

SIMPLIFYING, STABILIZING & SCALING CONTINUOUS-TIME CONSISTENCY MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Consistency models (CMs) are a powerful class of diffusion-based generative models optimized for fast sampling. Most existing CMs are trained using discretized timesteps, which introduce additional hyperparameters and are prone to discretization errors. While continuous-time formulations can mitigate these issues, their success has been limited by training instability. To address this, we propose a simplified theoretical framework that unifies previous parameterizations of diffusion models and CMs, identifying the root causes of instability. Based on this analysis, we introduce key improvements in diffusion process parameterization, network architecture, and training objectives. These changes enable us to train continuous-time CMs at an unprecedented scale, reaching 1.5B parameters on ImageNet 512×512. Our proposed training algorithm, using only two sampling steps, achieves FID scores of 2.06 on CIFAR-10, 1.48 on ImageNet 64×64, and 1.88 on ImageNet 512×512, narrowing the gap in FID scores with the best existing diffusion models to within 10%.

1 INTRODUCTION

Diffusion models (Sohl-Dickstein et al., 2015; Song & Ermon, 2019; Ho et al., 2020; Song et al., 2021b) have revolutionized generative AI, achieving remarkable results in image (Rombach et al., 2022; Ramesh et al., 2022; Ho et al., 2022), 3D (Poole et al., 2022; Wang et al., 2024; Liu et al., 2023b), audio (Liu et al., 2023a; Evans et al., 2024), and video generation (Blattmann et al., 2023; Brooks et al., 2024). Despite their success, a significant drawback is their slow sampling speed, often requiring dozens to hundreds of steps to generate a single sample. Various diffusion distillation techniques have been proposed, including direct distillation (Luhman & Luhman, 2021; Zheng et al., 2023b), adversarial distillation (Wang et al., 2022; Sauer et al., 2023), progressive distillation (Salimans & Ho, 2022), and variational score distillation (VSD) (Wang et al., 2024; Yin et al., 2024b;a; Luo et al., 2024; Xie et al., 2024b; Salimans et al., 2024). However, these methods come with challenges: direct distillation incurs extensive computational cost due to the need for numerous diffusion model samples; adversarial distillation introduces complexities associated with GAN training; progressive distillation requires multiple training stages and is less effective for one or two-step generation; and VSD can produce overly smooth samples with limited diversity and struggles at high guidance levels.

Consistency models (CMs) (Song et al., 2023; Song & Dhariwal, 2023) offer significant advantages in addressing these issues. They eliminate the need for supervision from diffusion model samples, avoiding the computational cost of generating synthetic datasets. CMs also bypass adversarial training, sidestepping its inherent difficulties. Aside from distillation, CMs can be trained from scratch with consistency training (CT), without relying on pre-trained diffusion models. Previous work (Song & Dhariwal, 2023; Geng et al., 2024; Luo et al., 2023; Xie et al., 2024a) has demonstrated the




Figure 1: Sample quality vs. effective sampling compute (billion parameters \times number of function evaluations during sampling). We compare the sample quality of different models on ImageNet 512×512, measured by FID (\downarrow). Our 2-step sCM achieves sample quality comparable to the best previous generative models while using less than 10% of the effective sampling compute.



Figure 2: Selected 2-step samples from a continuous-time consistency model trained on ImageNet 512×512 .

effectiveness of CMs in few-step generation, especially in one or two steps. However, these results are all based on discrete-time CMs, which introduces discretization errors and requires careful scheduling of the timestep grid, potentially leading to suboptimal sample quality. In contrast, continuous-time CMs avoid these issues but have faced challenges with training instability (Song et al., 2023; Song & Dhariwal, 2023; Geng et al., 2024).

In this work, we introduce techniques to simplify, stabilize, and scale up the training of continuous-time CMs. Our first contribution is TrigFlow, a new formulation that unifies EDM (Karras et al., 2022; 2024) and Flow Matching (Peluchetti, 2022; Lipman et al., 2022; Liu et al., 2022; Albergo et al., 2023; Heitz et al., 2023), significantly simplifying the formulation of diffusion models, the associated probability flow ODE and CMs. Building on this foundation, we analyze the root causes of instability in CM training and propose a complete recipe for mitigation. Our approach includes improved time-conditioning and adaptive group normalization within the network architecture. Additionally, we re-formulate the training objective for continuous-time CMs, incorporating adaptive weighting and normalization of key terms, and progressive annealing for stable and scalable training.

