

Dual-End Consistency Model

Linwei Dong^{1,2} Ruoyu Guo^{2†} Ge Bai² Zehuan Yuan²
Yawei Luo^{1*} Changqing Zou^{1,3}

¹Zhejiang University ²Bytedance Inc. ³Zhejiang Lab

Abstract. The slow iterative sampling nature remains a major bottleneck for the practical deployment of diffusion and flow-based generative models. While consistency models (CMs) represent a state-of-the-art distillation-based approach for efficient generation, their large-scale application is still limited by two key issues: training instability and inflexible sampling. Existing methods seek to mitigate these problems through architectural adjustments or regularized objectives, yet overlook the critical reliance on trajectory selection. In this work, we first conduct an analysis on these two limitations: training instability originates from loss divergence induced by unstable self-supervised term, whereas sampling inflexibility arises from error accumulation. Based on these insights and analysis, we propose the **Dual-End Consistency Model (DE-CM)** that selects vital sub-trajectory clusters to achieve stable and effective training. DE-CM decomposes the PF-ODE trajectory and selects three critical sub-trajectories as optimization targets. Specifically, our approach leverages continuous-time CMs objectives to achieve few-step distillation and utilizes flow matching as a boundary regularizer to stabilize the training process. Furthermore, we propose a novel noise-to-noisy (N2N) mapping that can map noise to any point, thereby alleviating the error accumulation in the first step. Extensive experimental results show the effectiveness of our method: it achieves a state-of-the-art FID score of 1.70 in one-step generation on the ImageNet 256×256 dataset, outperforming existing CM-based one-step approaches.

1 Introduction

The rapid advancement of diffusion [17, 46, 50, 51] and flow-based [7, 26, 29] generative models has revolutionized data generation across image [40, 42, 62], 3D [28, 38, 55], audio [9, 27] and video generation [1, 53], yet their reliance on slow iterative sampling remains a fundamental bottleneck. While higher-order solvers [32, 33, 47] offer limited acceleration, distillation techniques [44, 59, 60], particularly consistency models (CMs) [23, 49], have emerged as the most effective path toward few-step generation. CMs achieve efficient distillation by learning consistency function that directly maps noise to data along predefined trajectories. However, early discrete-time implementations suffered from discretization errors that potentially compromised output quality [13, 48], ultimately driving the development of continuous-time variants [31, 43].

* Corresponding Authors. † Project Leader.



Fig. 1: Left: Selected samples from DE-CM trained on ImageNet at 256×256 resolution with 2 NFE. **Right:** Images generated by our text-to-image model across 2 to 50 NFE. The prompts used are provided in the Appendix.

Continuous-time CMs, while circumventing discrete error accumulation, do not address the fundamental limitations (including training instability and inflexible sampling) of CMs, thus impeding their wider deployment. Recent main efforts on stabilization, efficiency, and flexibility fall into three categories: 1) diffusion modeling. Methods like sCM [31] that employ TrigFlow to simplify the formulation of diffusion models, which unify the flow ODE and CMs. MeanFlow [11], AYP [43] and IMM [63] introduce the flow maps [2] to map noisy inputs directly to any point along the PF-ODE trajectory. 2) network architecture. Approaches [31] tailor the architecture design to achieve more stable and efficient training. 3) training objective with regularization. Approaches like IMM, MeanFlow, SplitMeanFlow [15], sCoT [56] and FACM [37] reformulate the learning objective with additional constraints or flow modeling. While these efforts facilitate continuous-time CMs training, they remain fundamentally limited by treating suboptimal sub-trajectories as their optimization objectives. Methods like MeanFlow and AYP couple the optimization throughout the whole trajectories by learning mappings between any two points, as shown in Fig. 2 (b). These entangled learning objectives are particularly problematic given that CMs impose consistency as a network behavior property without an explicit boundary condition from the ground-truth flow characteristics, leading to slower convergence or even training collapse. Furthermore, approaches like sCMs treat vanilla CMs trajectory alone, potentially increasing instability due to the absence of the critical path. By explicitly decoupling the instantaneous velocity objective along the PF-ODE trajectory, FACM achieves improved training stability. This indicates that the core challenge lies in selecting the appropriate candidate trajectories to achieve both training stability and strong performance.

In this work, we begin with an analysis of two inherent limitations in vanilla CMs: unstable training and inflexible sampling. Our findings indicate that training instability arises from loss divergence, which is triggered by unstable self-




Fig. 2: Left (a): Comparison of FID scores across different models under various NFE settings, showing the superior performance of our method in both few-step and multi-step sampling. **Right (b):** Comparison of learning objectives. CMs Distillation [49] ($s = T$), Flow Matching [26] ($s = t$), MeanFlow [11] / AYF [43] (upper triangle). DE-CM (triangular boundary).

supervised objectives and the catastrophic forgetting of original instantaneous velocity features. Furthermore, we attribute inflexible sampling to error accumulation caused by the CMs sampler [23]. Based on these analysis, we propose a novel method, **Dual-End Consistency Model**, to achieve efficient, stable, and flexible model distillation under critical sub-trajectory clusters selection. We employ flow maps to incorporate entire transport space and identify the significant sub-trajectory clusters for optimization. Taking the intermediate time t as the dividing point, we focus on the following three trajectory segments: 1) Consistency Trajectory. 2) Instantaneous Trajectory. 3) Noise-to-Noise Trajectory. The comparative illustration of our optimization trajectories is provided in Fig. 2 (b). Consistency trajectories enable the model to perform inference over few-step via consistency property, while instantaneous trajectories serve as a boundary regularization to stabilize the training. By incorporating noise-to-noisy trajectories, we formulate a training objective that enables our model to alleviate the error accumulation imposed by the CMs sampler. DE-CM excels in few-step sampling on C2I and T2I tasks while maintaining performance across multi-steps, as show in Fig. 1. On ImageNet 256×256 [5], DE-CM achieves 1.70 FID using 1-NFE generation, as shown in Fig. 2 (a).

The key contributions are summarized as follows:

- We analyze the potential constraints imposed on CMs and elucidate the underlying causes of instability and inflexible sampling.
- We propose DE-CM, a novel method for stable and effective training through the selection of vital trajectories.
- We derive a new objective noise-to-noisy mapping to learn a mapping from pure noise to arbitrary points.
- Our experiments on C2I and T2I demonstrate the effectiveness of DE-CM, achieving state-of-the-art performance on the ImageNet dataset with 250 training epoch.

