], FFHQ (64×64) [49], and AFHQv2 [50] (64×64). For latent-space text-to-image generation, we employ Stable Diffusion [27] checkpoints for LSUN-bedroom and the guidance-distilled version of Flux (FLUX.1-dev) [6]. For videos, we use LTX-Video [28].

Method	CIFAR10 32×32			FFHQ 64×64			AFHQv2 64×64			ImageNet 64×64		
	NFE=3	NFE=5	NFE=7	NFE=3	NFE=5	NFE=7	NFE=3	NFE=5	NFE=7	NFE=3	NFE=5	NFE=7
Best Heu.	57.39	17.35	7.61	72.29	17.52	8.76	40.24	9.01	4.73	44.93	15.53	8.64
DMN	77.69	12.93	5.42	178.09	20.93	10.17	178.76	26.11	13.03	33.72	10.47	5.39
AMED	18.49	7.59	4.36	26.87	12.49	6.64	31.82	7.34	4.39	28.06	10.74	6.66
GITS	25.98	6.77	3.43	26.41	8.85	5.36	24.17	8.72	5.48	26.41	9.85	6.44
LD3	16.52	5.32	3.37	23.86	8.56	4.69	17.94	6.09	2.77	27.82	11.55	5.63
INDIS	9.26	3.31	2.60	17.72	6.91	3.90	10.15	3.48	2.37	18.96	7.28	4.94

Table 1. FID Comparison for pixel-space DPMs on CIFAR10, FFHQ, AFHQv2 and class conditional ImageNet64, reporting for NFE=3, 5, and 7. Best Heu. represent the best heuristics schedule among uniform, logSNR and polynomial schedules. The complete results are provided in the appendix.




Figure 5. Ablation study on FFHQ, LSUN-Bedroom and FLUX.1-dev, instance condition is observed to be the most contributing factor, while the effect of shifted factors varies across pretrained models.

Baseline methods. We compare our method against recent open-source discretization techniques for diffusion models, specifically DMN [22], GITS [21] and LD3 [18]. We also include AMED [33] to validate the effectiveness of our instance-aware approach. Since AFS [51] can be considered an orthogonal strategy that saves one NFE compared to discretization (in small scale dataset), we report the best result from with and without AFS for fair comparison (for AMED, we apply AFS to odd NFEs). For FLUX.1-dev and LTX-video, we don't use AFS. We also include results from the best-performing heuristic schedules among previously proposed manually-crafted options based on LogSNR, Uniform, and polynomial schedules. We keep our solver choice consistent with iPNDM, we report the best result from DPM-Solver++ [15], Uni_PC [14] and iPNDM [13] for DMN, GITS and LD3. (A further evaluation on varying solvers can be referred in the appendix.) For FLUX.1-dev and LTX-video, we set resolution dependent shifted timesteps (RDS) [6] and globally optimized discretization(GOD, *i.e.*, optimizing a single set of parameters) as our baseline.

Evaluation metrics. Our primary quantitative evaluation metric is the Fréchet Inception Distance (FID) [52] over 50k generated images. For text-to-image model (FLUX.1-dev), we calculate both FID and CLIP scores [53]. These metrics are computed on a set of 10k generated images, using prompts randomly sampled from MS-COCO validation [54] set, adhering to [55]. Recognizing the potential limitations of FID for high-resolution text-to-image generation, we supplement our evaluation with CMMD [56] and provide qualitative analyses to ensure a comprehensive comparison. All quantitative metrics are averaged across three runs.

Training settings. We pre-generated fixed teacher datasets for distillation. For pixel-space and latent-space DPMs, this comprised 10,000 images and corresponding random generator states of noise produced by a 30-step iPNDM solver, serving as targets for student models across various NFEs. For FLUX.1-dev, 10,000 prompt-conditioned images were generated using a 10-step guided iPNDM solver. We also use gradient-checkpointing for reducing memory cost on Flux. For LTX-video, we use 5000 prompt-conditioned videos generated by a 7-step euler solver. Additional implementation details regarding efficiency can be referred to the appendix.

Method	CIFAR10 (pixel-space)			LSUN 256×256 (latent-space)		
	NFE=3	NFE=5	NFE=7	NFE=3	NFE=5	NFE=7
EPD	10.40	4.33	2.82	13.21	7.52	5.97
AdaSDE	12.62	4.18	2.88	18.03	6.96	5.16
INDIS	9.26	3.31	2.60	12.44	4.99	3.81

Table 3. Comparison with solver distillation on FID.