With these improvements, we elevate the performance of consistency models in both consistency training and distillation, achieving comparable or better results compared to previous discrete-time formulations. Our models, referred to as sCMs, demonstrate success across various datasets and model sizes. We train sCMs on CIFAR-10, ImageNet 64×64 , and ImageNet 512×512 , reaching an unprecedented scale with 1.5 billion parameters—the largest CMs trained to date (samples in Figure 2). We show that sCMs scale effectively with increased compute, achieving better sample quality in a predictable way. Moreover, when measured against state-of-the-art diffusion models, which require significantly more sampling compute, sCMs narrow the FID gap to within 10% using two-step generation. In addition, we provide a rigorous justification for the advantages of continuous-

time CMs over discrete-time variants by demonstrating that sample quality improves as the gap between adjacent timesteps narrows to approach the continuous-time limit. Furthermore, we examine the differences between sCMs and VSD, finding that sCMs produce more diverse samples and are more compatible with guidance, whereas VSD tends to struggle at higher guidance levels.

2 PRELIMINARIES

2.1 DIFFUSION MODELS

Given a training dataset, let p_d denote its underlying data distribution and σ_d its standard deviation. Diffusion models generate samples by learning to reverse a noising process that progressively perturbs a data sample $\mathbf{x}_0 \sim p_d$ into a noisy version $\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \mathbf{z}$, where $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is standard Gaussian noise. This perturbation increases with $t \in [0, T]$, where larger t indicates greater noise.

We consider two recent formulations for diffusion models.

EDM (Karras et al., 2022; 2024). The noising process simply sets $\alpha_t = 1$ and $\sigma_t = t$, with the training objective given by $\mathbb{E}_{\mathbf{x}_0, \mathbf{z}, t} [w(t) \|\mathbf{f}_\theta^{\text{DM}}(\mathbf{x}_t, t) - \mathbf{x}_0\|_2^2]$, where $w(t)$ is a weighting function. The diffusion model is parameterized as $\mathbf{f}_\theta^{\text{DM}}(\mathbf{x}_t, t) = c_{\text{skip}}(t)\mathbf{x}_t + c_{\text{out}}(t)\mathbf{F}_\theta(c_{\text{in}}(t)\mathbf{x}_t, c_{\text{noise}}(t))$, where \mathbf{F}_θ is a neural network with parameters θ , and c_{skip} , c_{out} , c_{in} , and c_{noise} are manually designed coefficients that ensure the training objective has the unit variance across timesteps at initialization. For sampling, EDM solves the *probability flow ODE (PF-ODE)* (Song et al., 2021b), defined by $\frac{d\mathbf{x}_t}{dt} = [\mathbf{x}_t - \mathbf{f}_\theta^{\text{DM}}(\mathbf{x}_t, t)]/t$, starting from $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, T^2 \mathbf{I})$ and stopping at \mathbf{x}_0 .

Flow Matching. The noising process uses differentiable coefficients α_t and σ_t , with time derivatives denoted by α'_t and σ'_t (typically, $\alpha_t = 1 - t$ and $\sigma_t = t$). The training objective is given by $\mathbb{E}_{\mathbf{x}_0, \mathbf{z}, t} [w(t) \|\mathbf{F}_\theta(\mathbf{x}_t, t) - (\alpha'_t \mathbf{x}_0 + \sigma'_t \mathbf{z})\|_2^2]$, where $w(t)$ is a weighting function and \mathbf{F}_θ is a neural network parameterized by θ . The sampling procedure begins at $t = 1$ with $\mathbf{x}_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and solves the probability flow ODE (PF-ODE), defined by $\frac{d\mathbf{x}_t}{dt} = \mathbf{F}_\theta(\mathbf{x}_t, t)$, from $t = 1$ to $t = 0$.