2 Preliminaries

Diffusion Models and Flow Matching. Given a dataset with underlying distribution p_{data} and standard deviation σ , diffusion models learn to reverse a noising process that transforms a clean sample $x_1 \sim p_{data}$ into $x_t = \alpha_t x_1 + \sigma_t z$, where $z \sim \mathcal{N}(0, I)$, and the noise level increases as the timestep t decreases from T to 0. EDM [22] defines $\alpha_t = 1, \sigma_t = 1 - t$, and optimizes:

$$\mathbb{E}_{x_1, z, t} [w(t) \|f_{\theta}^{\text{DM}}(x_t, t) - x_1\|_2^2], \quad (1)$$

where $w(t)$ is the weighting and $f_{\theta}(x_t, t)$ is defined as:

$$f_{\theta}(x_t, t) = c_{\text{skip}}(t)x_t + c_{\text{out}}(t)F_{\theta}(c_{\text{in}}(t)x_t, c_{\text{noise}}(t)), \quad (2)$$

Here, F_{θ} is a neural network, and these terms $c(t)$ are scaling coefficients that stabilize training.

Flow matching (FM) uses differentiable schedules α_t, σ_t (e.g, $\alpha_t = t, \sigma_t = 1 - t$) and minimizes:

$$\mathbb{E}_{x_0, z, t} [w(t) \|F_{\theta}(x_t, t) - (\dot{\alpha}_t x_0 + \dot{\sigma}_t z)\|_2^2], \quad (3)$$

where $\dot{\alpha}_t$ and $\dot{\sigma}_t$ denote the derivatives with respect to t . Sampling solves the probability flow ODE (PF-ODE):

$$\frac{dx_t}{dt} = F_{\theta}(x_t, t), \quad (4)$$

from $x_0 \sim \mathcal{N}(0, I)$ at $t = 0$ to $t = 1$.

Consistency Models. Consistency Models (CMs) are designed to map any point x_t along a PF-ODE trajectory directly to its endpoint x_1 in a single step. This behavior is governed by a self-consistency property: for any pair of points x_t and $x_{t'}$ on the same trajectory, the model must satisfy $f_{\theta}(x_t, t) = f_{\theta}(x_{t'}, t')$. This implies $f_{\theta}(x_t, t) = x_1$ for all $t \in [0, 1]$ with the boundary condition $f_{\theta}(x_1, 1) = x_1$.

Discrete-time CMs follow the Eq. (2) with $c_{\text{skip}}(1) = 1, c_{\text{out}}(1) = 0$ and learn the consistency between two adjacent timesteps. The training objective is defined as follows:

$$\mathbb{E}_{x_1 \sim p_{data}} [w(t) \cdot d(f_{\theta}(x_t, t), f_{\theta^-}(x_{t+\Delta t}, t + \Delta t))], \quad (5)$$

where $w(t)$ is the time-aware weighting, $\Delta t > 0$ is a minimal step, θ^- denotes $\text{stopgrad}(\theta)$ and $d(\cdot, \cdot)$ is a distance metric. Discrete-time CMs are sensitive to the discretization step size Δt , typically requiring carefully designed schedules to ensure convergence. Furthermore, the sample $x_{t+\Delta t}$ is commonly generated from x_t by numerically solving the PF-ODE process that introduces additional discretization error, potentially leading to suboptimal sample quality.

Continuous-time CMs avoid the accumulation of discrete errors by taking the limit $\Delta t \rightarrow 0$. When employing $d(x, y) = \|x - y\|_2^2$, the gradient of Eq. (5) converges to

$$\nabla_{\theta} \mathbb{E}_{x_t, t} \left[-w(t) f_{\theta}^{\top}(x_t, t) \frac{df_{\theta^-}(x_t, t)}{dt} \right], \quad (6)$$




Fig. 3: **Left (a):** Gradient stability comparison between supervised and self-supervised loss functions. Self-supervised loss exhibits more unstable gradients. **Right (b):** Visualization results of different NFE sampling using the CM sampler.

where $\frac{df_{\theta-}(\mathbf{x}_t, t)}{dt} = \nabla_{\mathbf{x}_t} f_{\theta-}(\mathbf{x}_t, t) \frac{d\mathbf{x}_t}{dt} + \partial_t f_{\theta-}(\mathbf{x}_t, t)$ is the tangent of $f_{\theta-}$ at $(\mathbf{x}_t, t)$ along the trajectory of the PF-ODE $\frac{d\mathbf{x}_t}{dt}$ and $w(t)$ is the weighting function. This modeling approach eliminates the need for a predefined sampling interval Δt , thus avoiding discretization errors.

3 Limitations of Consistency Models

3.1 Training instability

Similar to discrete versions, continuous-time CMs face the challenge with training instability [13, 48]. To better understand the causes of instability, we analyze its ultimate training objective. Following the naive discrete-time CMs based on OT-FM, we set $c_{skip}(t) = 1$ and $c_{out}(t) = 1 - t$. Thus the training objective of Eq. (6) can be derived as follows:

$$\mathbb{E}_{x_t, t} [\|F_{\theta}(x_t, t) - F_{\theta-}(x_t, t) + w'(t) \cdot g_{\theta-}\|_2^2], \quad (7)$$

where $g_{\theta-} = F_{\theta-}(x_t, t) - v_{\phi} - (1 - t) \frac{dF_{\theta-}(x_t, t)}{dt}$, v_{ϕ} denotes the desired trajectory tangent $\frac{dx}{dt}$ and it can be directly derived from data ($v_{\phi} = \dot{x}_1 + \dot{\sigma}x_0$) or predicted by a teacher model ($v_{\phi} = v_{\text{tea}}(x_t, t)$). ϵ is equal to Δt , $w'(t) = \frac{1}{2}w(t)\epsilon(1 - t)$ and $w(t)$ is a custom weighting function with respect to t and commonly set to $\frac{1}{\epsilon}$ to accelerate convergence. To clearly delineate the training objectives, we further decompose Eq. (7):

$$\mathbb{E}_{x_t, t} [\|F_{\theta}(x_t, t) - v_{\phi}\|_2^2 + \|F_{\theta}(x_t, t) - \tilde{v}\|_2^2], \quad (8)$$

where $\tilde{v} = F_{\theta-}(x_t, t) + (1 - t) \frac{dF_{\theta-}}{dt}$. We train the two losses independently, with the gradients generated during training shown in Fig. 3 (a). Compared to

full-supervised term, self-supervised loss is more prone to instability and volatility. This reveals that: 1) A divergence exists between the fully supervised and self-supervised components when F_θ has not converged. 2) The emphasis on consistency in F results in the overshadowing of the pretrained instantaneous velocity concept. These suggest inherent limitations in the optimization trajectory predefined by the vanilla CMs.