5.2. Main results

Pixel-Space DPMs. We first evaluate our instance-aware method against various discretization acceleration techniques, including those with discretization tuning and hand-crafted heuristics, on low-resolution pixel-space DPMs. As presented in Section 4.3, our approach, which conditions the discretization strategy on the initial prior sample, consistently outperforms the previous state-of-the-art methods. Specifically, compared to the strongest baseline, our method achieves average FID improvements of 35.33%, 31.50%, and 15.62% for NFE=3, 5, and 7, respectively, across four datasets. This trend indicates that the performance advantage of instance-aware discretization is more pronounced at lower NFEs, an observation consistent with our statistical analysis on the 2D synthetic examples in Fig. 3.

Latent-Space

DPMs. We further validate our instance-aware method on latent-space diffusion models (Stable Diffusion checkpoints on LSUN-Bedroom 256x256). The performance

Method	NFE=3	NFE=5	NFE=7
best heu.	41.99	6.38	4.39
DMN	28.11	6.15	5.16
AMED	58.21	13.20	7.00
GITS	44.78	17.29	9.59
LD3	14.62	5.93	4.31
INDIS	12.44	4.99	3.81

Table 2. FID on latent space LSUN 256x256.

trends observed are consistent with those from pixel-space DPMs. For LSUN-Bedroom, we achieve an average FID improvement of 14.12% across NFE=3, 5, and 7 when compared to the strongest baseline results.

Comparison with solver distillation. We also provide comparison with solver based distillation [36, 37] on pixel and latent domain.

FLUX.1-dev and LTX-Video. For FLUX.1-dev, we observe



Figure 6. Qualitative Results on FLUX.1-dev (NFE=7) of our instance-level INDIS method, compared with Global Heuristics (RDS) and globally optimized discretization (GOD). Corresponding prompts could be referred in the appendix.

its inherent robustness to hyperparameter variations, leading to a diminishing impact of different discretization strategies as NFE increases. In few-NFE settings, our approach demonstrates its effectiveness compared to global counterparts. We also verify the effectiveness of our instance discretizations on LTX-Video [28] measured on VBench [57]. It’s observed that given the latent prior of video, INDIS is able to capture instance-level benefits in discretizations, improving aesthetic, imaging quality, and subject consistency.

5.3. Ablations

We ablate key framework components: *w/o* instance-level conditioning, *w/o* shift factors, and *w/o* textual guidance (for

Metrics	Method	NFE=3	NFE=5	NFE=7
FID (↓)	RDS	64.50	30.12	22.58
	GOD	56.82	28.52	22.77
	INDIS	44.35	24.89	22.70
CLIP (↑)	RDS	23.29	29.66	30.76
	GOD	24.41	29.70	30.80
	INDIS	26.33	30.01	30.86
CMMD (↓)	RDS	1.75	0.86	0.89
	GOD	1.72	0.79	0.75
	INDIS	1.69	0.75	0.73

Table 4. Performance comparison on FLUX.1-dev (NFE=3, 5, 7).

Method	Aes. Quality	Imaging Quality	Subject Consis.
RDS	0.579	0.597	0.963
GOD	0.583	0.603	0.963
INDIS	0.593	0.613	0.964

Table 5. Performance comparison on LTX-Video (NFE=5).




Figure 7. Visualization of instance-level discretizations on FLUX.

FLUX.1-dev). As shown in Fig. 5, removing the instance-specific conditioning consistently causes the most severe performance drop. The effectiveness of shift and scale factors varies across base models. we attribute this to their differing noise schedules and inherent exposure bias severities (detailed in the appendix).

Qualitative results. Fig. 6 demonstrates that our instance-aware discretization (NFE=7) yields perceptibly stronger visual results on FLUX.1-dev, enhancing both detail and coherence over global heuristics. Additionally, Fig. 7 contrasts the distribution of our instance-level timesteps against global baselines on FLUX.1-dev.

6. Conclusion

This work demonstrates the advantage of moving beyond globally fixed discretizations for diffusion ODE sampling. Motivated by synthetic experiments, we propose an instance-aware strategy that dynamically tailors timestep schedules to the initial noise and available guidance. Extensive evaluations across diverse diffusion and flow matching models for images and videos confirm that our approach consistently improves few-step sampling performance.

Limitations and future work. Relying on gradient checkpointing for scalability (*e.g.*, on FLUX.1-dev) introduces computational overhead. Future work will explore integrating adjoint matching to optimize efficiency under specific solver parameterizations.

Acknowledgement

This work was supported by the National Natural Science Foundation of China (No. 6250070674) and the Zhejiang Leading Innovative and Entrepreneur Team Introduction Program (2024R01007).