2.2 CONSISTENCY MODELS

A consistency model (CM) (Song et al., 2023; Song & Dhariwal, 2023) is a neural network $\mathbf{f}_\theta(\mathbf{x}_t, t)$ trained to map the noisy input \mathbf{x}_t directly to the corresponding clean data \mathbf{x}_0 in one step, by following the sampling trajectory of the PF-ODE starting at \mathbf{x}_t . A valid \mathbf{f}_θ must satisfy the *boundary condition*, $\mathbf{f}_\theta(\mathbf{x}, 0) \equiv \mathbf{x}$. One way to meet this condition is to parameterize the consistency model as $\mathbf{f}_\theta(\mathbf{x}_t, t) = c_{\text{skip}}(t)\mathbf{x}_t + c_{\text{out}}(t)\mathbf{F}_\theta(c_{\text{in}}(t)\mathbf{x}_t, c_{\text{noise}}(t))$ with $c_{\text{skip}}(0) = 1$ and $c_{\text{out}}(0) = 0$. CMs are trained to have consistent outputs at adjacent time steps. Depending on how nearby time steps are selected, there are two categories of consistency models, as described below.

Discrete-time CMs. The training objective is defined at two adjacent time steps with finite distance:

$$\mathbb{E}_{\mathbf{x}_t, t} [w(t)d(\mathbf{f}_\theta(\mathbf{x}_t, t), \mathbf{f}_\theta^-(\mathbf{x}_{t-\Delta t}, t - \Delta t))], \quad (1)$$

where θ^- denotes $\text{stopgrad}(\theta)$, $w(t)$ is the weighting function, $\Delta t > 0$ is the distance between adjacent time steps, and $d(\cdot, \cdot)$ is a metric function; common choices are ℓ_2 loss $d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_2^2$, Pseudo-Huber loss $d(\mathbf{x}, \mathbf{y}) = \sqrt{\|\mathbf{x} - \mathbf{y}\|_2^2 + c^2} - c$ for $c > 0$ (Song & Dhariwal, 2023), and LPIPS loss (Zhang et al., 2018). Discrete-time CMs are sensitive to the choice of Δt , and therefore require manually designed annealing schedules (Song & Dhariwal, 2023; Geng et al., 2024) for fast convergence. The noisy sample $\mathbf{x}_{t-\Delta t}$ at the preceding time step $t - \Delta t$ is often obtained from \mathbf{x}_t

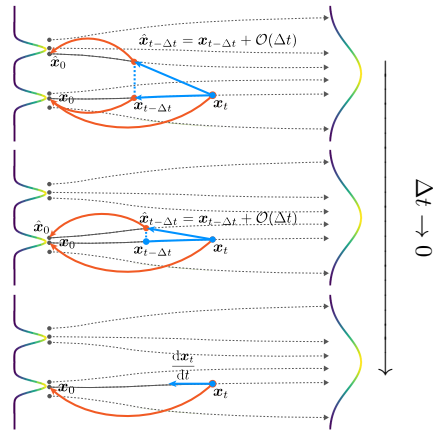


Figure 3: Discrete-time CMs (top & middle) vs. continuous-time CMs (bottom). Discrete-time CMs suffer from discretization errors from numerical ODE solvers, causing imprecise predictions during training. In contrast, continuous-time CMs stay on the ODE trajectory by following its tangent direction with infinitesimal steps.

by solving the PF-ODE with numerical ODE solvers using step size Δt , which can cause additional discretization errors.

Continuous-time CMs. When using $d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_2^2$ and taking the limit $\Delta t \rightarrow 0$, Song et al. (2023, Remark 10) show that the gradient of Eq. (1) with respect to θ converges to

$$\nabla_{\theta} \mathbb{E}_{\mathbf{x}_t, t} \left[w(t) \mathbf{f}_{\theta}^{\top}(\mathbf{x}_t, t) \frac{d\mathbf{f}_{\theta}(\mathbf{x}_t, t)}{dt} \right], \quad (2)$$