3.2 Inflexible Sampling Mechanism

The challenges of inflexible sampling have been thoroughly analyzed in Consistency Trajectory Models (CTM) [23]. Let $t \in (0, T)$ and $\gamma \in [0, 1]$. We denote p_θ as the density derived from a sampling trajectory following the transition sequence $T \rightarrow \sqrt{1 - \gamma^2}t \rightarrow t \rightarrow 0$, initiated from the prior distribution p_T . Then, we can obtain that

$$D_{TV}(p_{\text{data}}, p_\theta) = \mathcal{O}(\sqrt{T - \sqrt{1 - \gamma^2}t + t}), \quad (9)$$

where $D_{TV}(p, q) := \frac{1}{2} \int |p(\mathbf{x}) - q(\mathbf{x})| d\mathbf{x}$. When generalized to N steps, the sampling behavior is critically governed by the interpolation parameter γ . For $\gamma = 1$, the sampler iteratively performs long jumps from each t_n to 0, which leads to significant time overlaps between jumps and causes the approximation error to aggregate to $\mathcal{O}(\sqrt{T + t_1 + \dots + t_N})$. In contrast, setting $\gamma = 0$ eliminates these temporal overlaps and prevents error accumulation, resulting in a stable error bound of only $\mathcal{O}(\sqrt{T})$. Consequently, by leveraging the $\gamma = 0$ regime (e.g. Euler Solver), diffusion model can effectively addresses the performance degradation associated with high NFE in distillation models while simultaneously bypassing the discretization errors inherent in score-based models. However, to satisfy the consistency function and boundary conditions, the sampling mechanism of CMs is based on $\gamma = 1$. The absence of intermediate state transitions imposes the following limitations [23, 49] on naive CM sampling, explaining its suboptimal performance: 1) Inflexible Sampling Steps. The model is constrained to a fixed step count, limiting adaptability. 2) Artifact Accumulation. Multi-step sampling leads to an accumulation of errors, which often manifests as visual artifacts. 3) Sensitivity to Initial Noise. The generation quality is highly sensitive to the choice of initial noise, leading to poor robustness and high output variance. As shown in Fig. 3 (b), due to the accumulation of errors, the generated quality first increases and then decreases as the NFE gradually increases.

4 Dual-End Consistency Model

MeanFlow [11] and AYF [43] indicate that the aforementioned limitation of CMs lies in the exclusive dependence on the time condition t , which is inherently insufficient to characterize the entire PF-ODE trajectory. Consequently, they augment the temporal specification with an auxiliary endpoint embedding s ,




Fig. 4: We select significant trajectories from the whole $\{(t, s) | t < s\}$ space and treat these selected trajectories as our optimization targets. Specifically, we employ the continuous-time consistency distillation trajectory to optimize the mapping from arbitrary time points to data, thereby achieving few-step distillation. We leverage the proposed noise-to-noisy (N2N) mapping objective to eliminate the constraint of predicting only x_1 points. We utilize flow matching loss to overcome the instability in training. We find that incorporating these three trajectory clusters is enough to enable stable and efficient optimization for few-step distillation.

enabling the estimation of the mean velocity along the trajectory. The trajectory consistency function is formulated as follows:

$$f_{\theta}(x_t, t, s) = x_t + (s - t)F_{\theta}(x_t, t, s) \quad (10)$$

where F_{θ} is a neural network. Eq. (10) encompasses the entire trajectories when given $t, s \in [0, 1]$ and $t < s$, thus satisfying both CMs and FM parameterization. This formulation will optimize any sub-trajectory defined by t and s and have access to ODE solver (e.g. Euler Sampler). However, unlike the consistency function that exhibits explicit boundary conditions $f_{\theta}(x_1, 1) = x_1$, the boundary conditions of Eq. (10) is not strictly constrained as t, s are both variables. This will introduces greater instability risks during training. Moreover, optimizing the entire sub-trajectory space also leads to inefficient training as the redundant trajectory optimization degenerates convergence. The computational complexity, $O(n^2)$, became a primary bottleneck, leading to slow optimization. Fortunately, we find that optimizing only specific sub-trajectory clusters is sufficient to ensure stable training and high-quality samples. We focus on three key trajectory clusters (Fig. 4): consistency trajectories, instantaneous trajectories and noise-to-noisy trajectories.

4.1 Few-Step Models with Consistency Trajectory

Given $t \in [0, 1]$ and $s = 1$, Eq. (10) can be equated to the consistency function of standard CMs. By analogy with the derivation of Eq. (7), we can further simplify:

$$\mathcal{L}_{\text{cm}} = \mathbb{E}_{x_t, t, s=1} \left[\|F_{\theta}(x_t, t, s) - v_{\text{cm}}^{\text{tar}}\|_2^2 \right], \quad (11)$$

where $v_{\text{cm}}^{\text{tar}} = w'(t)\bar{v} + (1-w'(t))F_{\theta^-}(x_t, t, 1)$ and $\bar{v} = v_\phi + (1-t)\frac{dF_{\theta^-}(x_t, t, 1)}{dt}$. It can be seen that the final loss is a weighted combination of the model prediction and another velocity $\bar{v}$. Moreover, $\bar{v}$ is defined as the instantaneous velocity $v_\phi(x_t, t)$ corrected by the average acceleration $\frac{dF_{\theta^-}(x_t, t, 1)}{dt}$, thus representing the average velocity [11]. Given that $\frac{dF_{\theta^-}(x_t, t, 1)}{dt}$ serves as the training objective but is not available in advance, the standard approach is to utilize Jacobian-Vector product (JVP) to dynamically compute it based on current neural network states. The total derivative, expanded via the chain rule, is:

$$\frac{dF_{\theta^-}}{dt} = [\nabla_{x_t} F_{\theta^-}, \frac{\partial F_{\theta^-}}{\partial t}, \frac{\partial F_{\theta^-}}{\partial s}] \cdot [v_\phi, 1, 0]^\top, \quad (12)$$

where $[v_\phi, 1, 0]$ is drawn from $[\frac{dx_t}{dt}, \frac{dt}{dt}, \frac{ds}{dt}]$, and $\frac{dx_t}{dt} = v_\phi$ represents the direction given by the pretrained diffusion or flow model along the PF-ODE. The consistency properties enable the model to learn the mapping from noisy input x_t to data x_1 . This is analogous to reflow that construct paired data, alleviating the difficulty of learning few-step models.