References

- [1] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proc. NeurIPS*, 2020. 1, 2
- [2] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *Proc. ICLR*, 2021. 1, 2
- [3] Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv:2303.08797*, 2023. 1, 3
- [4] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *Proc. ICLR*, 2023.
- [5] Xingchao Liu, Chengyue Gong, and qiang liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *Proc. ICLR*, 2023. 1, 3
- [6] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. 1, 2, 4, 6, 7
- [7] Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. Diffusion-lm improves controllable text generation. In *Proc. NeurIPS*, 2022.
- [8] Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress: Discrete denoising diffusion for graph generation. In *Proc. ICLR*, 2023.
- [9] Haohe Liu, Zehua Chen, Yi Yuan, Xinhao Mei, Xubo Liu, Danilo Mandic, Wenwu Wang, and Mark D Plumbley. Audioldm: Text-to-audio generation with latent diffusion models. *arXiv:2301.12503*, 2023.
- [10] Abul Ehtesham, Saket Kumar, Aditi Singh, and Tala Talaie Khoei. Movie gen: Swot analysis of meta’s generative ai foundation model for transforming media generation, advertising, and entertainment industries. In *Proc. CCWC*, 2025. 1
- [11] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *Proc. ICLR*, 2022. 1
- [12] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *Proc. ICML*, 2023. 1, 6
- [13] Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator. *arXiv:2204.13902*, 2022. 1, 2, 3, 7
- [14] Wenliang Zhao, Lujia Bai, Yongming Rao, Jie Zhou, and Jiwen Lu. Unipc: A unified predictor-corrector framework for fast sampling of diffusion models. In *Proc. NeurIPS*, 2023. 2, 7
- [15] 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 *Proc. NeurIPS*, 2022. 1, 2, 3, 7
- [16] Feng Liu, Shiwei Zhang, Xiaofeng Wang, Yujie Wei, Haonan Qiu, Yuzhong Zhao, Yingya Zhang, Qixiang Ye, and Fang Wan. Timestep embedding tells: It’s time to cache for video diffusion model. In *Proc. CVPR*, 2025.
- [17] Xinyin Ma, Gongfan Fang, and Xinchao Wang. Deepcache: Accelerating diffusion models for free. In *Proc. CVPR*, 2024. 1
- [18] Vinh Tong, Dung Trung Hoang, Anji Liu, Guy Van den Broeck, and Mathias Niepert. Learning to discretize denoising diffusion ODEs. In *Proc. ICLR*, 2025. 1, 2, 3, 5, 6, 7
- [19] Sanghwan Kim, Hao Tang, and Fisher Yu. Distilling ode solvers of diffusion models into smaller steps. In *Proc. CVPR*, 2024. 2, 3
- [20] Yong-Hyun Park, Chieh-Hsin Lai, Satoshi Hayakawa, Yuhta Takida, and Yuki Mitsufuji. Jump your steps: Optimizing sampling schedule of discrete diffusion models. In *Proc. ICLR*, 2025. 1, 2
- [21] Defang Chen, Zhenyu Zhou, Can Wang, Chunhua Shen, and Siwei Lyu. On the trajectory regularity of ODE-based diffusion sampling. In *Proc. ICML*, 2024. 1, 2, 3, 7
- [22] Shuchen Xue, Zhaoqiang Liu, Fei Chen, Shifeng Zhang, Tianyang Hu, Enze Xie, and Zhenguo Li. Accelerating diffusion sampling with optimized time steps. In *Proc. CVPR*, 2024. 2, 3, 7
- [23] Amirmojtaba Sabour, Sanja Fidler, and Karsten Kreis. Align your steps: Optimizing sampling schedules in diffusion models. In *Proc. ICML*, 2024. 1, 2
- [24] Nanye Ma, Shangyuan Tong, Haolin Jia, Hexiang Hu, Yu-Chuan Su, Mingda Zhang, Xuan Yang, Yandong Li, Tommi Jaakkola, Xuhui Jia, et al. Inference-time scaling for diffusion models beyond scaling denoising steps. *arXiv:2501.09732*, 2025. 1, 6
- [25] Zikai Zhou, Shitong Shao, Lichen Bai, Zhiqiang Xu, Bo Han, and Zeke Xie. Golden noise for diffusion models: A learning framework. *arXiv:2411.09502*, 2024. 1, 6
- [26] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Proc. NeurIPS*, 2022. 2, 3, 4, 6
- [27] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution im-