where $\frac{d\mathbf{f}_{\theta}(\mathbf{x}_t, t)}{dt} = \nabla_{\mathbf{x}_t} \mathbf{f}_{\theta}(\mathbf{x}_t, t) \frac{d\mathbf{x}_t}{dt} + \partial_t \mathbf{f}_{\theta}(\mathbf{x}_t, t)$ is the *tangent* of \mathbf{f}_{θ} at (\mathbf{x}_t, t) along the trajectory of the PF-ODE $\frac{d\mathbf{x}_t}{dt}$. Notably, continuous-time CMs do not rely on ODE solvers, which avoids discretization errors and offers more accurate supervision signals during training. However, previous work (Song et al., 2023; Geng et al., 2024) found that training continuous-time CMs, or even discrete-time CMs with an extremely small Δt , suffers from severe instability in optimization. This greatly limits the empirical performance and adoption of continuous-time CMs.

Consistency Distillation and Consistency Training. Both discrete-time and continuous-time CMs can be trained using either *consistency distillation* (CD) or *consistency training* (CT). In consistency distillation, a CM is trained by distilling knowledge from a pretrained diffusion model. This diffusion model provides the PF-ODE, which can be directly plugged into Eq. (2) for training continuous-time CMs. Furthermore, by numerically solving the PF-ODE to obtain $\mathbf{x}_{t-\Delta t}$ from \mathbf{x}_t , one can also train discrete-time CMs via Eq. (1). Consistency training (CT), by contrast, trains CMs from scratch without the need for pretrained diffusion models, which establishes CMs as a standalone family of generative models in their own right. Specifically, CT approximates $\mathbf{x}_{t-\Delta t}$ in discrete-time CMs as $\mathbf{x}_{t-\Delta t} = \alpha_{t-\Delta t} \mathbf{x}_0 + \sigma_{t-\Delta t} \mathbf{z}$, reusing the same data \mathbf{x}_0 and noise \mathbf{z} when sampling $\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \mathbf{z}$. In the continuous-time limit, as $\Delta t \rightarrow 0$, this approach yields an unbiased estimate of the PF-ODE $\frac{d\mathbf{x}_t}{dt} \rightarrow \alpha'_t \mathbf{x}_0 + \sigma'_t \mathbf{z}$, leading to an unbiased estimate of Eq. (2) for training continuous-time CMs.

3 SIMPLIFYING CONTINUOUS-TIME CONSISTENCY MODELS

Previous consistency models (CMs) adopt the model parameterization and diffusion process formulation in EDM (Karras et al., 2022). Specifically, the CM is parameterized as $\mathbf{f}_{\theta}(\mathbf{x}_t, t) = c_{\text{skip}}(t) \mathbf{x}_t + c_{\text{out}}(t) \mathbf{F}_{\theta}(c_{\text{in}}(t) \mathbf{x}_t, c_{\text{noise}}(t))$, where \mathbf{F}_{θ} is a neural network with parameters θ . The coefficients $c_{\text{skip}}(t)$, $c_{\text{out}}(t)$, $c_{\text{in}}(t)$ are fixed to ensure that the variance of the diffusion objective is equalized across all time steps at initialization, and $c_{\text{noise}}(t)$ is a transformation of t for better time conditioning. Since EDM diffusion process is variance-exploding (Song et al., 2021b), meaning that $\mathbf{x}_t = \mathbf{x}_0 + t\mathbf{z}$, we can derive that $c_{\text{skip}}(t) = \sigma_d^2/(t^2 + \sigma_d^2)$, $c_{\text{out}}(t) = \sigma_d \cdot t/\sqrt{\sigma_d^2 + t^2}$, and $c_{\text{in}}(t) = 1/\sqrt{t^2 + \sigma_d^2}$ (see Appendix B.6 in Karras et al. (2022)). Although these coefficients are important for training efficiency, their complex arithmetic relationships with t and σ_d complicate theoretical analyses of CMs.