4.2 Flexible Models with Noise-to-Noisy Mapping

By optimizing the time interval $[0, t]$, we enable the model to map from pure noise x_0 to any intermediate noisy state x_t , rather than returning only to the data x_1 . This helps mitigate the error accumulation caused by learning only long jumps, by exposing the model to noise inputs similar to those encountered during inference. To preserve the meaning of the random variable t , we reset the left endpoint in Eq. (10) by $r \rightarrow 0$, while the right remains t . Following Eq. (5), for timestep t , under the L_2 norm, we have:

$$\mathbb{E}_{x_r, r, t | r \rightarrow 0} [w(r, t) \|f_\theta(x_r, r, t), \mathcal{S}_\psi(f_{\theta^-}(x_r, r, t'))\|_2^2], \quad (13)$$

where $\mathcal{S}_\psi$ refers to *Euler Solver* on the PF-ODE starting from $x_{t'}$ to x_t ($t > t'$). Inspired by [43, 56], its gradient can be derived to converge to

$$\nabla_\theta \mathbb{E}_{x_r, r, t | r \rightarrow 0} \left[w(r, t)' f_{\theta^-}^\top(x_r, r, t) \left(\frac{df_{\theta^-}(x_r, r, t)}{dt} - v_\psi \right) \right] \quad (14)$$

where $v_\psi = v_{\text{tea}}(f_{\theta^-}(x_r, r, t), t)$ denotes the desired trajectory tangent $\frac{d\hat{x}}{dt}$. Eq. (14) reveals the gradient convergence objective of the self-supervised loss Eq. (13) with respect to the right endpoint t . Considering the left endpoint r , we can regard it as a consistency function based on boundary $f_\theta(x_t, t, t) = x_t$. Then, its gradient objective can be derived from Eq. (6). The two optimization objectives can be expressed in weighted form, and their gradients can be proven to converge to

$$\nabla_\theta \mathbb{E}_{x_r, r, t | r \rightarrow 0} \left[w(r, t)' f_{\theta^-}^\top(x_r, r, t) \left(\lambda \frac{df_{\theta^-}(x_r, r, t)}{dt} - \gamma \frac{df_{\theta^-}(x_r, r, t)}{dr} - \lambda \cdot v_{\text{tea}}(f_{\theta^-}(x_r, r, t), t) \right) \right], \quad (15)$$




Fig. 5: Left (a): Comparison of gradient norm curves with and without flow matching boundary constraints. **Right (b):** Training Efficiency Comparison of Different Methods on 8-GPU. Our method achieves a more efficient convergence rate.

where λ, γ are weighting factors. The final optimization loss $\mathcal{L}_{n2n}$ can be expressed as:

$$\mathbb{E}_{x_r, r, t | r \rightarrow 0} [\|F_\theta(x_t, t, s) - F_{\theta-}(x_t, t, s) + w'(r, t)g\|_2^2], \quad (16)$$

where $g = (\lambda + \gamma)F_{\theta-} - (\gamma v_\phi + \lambda v_\psi) - (t - r)\dot{F}_{\theta-}$ and $\dot{F}_{\theta-} = [\nabla_{x_t} F_{\theta-}, \frac{\partial F_{\theta-}}{\partial t}, \frac{\partial F_{\theta-}}{\partial s}] \cdot [\lambda v_\phi, \lambda, -\gamma]^\top$. Eq. (16) utilizes the velocities on both sides of the endpoint r and t , along with the correction term from the neural network.

4.3 Stable Models with Instantaneous Velocity

Directly using Eq. (11) or Eq. (16) for optimization does not stabilize training. We find this is due to the absence of necessary boundary conditions. Following Eq. (8), We can restate the training objective when given t and s :

$$\mathbb{E}_{x_t, t, s} [\|F_\theta(x_t, t, s) - v_\phi\|_2^2 + \|F_\theta(x_t, t, s) - \tilde{v}\|_2^2], \quad (17)$$

where $\tilde{v} = F_{\theta-}(x_t, t, s) + (s - t)\frac{dF_{\theta-}}{dt}$. The dynamic calculation of $\frac{dF_{\theta-}}{dt}$ in Eq. (12) renders the self-supervised term susceptible to instability, which raises the risk that online models could diverge from a single erroneous update. The large gradient phenomena are shown in Fig. 5 (a).

However, when $s = t$, we observe that the original gradient update paradigm is driven solely by ground truth as the second term equals zero. This term also represents the left boundary condition. The single loss term makes the optimization objective deterministic, eliminating slack on the left-hand side. Similarly, for targets with loose right-hand constraints like Eq. (16), a right-hand regularization can be incorporated. This approach corresponds to selecting all instantaneous trajectory clusters and it makes training more stable (w/ flow in Fig. 5 (a)). The specific loss $\mathcal{L}_{fm}$ is formulated as:

$$\mathbb{E}_{x_t, t, s | s=t} [\|F_\theta(x_t, t, s) - v_\phi\|_2^2 + L_{cos}(F_\theta(x_t, t, s), v_\phi)], \quad (18)$$

Beyond its role as a robust regularizer, this loss function is designed to safeguard the pre-trained model’s capabilities, thus supporting multi-step ODE sampling.

The overall training process is shown in Algorithm 1.