- age synthesis with latent diffusion models. In *Proc. CVPR*, 2022. 2, 4, 6
- [28] Yoav HaCohen, Nisan Chiprut, Benny Brazowski, Daniel Shalem, Dudu Moshe, Eitan Richardson, Eran Levin, Guy Shiran, Nir Zabari, Ori Gordon, et al. Ltx-video: Realtime video latent diffusion. *arXiv preprint arXiv:2501.00103*, 2024. 2, 6, 8
- [29] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Proc. ICLR*, 2021. 2
- [30] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv:2211.01095*, 2022. 2, 3
- [31] Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on manifolds. In *Proc. ICLR*, 2022. 2
- [32] Jinyoung Choi, Junoh Kang, and Bohyung Han. Enhanced diffusion sampling via extrapolation with multiple ODE solutions. In *Proc. ICLR*, 2025. 2
- [33] Zhenyu Zhou, Defang Chen, Can Wang, and Chun Chen. Fast ode-based sampling for diffusion models in around 5 steps. In *Proc. CVPR*, 2024. 2, 5, 7
- [34] Neta Shaul, Juan Perez, Ricky T. Q. Chen, Ali Thabet, Albert Pumarola, and Yaron Lipman. Bespoke solvers for generative flow models. In *Proc. ICLR*, 2024. 2
- [35] Neta Shaul, Uriel Singer, Ricky T. Q. Chen, Matthew Le, Ali Thabet, Albert Pumarola, and Yaron Lipman. Bespoke non-stationary solvers for fast sampling of diffusion and flow models. In *Proc. ICML*, 2024. 2
- [36] Beier Zhu, Ruoyu Wang, Tong Zhao, Hanwang Zhang, and Chi Zhang. Distilling parallel gradients for fast ode solvers of diffusion models. In *Proc. ICCV*, 2025. 7
- [37] Ruoyu Wang, Beier Zhu, Junzhi Li, Liangyu Yuan, and Chi Zhang. Adaptive stochastic coefficients for accelerating diffusion sampling. In *Proc. NeurIPS*, 2025. 7
- [38] Ruoyu Wang, Ziyu Li, Beier Zhu, Liangyu Yuan, Hanwang Zhang, Xun Yang, Xiaojun Chang, and Chi Zhang. Parallel diffusion solver via residual dirichlet policy optimization, 2025. URL <https://arxiv.org/abs/2512.22796>. 2
- [39] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proc. ICML*, 2015. 2
- [40] Tero Karras, Miika Aittala, Tuomas Kynkäänniemi, Jaakko Lehtinen, Timo Aila, and Samuli Laine. Guiding a diffusion model with a bad version of itself. In *Proc. NeurIPS*, 2024. 3
- [41] Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. Analyzing and improving the training dynamics of diffusion models. In *Proc. CVPR*, 2024. 4
- [42] Shanchuan Lin, Bingchen Liu, Jiashi Li, and Xiao Yang. Common diffusion noise schedules and sample steps are flawed. In *Proc. CVPR*, 2024. 5
- [43] Mang Ning, Mingxiao Li, Jianlin Su, Albert Ali Salah, and Itir Onal Ertugrul. Elucidating the exposure bias in diffusion models. In *Proc. ICLR*, 2024.
- [44] Mingxiao Li, Tingyu Qu, Ruicong Yao, Wei Sun, and Marie-Francine Moens. Alleviating exposure bias in diffusion models through sampling with shifted time steps. In *Proc. ICLR*, 2024. 5
- [45] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. CVPR*, 2018. 6
- [46] Luca Eyring, Shyamgopal Karthik, Karsten Roth, Alexey Dosovitskiy, and Zeynep Akata. Reno: Enhancing one-step text-to-image models through reward-based noise optimization. In *Proc. NeurIPS*, 2024. 6
- [47] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. 6
- [48] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Fei-Fei Li. Imagenet large scale visual recognition challenge. *IJCV*, 2015. 6
- [49] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proc. CVPR*, 2019. 6
- [50] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proc. CVPR*, 2020. 6
- [51] Tim Dockhorn, Arash Vahdat, and Karsten Kreis. Genie: Higher-order denoising diffusion solvers. In *Proc. NeurIPS*, 2022. 7
- [52] 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 *Proc. NeurIPS*, 2017. 7
- [53] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *Proc. ICML*, 2021. 7
- [54] 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 *Proc. ECCV*, 2014. 7
- [55] Jiangshan Wang, Junfu Pu, Zhongang Qi, Jiayi Guo, Yue Ma, Nisha Huang, Yuxin Chen, Xiu Li, and Ying

Shan. Taming rectified flow for inversion and editing. *arXiv:2411.04746*, 2024. 7

- [56] Sadeep Jayasumana, Srikumar Ramalingam, Andreas Veit, Daniel Glasner, Ayan Chakrabarti, and Sanjiv Kumar. Rethinking fid: Towards a better evaluation metric for image generation. In *Proc. CVPR*, 2024. 7
- [57] Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin, Nattapol Chanpaisit, et al. Vbench: Comprehensive benchmark suite for video generative models. In *Proc. CVPR*, 2024. 8