To simplify EDM and subsequently CMs, we propose *TrigFlow*, a formulation of diffusion models that keep the EDM properties but satisfy $c_{\text{skip}}(t) = \cos(t)$, $c_{\text{out}}(t) = -\sigma_d \sin(t)$, and $c_{\text{in}}(t) \equiv 1/\sigma_d$ (proof in Appendix B). TrigFlow is a special case of flow matching (also known as stochastic interpolants or rectified flows) and v-prediction parameterization (Salimans & Ho, 2022). It closely resembles the trigonometric interpolant proposed by Albergo & Vanden-Eijnden (2023); Albergo et al. (2023); Ma et al. (2024), but is modified to account for σ_d , the standard deviation of the data distribution p_d . Since TrigFlow is a special case of flow matching and simultaneously satisfies EDM principles, it combines the advantages of both formulations while allowing the diffusion process, diffusion model parameterization, the PF-ODE, the diffusion training objective, and the CM parameterization to all have simple expressions, as provided below.

Diffusion Process. Given $\mathbf{x}_0 \sim p_d(\mathbf{x}_0)$ and $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \sigma_d^2 \mathbf{I})$, the noisy sample is defined as $\mathbf{x}_t = \cos(t) \mathbf{x}_0 + \sin(t) \mathbf{z}$ for $t \in [0, \frac{\pi}{2}]$. As a special case, the prior sample $\mathbf{x}_{\frac{\pi}{2}} \sim \mathcal{N}(\mathbf{0}, \sigma_d^2 \mathbf{I})$.

Diffusion Models and PF-ODE. We parameterize the diffusion model as $\mathbf{f}_{\theta}^{\text{DM}}(\mathbf{x}_t, t) = \mathbf{F}_{\theta}(\mathbf{x}_t/\sigma_d, c_{\text{noise}}(t))$, where \mathbf{F}_{θ} is a neural network with parameters θ , and $c_{\text{noise}}(t)$ is a transformation of t to facilitate time conditioning. The corresponding PF-ODE is given by

$$\frac{d\mathbf{x}_t}{dt} = \sigma_d \mathbf{F}_{\theta} \left(\frac{\mathbf{x}_t}{\sigma_d}, c_{\text{noise}}(t) \right). \quad (3)$$

Diffusion Objective. In TrigFlow, the diffusion model is trained by minimizing

$$\mathcal{L}_{\text{Diff}}(\theta) = \mathbb{E}_{\mathbf{x}_0, \mathbf{z}, t} \left[\left\| \sigma_d \mathbf{F}_\theta \left(\frac{\mathbf{x}_t}{\sigma_d}, c_{\text{noise}}(t) \right) - \mathbf{v}_t \right\|_2^2 \right], \quad (4)$$

where $\mathbf{v}_t = \cos(t)\mathbf{z} - \sin(t)\mathbf{x}_0$ is the training target.

Consistency Models. As mentioned in Sec. 2.2, a valid CM must satisfy the boundary condition $\mathbf{f}_\theta(\mathbf{x}, 0) \equiv \mathbf{x}$. To enforce this condition, we parameterize the CM as the single-step solution of the PF-ODE in Eq. (3) using the first-order ODE solver (see Appendix B.1 for derivations). Specifically, CMs in TrigFlow take the form of

$$\mathbf{f}_\theta(\mathbf{x}_t, t) = \cos(t)\mathbf{x}_t - \sin(t)\sigma_d \mathbf{F}_\theta \left(\frac{\mathbf{x}_t}{\sigma_d}, c_{\text{noise}}(t) \right), \quad (5)$$

where $c_{\text{noise}}(t)$ is a time transformation for which we defer the discussion to Sec. 4.1.

4 STABILIZING CONTINUOUS-TIME CONSISTENCY MODELS

Training continuous-time CMs has been highly unstable (Song et al., 2023; Geng et al., 2024). As a result, they perform significantly worse compared to discrete-time CMs in prior works. To address this issue, we build upon the TrigFlow framework and introduce several theoretically motivated improvements to stabilize continuous-time CMs, with a focus on parameterization, network architecture, and training objectives.