Algorithm 1: DE-CM Training Procedure

Input: Pretrained diffusion model F_{tea} , online model F_θ , paired dataset $\mathcal{D} = \{z_{\text{ref}}, y_{\text{ref}}\}$, hyperparameters λ, γ , weighting function $w(r, t)$.
Output: Trained model F_θ , ema model F_{ema} .

```

// Initialize model
1  $F_\theta \leftarrow \text{copyWeights}(F_{\text{tea}})$ ,  $F_{\text{ema}} \leftarrow \text{copyWeights}(F_\theta)$ 
2 while train do
    // Prepare Data
    3 Sample batch  $z \sim \mathcal{N}(0, \mathbf{I})$  and  $(z_{\text{ref}}, y_{\text{ref}}) \sim \mathcal{D}$ 
    4 Sample timestep  $t \sim \mathcal{U}(0, 1)$ ,  $r \rightarrow 0$ 
    5  $z_t \leftarrow (1 - t)z + tz_{\text{ref}}$ ,  $v_t^{\text{cond}} \leftarrow F_{\text{tea}}(z_t, t; y_{\text{ref}})$ ,  $v_t^{\text{uncond}} \leftarrow F_{\text{tea}}(z_t, t; \phi)$ 
    6  $v_t \leftarrow v_t^{\text{uncond}} + w_{\text{cfg}}(v_t^{\text{cond}} - v_t^{\text{uncond}})$ 
    7
    // Continuous-time CD
    8  $\mathcal{L}_{\text{cm}} \leftarrow \text{ContinuousCDLoss}(F_\theta(z_t, t, 1; y_{\text{ref}}), v_{\text{cm}}^{\text{tar}})$  // Eq 11
    9
    // Flow Matching
    10  $\mathcal{L}_{\text{fm}} \leftarrow \text{FlowMatchingLoss}(F_\theta(z_t, t, t; y_{\text{ref}}), v_t)$  // Eq 18
    11
    // Noise-to-Noisy Mapping
    12  $z_r \leftarrow (1 - r)z + rz_{\text{ref}}$ 
    13  $v_r^{\text{cond}} \leftarrow F_{\text{tea}}(z_r, r; y_{\text{ref}})$ ,  $v_r^{\text{uncond}} \leftarrow F_{\text{tea}}(z_r, r; \phi)$ 
    14  $v_r \leftarrow v_r^{\text{uncond}} + w_{\text{cfg}}(v_r^{\text{cond}} - v_r^{\text{uncond}})$ 
    15  $F_{\text{n2n}}, \nabla_t F_{\text{n2n}} \leftarrow \text{JVP}(F_\theta, (x_r, r, t; y_{\text{ref}}), (\lambda v_r, \lambda, -\gamma))$ 
    16  $g_{\text{n2n}} \leftarrow (\lambda + \gamma) \cdot \text{sg}(F_{\text{n2n}}) - \lambda(v_r + (t - r)\nabla_t F_{\text{n2n}}) - \gamma v_t$ 
    17  $v_{\text{n2n}}^{\text{tar}} \leftarrow \text{sg}(F_{\text{n2n}}) - w(r, t) \cdot g_{\text{n2n}}$ 
    18  $\mathcal{L}_{\text{n2n}} \leftarrow \text{MSELoss}(F_{\text{n2n}}, v_{\text{n2n}}^{\text{tar}})$  // Eq 16
    19  $F_\theta, F_{\text{ema}} \leftarrow \text{update}(F_\theta, \mathcal{L}_{\text{n2n}} + \mathcal{L}_{\text{cm}} + \mathcal{L}_{\text{fm}})$ 
20 end

```

5 Experiments

5.1 Setup

For class-to-image (C2I) generation, we use ImageNet-256 $\times$ 256 datasets [5]. We implement our models in the latent space of a pre-trained VAE tokenizer [58], which produces a latent dimensionality of 32 $\times$ 16 $\times$ 16. Building upon the pre-trained LightningDiT [58], we introduce an additional input s to F_θ to condition the model on the right timestep. We evaluate using Fréchet Inception Distance (FID) [16] on 50K generated images. And we train the C2I model with learning rate 1e-4 and AdamW [30] optimizer.

For text-to-image (T2I) generation, we use 100K data from text-to-image-2M datasets [20]. We utilize the default SD3-VAE tokenizer. We leverage a reward model [57] to assess generation quality, incorporating CLIP score [39], BLIP




Fig. 6: Selected samples from DE-CM trained on ImageNet at 256×256 resolution with 1 NFE.

score [25], and ImageReward [57]. We adopt LoRA [18] with a rank of 64, a learning rate of $5e-4$, and the AdamW optimizer.

5.2 Qualitative and Quantitative Comparison

Class-to-Image Generation. In Tab. 1 we show ImageNet 256×256 results, reporting FID scores along with the NFE (Number of Function Evaluations). Our model establishes a new state-of-the-art at 1 NFE among comparable-scale models and even surpasses the vast majority of multi-step models with just 2 NFE, demonstrating robust and superior performance. Furthermore, our model maintains stable performance at large NFE using a single set of parameters. Fig. 6 presents the visual results generated by our model using 1 NFE sampling.

Text-to-Image Generation. Fig. 7 provides a visualization of our method against competing models across various inference budgets (1, 2, 4, 32 and 50 NFE). While pre-trained models like Flux [24] and SD3.5 [7] are inherently incapable of few-step generation, existing CMs like LCM [34] and CTM [23] suffer from blurry results at low NFE and oversaturation at large NFE. Methods including PCM [54] and Hyper-SD [41], which require dedicated training for each NFE, face challenges with oversaturation and poor instruction alignment (e.g. fail to render Van Gogh style at 2 NFE). DE-CM demonstrates compelling performance, achieving appealing results in few-step inference and realistic quality

Table 1: Class-conditional generation on ImageNet-256×256. ‘× 2’ indicates that CFG doubles the NFE per step. † indicates distillation reproduction (250 epoch) based on our configuration.

Method	NFE (↓)	#Params (↓)	Time (↓)	FID (↓)
GAN				
GigaGAN [21]	1	569M	-	3.45
StyleGAN-XL [45]	1	166M	0.3	2.30
Masked & AR				
VQGAN [8]	256	227M	19	3.04
MaskGIT [3]	8	227M	0.5	6.18
VAR-d30 [52]	10	2B	1	1.92
Diffusion & Flow				
ADM [6]	250	280M	16	10.94
DiT-XL/2 ($w=1.5$) [36]	250 × 2	675M	>3.75	2.27
SiT-XL/2 ($w=1.5$) [35]	250 × 2	675M	> 2.89	2.15
LightningDiT-XL/1 ($w=6.7$)	250 × 2	675M	> 2.82	1.35
Consistency Model				
MeanFlow [11]	1	675M	0.011	3.43
	2		0.024	2.20
MeanFlow-† [11]	1	675M	0.041	2.79
	2		0.048	1.74
IMM [63]	1 × 2	675M	0.057	7.77
	2 × 2		0.068	3.99
FACM-† [37]	1	675M	0.011	1.76
iCT [48]	1	675M	-	34.24
	2		-	20.30
sCMs-† [31]	1	675M	0.011	3.34
	2		0.024	1.94
Shortcut [10]	1	675M	0.048	10.60
DE-CM (Ours)	1	675M	0.011	1.70
	2		0.024	1.33
	50		0.58	1.26

Table 2: Quantity comparison on pre-trained and CM distillation-based text-to-image model. ‘× 2’ indicates that CFG doubles the NFE per step. The best results of each metric are highlighted in **bold**.

Method	NFE (↓)	CLIP (↑)	BLIP (↑)	ImageRewards (↑)
Multi-step Models				
SD3.5-Medium [7]	4	0.2395	0.4457	-1.6017
	50 × 2	0.2998	0.5451	0.8750
SD3.5-Large [7]	4	0.2385	0.4221	-1.582
	50 × 2	0.3024	0.5465	0.8349
Flux.1-Dev [24]	4	0.2543	0.4523	-1.2984
	50 × 2	0.2941	0.5384	0.9634
DE-CM (Ours)	50	0.2993	0.5490	0.9712
Few-step Models				
SD3.5-Turbo [9]	1	0.2579	0.4974	-0.8519
	4	0.2953	0.5418	0.6844
	8	0.2935	0.5334	0.6513
LCM [34]	1	0.2508	0.4560	-1.0622
	4	0.2948	0.5306	0.4910
	8	0.2993	0.5318	0.5871
CTM [23]	1	0.2624	0.4764	-0.7639
	4	0.2986	0.5284	0.5703
	8	0.2993	0.5278	0.6762
Hyper-SD [41]	1	0.2939	0.5161	0.5704
	4	0.2949	0.5307	0.7931
	8	0.3029	0.5293	0.8447
PCM [54]	1	0.2726	0.4764	-0.6584
	4	0.2981	0.5265	0.6025
	8	0.2996	0.5286	0.6195
DE-CM (Ours)	1	0.2996	0.5398	0.5758
	4	0.2999	0.5474	0.8117
	8	0.3000	0.5479	0.8671

at large NFE. Fig. 9 further demonstrates the visualization results of our 1 NFE. Compared to other methods, our approach exhibits finer textures and enhanced realism. Tab. 2 presents a quantitative evaluation using feedback scores. Our model achieves strong performance across various metrics, notably BLIP scores, and exhibits superior sample quality universally across all NFE settings.

Efficiency Comparison. Fig. 8 presents 1 NFE visual comparisons under 16 GPU hours training. Compared to MeanFlow, DE-CM yields clearer images, while compared to sCMs, it produces fewer artifacts. This demonstrates the powerful training efficiency of DE-CM. Fig. 5 (b) further illustrates the convergence performance of DE-CM under 8 GPUs resources, demonstrating strong efficiency and performance.

5.3 Ablation Studies

Tab. 4 presents ablation studies under different training objectives. We can draw the following conclusions: 1) The absence of CD loss prevents the achievement of few-step inference, as the distillation process is consistent-function-driven. 2) The absence of flow matching loss leads to suboptimal results due to unstable training. 3) The absence of N2N loss affects multi-step sampling results due to the absence of initial noise perception. We further investigate the impact of various hyperparameters and experimental configurations on the results. Covering all




Fig. 7: Qualitative visualization reveals DE-CM superiority in quality and text alignment across all NFE. Existing competing models exhibit significant step-quality trade-offs, failing at low NFE (SD3.5-Medium, FLUX.1-Dev) or producing oversaturated samples at high NFE (LCM, CTM, PCM and Hyper-SD). DE-CM maintains robust performance throughout the efficiency-quality spectrum.




Fig. 8: Comparison of 1 NFE results under 8 GPUs resources. DE-CM demonstrated exceptional convergence efficiency at 16 GPU hours compared to sCMs and MeanFlow.

losses enables all steps to achieve significant gains. Detailed implementation details are provided in the appendix.

5.4 Discussion with Recent Efforts

Our method is related to recent efforts on consistency-style training, but differs from them in both motivation and formulation.

Table 3: Comparison with recent methods. † indicates our repro.

Method	1NFE-FID	2NFE-FID	1NFE-FID †	2NFE-FID †
AlphaFlow	2.58	2.15	2.24	1.71
CMT	3.34	-	2.71	1.68
DE-CM	-	-	1.70	1.33

Discussion on stability. Shortcut-style methods [10] rely on flow matching (FM) to anchor the regime near $d = 0$, making stability highly dependent on the model behavior as $d \rightarrow 0$. MeanFlow [11] implicitly includes FM through its flow-ratio formulation, and sLCT [4] also employs FM, though only in the low-noise regime. In contrast, DE-CM identifies the source of CMs instability by decomposing the continuous-time objective in Eq. (8). From this perspective, FM acts as a left-boundary regularizer that offsets the instability introduced by the self-supervised term, rather than serving as the primary target of training. This explicit formulation also enables more flexible loss design, e.g., the additional $L_{\cos}$ term in Eq. (18), and extends naturally to the full noise range.

Further discussion on BOOT. Our method is also related to BOOT [14], since both optimize the segment from pure noise to x_t along the PF-ODE trajectory. However, BOOT is proposed for data-free distillation through a bootstrapping mechanism based on a discrete consistency property, whereas our N2N trajectory (Eq. (16)) is introduced to mitigate the error accumulation caused by long-jump CM sampling and is derived from the continuous-time consistency model. N2N therefore represents a theoretically grounded trajectory family within DE-CM. The full derivation is provided in the appendix.

Comparison to recent efforts. Recent methods [12, 19, 61] improve training from different perspectives. iMF [12] reformulates MeanFlow as a v -prediction objective; Alphaflow [61] progressively adjusts α during training, transitioning from FM to shortcut-style training and eventually to MeanFlow; and CMT [19] introduces an intermediate training stage to provide a more stable initialization. In contrast, DE-CM is derived from an objective-level diagnosis of CM instability and provides a corresponding correction mechanism. Quantitative comparisons are summarized in Tab. 3.

6 Conclusion and Limitations

In this work, we propose **Dual-End Consistency Model (DE-CM)**, a novel framework that effectively addresses the dual challenges of training instability and sampling inflexibility in CMs. DE-CM integrates continuous-time CMs objectives for efficient few-step distillation while employing flow matching as a boundary regularizer to ensure training stability. Furthermore, we introduce an innovative noise-to-noisy (N2N) mapping mechanism that enables flexible mapping from noise to any trajectory point, thus supporting to alleviate error accumulation. Extensive experiments demonstrate that our method sets a




Fig. 9: Visual quality comparison between DE-CM and existing methods under 1 NFE. DE-CM achieves superior naturalness in detail rendering, outperforming existing methods. The prompts used are provided in the Appendix.

Table 4: Ablation study on ImageNet-256 (FID $\downarrow$) analyzing the impact of training objective across different inference steps. We compared the model’s performance under missing FM, CD, and N2N losses respectively.

Method	STEP=1	STEP=2	STEP=4	STEP=50
Training Objectives				
Baseline Teacher [58]	309.51	230.79	111.59	2.77
Baseline Teacher ($w = 1.75$)	262.60	221.03	87.56	1.59
Flow Matching (FM)	263.86	220.55	87.93	1.45
CMs Distillation (CD)	3.08	2.50	2.54	6.52
FM + CD	1.78	1.41	1.62	1.78
FM + Noise-to-Noisy (N2N) Transmission	441.66	297.28	132.60	1.41
FM + CD + N2N	1.70	1.33	1.41	1.26

new state-of-the-art FID score of 1.70 for one-step generation on the ImageNet 256 $\times$ 256 benchmark with 250 epoch training. Due to the incompatibility of the JVP operator with the FSDP framework and Flash Attention, the optimization process consumes a substantial amount of GPU memory, which limits the training of larger-scale models. Existing approaches try to address these problems, and it is hoped that more effective alternatives will emerge in the future.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grants No. 62293554 and U2336212; the “Pioneer” and “Leading Goose” R&D Program of Zhejiang under Grants No. 2025C02022 and 2024C01161; the Zhejiang Provincial Natural Science Foundation of China under Grants No. LZ24F020002 and LD24F020007; the Ningbo Innovation “Yongjiang 2035” Key Research and Development Programme under Grant No. 2024Z292; and the Young Elite Scientists Sponsorship Program by CAST under Grant No. 2023QNRC001. The authors also gratefully acknowledge the support of the Zhejiang University Education Foundation Qizhen Scholar Foundation.

References

1. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023)
2. Boffi, N.M., Albergo, M.S., Vanden-Eijnden, E.: Flow map matching. arXiv preprint arXiv:2406.07507 **2**(3), 9 (2024)
3. Chang, H., Zhang, H., Jiang, L., Liu, C., Freeman, W.T.: Maskgit: Masked generative image transformer. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11315–11325 (2022)
4. Dao, Q., Doan, K., Liu, D., Le, T., Metaxas, D.: Improved training technique for latent consistency models. arXiv preprint arXiv:2502.01441 (2025)
5. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
6. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
7. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., et al.: Scaling rectified flow transformers for high-resolution image synthesis. In: Forty-first international conference on machine learning (2024)
8. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12873–12883 (2021)
9. Evans, Z., Carr, C., Taylor, J., Hawley, S.H., Pons, J.: Fast timing-conditioned latent audio diffusion. In: Forty-first International Conference on Machine Learning (2024)
10. Frans, K., Hafner, D., Levine, S., Abbeel, P.: One step diffusion via shortcut models. arXiv preprint arXiv:2410.12557 (2024)
11. Geng, Z., Deng, M., Bai, X., Kolter, J.Z., He, K.: Mean flows for one-step generative modeling. arXiv preprint arXiv:2505.13447 (2025)
12. Geng, Z., Lu, Y., Wu, Z., Shechtman, E., Kolter, J.Z., He, K.: Improved mean flows: On the challenges of fastforward generative models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 30467–30476 (2026)

13. Geng, Z., Pokle, A., Luo, W., Lin, J., Kolter, J.Z.: Consistency models made easy. arXiv preprint arXiv:2406.14548 (2024)
14. Gu, J., Zhai, S., Zhang, Y., Liu, L., Susskind, J.M.: Boot: Data-free distillation of denoising diffusion models with bootstrapping. In: ICML 2023 Workshop on Structured Probabilistic Inference & Generative Modeling (2023)
15. Guo, Y., Wang, W., Yuan, Z., Cao, R., Chen, K., Chen, Z., Huo, Y., Zhang, Y., Wang, Y., Liu, S., et al.: Splitmeanflow: Interval splitting consistency in few-step generative modeling. arXiv preprint arXiv:2507.16884 (2025)
16. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* **30** (2017)
17. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
18. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: Lora: Low-rank adaptation of large language models. *ICLR* **1**(2), 3 (2022)
19. Hu, Z., Lai, C.H., Mitsufuji, Y., Ermon, S.: Cmt: Mid-training for efficient learning of consistency, mean flow, and flow map models. arXiv preprint arXiv:2509.24526 (2025)
20. Jackyhate: Text-to-image-2m. <https://huggingface.co/datasets/jackyhate/text-to-image-2M> (2025). <https://doi.org/10.57967/hf/3066>
21. Kang, M., Zhu, J.Y., Zhang, R., Park, J., Shechtman, E., Paris, S., Park, T.: Scaling up gans for text-to-image synthesis. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10124–10134 (2023)
22. Karras, T., Aittala, M., Aila, T., Laine, S.: Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems* **35**, 26565–26577 (2022)
23. Kim, D., Lai, C.H., Liao, W.H., Murata, N., Takida, Y., Uesaka, T., He, Y., Mitsufuji, Y., Ermon, S.: Consistency trajectory models: Learning probability flow ode trajectory of diffusion. arXiv preprint arXiv:2310.02279 (2023)
24. black-forest labs: Flux.1-dev. <https://huggingface.co/black-forest-labs/FLUX.1-dev> (2024)
25. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: *International conference on machine learning*. pp. 12888–12900. PMLR (2022)
26. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. arXiv preprint arXiv:2210.02747 (2022)
27. Liu, H., Chen, Z., Yuan, Y., Mei, X., Liu, X., Mandic, D., Wang, W., Plumbley, M.D.: Audioldm: Text-to-audio generation with latent diffusion models. arXiv preprint arXiv:2301.12503 (2023)
28. Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 9298–9309 (2023)
29. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. arXiv preprint arXiv:2209.03003 (2022)
30. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
31. Lu, C., Song, Y.: Simplifying, stabilizing and scaling continuous-time consistency models. arXiv preprint arXiv:2410.11081 (2024)
32. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in neural information processing systems* **35**, 5775–5787 (2022)

33. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *Machine Intelligence Research* pp. 1–22 (2025)
34. Luo, S., Tan, Y., Huang, L., Li, J., Zhao, H.: Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378* (2023)
35. Ma, N., Goldstein, M., Albergo, M.S., Boffi, N.M., Vanden-Eijnden, E., Xie, S.: Sit: Exploring flow and diffusion-based generative models with scalable interpolant transformers. In: *European Conference on Computer Vision*. pp. 23–40. Springer (2024)
36. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 4195–4205 (2023)
37. Peng, Y., Zhu, K., Liu, Y., Wu, P., Li, H., Sun, X., Wu, F.: Flow-anchored consistency models. *arXiv preprint arXiv:2507.03738* (2025)
38. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988* (2022)
39. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning*. pp. 8748–8763. PmLR (2021)
40. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* **1**(2), 3 (2022)
41. Ren, Y., Xia, X., Lu, Y., Zhang, J., Wu, J., Xie, P., Wang, X., Xiao, X.: Hyper-sd: Trajectory segmented consistency model for efficient image synthesis. *Advances in Neural Information Processing Systems* **37**, 117340–117362 (2024)
42. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)
43. Sabour, A., Fidler, S., Kreis, K.: Align your flow: Scaling continuous-time flow map distillation. *arXiv preprint arXiv:2506.14603* (2025)
44. Sauer, A., Lorenz, D., Blattmann, A., Rombach, R.: Adversarial diffusion distillation. In: *European Conference on Computer Vision*. pp. 87–103. Springer (2024)
45. Sauer, A., Schwarz, K., Geiger, A.: Stylegan-xl: Scaling stylegan to large diverse datasets. In: *ACM SIGGRAPH 2022 conference proceedings*. pp. 1–10 (2022)
46. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *International conference on machine learning*. pp. 2256–2265. pmlr (2015)
47. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020)
48. Song, Y., Dhariwal, P.: Improved techniques for training consistency models. *arXiv preprint arXiv:2310.14189* (2023)
49. Song, Y., Dhariwal, P., Chen, M., Sutskever, I.: Consistency models (2023)
50. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems* **32** (2019)
51. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456* (2020)
52. Tian, K., Jiang, Y., Yuan, Z., Peng, B., Wang, L.: Visual autoregressive modeling: Scalable image generation via next-scale prediction. *Advances in neural information processing systems* **37**, 84839–84865 (2024)

53. Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., Zeng, J., Wang, J., Zhang, J., Zhou, J., Wang, J., Chen, J., Zhu, K., Zhao, K., Yan, K., Huang, L., Feng, M., Zhang, N., Li, P., Wu, P., Chu, R., Feng, R., Zhang, S., Sun, S., Fang, T., Wang, T., Gui, T., Weng, T., Shen, T., Lin, W., Wang, W., Wang, W., Zhou, W., Wang, W., Shen, W., Yu, W., Shi, X., Huang, X., Xu, X., Kou, Y., Lv, Y., Li, Y., Liu, Y., Wang, Y., Zhang, Y., Huang, Y., Li, Y., Wu, Y., Liu, Y., Pan, Y., Zheng, Y., Hong, Y., Shi, Y., Feng, Y., Jiang, Z., Han, Z., Wu, Z.F., Liu, Z.: Wan: Open and advanced large-scale video generative models. arXiv preprint arXiv:2503.20314 (2025)
54. Wang, F.Y., Huang, Z., Bergman, A., Shen, D., Gao, P., Lingelbach, M., Sun, K., Bian, W., Song, G., Liu, Y., et al.: Phased consistency models. *Advances in neural information processing systems* **37**, 83951–84009 (2024)
55. Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J.: Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *Advances in neural information processing systems* **36**, 8406–8441 (2023)
56. Wu, Z., Fan, X., Wu, H., Cao, L.: Traftlow: Trajectory distillation on pre-trained rectified flow. arXiv preprint arXiv:2502.16972 (2025)
57. Xu, J., Liu, X., Wu, Y., Tong, Y., Li, Q., Ding, M., Tang, J., Dong, Y.: Imagereward: learning and evaluating human preferences for text-to-image generation. In: *Proceedings of the 37th International Conference on Neural Information Processing Systems*. pp. 15903–15935 (2023)
58. Yao, J., Yang, B., Wang, X.: Reconstruction vs. generation: Taming optimization dilemma in latent diffusion models. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 15703–15712 (2025)
59. Yin, T., Gharbi, M., Park, T., Zhang, R., Shechtman, E., Durand, F., Freeman, B.: Improved distribution matching distillation for fast image synthesis. *Advances in neural information processing systems* **37**, 47455–47487 (2024)
60. Yin, T., Gharbi, M., Zhang, R., Shechtman, E., Durand, F., Freeman, W.T., Park, T.: One-step diffusion with distribution matching distillation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 6613–6623 (2024)
61. Zhang, H., Siarohin, A., Menapace, W., Vasilkovsky, M., Tulyakov, S., Qu, Q., Skorokhodov, I.: Alphaflow: Understanding and improving meanflow models. arXiv preprint arXiv:2510.20771 (2025)
62. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 3836–3847 (2023)
63. Zhou, L., Ermon, S., Song, J.: Inductive moment matching. arXiv preprint arXiv:2503.07565 (2025)