4.1 PARAMETERIZATION AND NETWORK ARCHITECTURE

Key to the training of continuous-time CMs is Eq. (2), which depends on the tangent function $\frac{d\mathbf{f}_{\theta-}(\mathbf{x}_t, t)}{dt}$. Under the TrigFlow formulation, this tangent function is given by

$$\frac{d\mathbf{f}_{\theta-}(\mathbf{x}_t, t)}{dt} = -\cos(t) \left(\sigma_d \mathbf{F}_{\theta-} \left(\frac{\mathbf{x}_t}{\sigma_d}, t \right) - \frac{d\mathbf{x}_t}{dt} \right) - \sin(t) \left(\mathbf{x}_t + \sigma_d \frac{d\mathbf{F}_{\theta-} \left(\frac{\mathbf{x}_t}{\sigma_d}, t \right)}{dt} \right), \quad (6)$$

where $\frac{d\mathbf{x}_t}{dt}$ represents the PF-ODE, which is either estimated using a pretrained diffusion model in consistency distillation, or using an unbiased estimator calculated from noise and clean samples in consistency training.

To stabilize training, it is necessary to ensure the tangent function in Eq. (6) is stable across different time steps. Empirically, we found that $\sigma_d \mathbf{F}_{\theta-}$, the PF-ODE $\frac{d\mathbf{x}_t}{dt}$, and the noisy sample \mathbf{x}_t are all relatively stable. The only term left in the tangent function now is $\sin(t) \frac{d\mathbf{F}_{\theta-}}{dt} = \sin(t) \nabla_{\mathbf{x}_t} \mathbf{F}_{\theta-} \frac{d\mathbf{x}_t}{dt} + \sin(t) \partial_t \mathbf{F}_{\theta-}$. After further analysis, we found $\nabla_{\mathbf{x}_t} \mathbf{F}_{\theta-} \frac{d\mathbf{x}_t}{dt}$ is typically well-conditioned, so instability originates from the time-derivative $\sin(t) \partial_t \mathbf{F}_{\theta-}$, which can be decomposed according to

$$\sin(t) \partial_t \mathbf{F}_{\theta-} = \sin(t) \frac{\partial c_{\text{noise}}(t)}{\partial t} \cdot \frac{\partial \text{emb}(c_{\text{noise}})}{\partial c_{\text{noise}}} \cdot \frac{\partial \mathbf{F}_{\theta-}}{\partial \text{emb}(c_{\text{noise}})}, \quad (7)$$

where $\text{emb}(\cdot)$ refers to the time embeddings, typically in the form of either positional embeddings (Ho et al., 2020; Vaswani, 2017) or Fourier embeddings (Song et al., 2021b; Tancik et al., 2020) in the literature of diffusion models and CMs.

Below we describe improvements to stabilize each component from Eq. (7) in turns.

Identity Time Transformation ($c_{\text{noise}}(t) = t$). Most existing CMs use the EDM formulation, which can be directly translated to the TrigFlow formulation as described in Appendix B.2. In particular, the time transformation becomes $c_{\text{noise}}(t) \propto \log(\sigma_d \tan t)$. Straightforward derivation shows that with this $c_{\text{noise}}(t)$, $\sin(t) \cdot \partial_t c_{\text{noise}}(t) = 1/\cos(t)$ blows up whenever $t \rightarrow \frac{\pi}{2}$. To mitigate numerical instability, we propose to use $c_{\text{noise}}(t) = t$ as the default time transformation.

Positional Time Embeddings. For general time embeddings in the form of $\text{emb}(c) = \sin(s \cdot 2\pi\omega \cdot c + \phi)$, we have $\partial_c \text{emb}(c) = s \cdot 2\pi\omega \cos(s \cdot 2\pi\omega \cdot c + \phi)$. With larger Fourier scale s , this derivative




Figure 4: **Stability of different formulations.** We show the norms of both terms in $\frac{d\mathbf{f}_{\theta-}}{dt} = \nabla_{\mathbf{x}} \mathbf{f}_{\theta-} \cdot \frac{d\mathbf{x}_t}{dt} + \partial_t \mathbf{f}_{\theta-}$ for diffusion models trained with the EDM ($c_{\text{noise}}(t) = \log(\sigma_d \tan(t))$) and TrigFlow ($c_{\text{noise}}(t) = t$) formulations using different time embeddings. We observe that large Fourier scales in Fourier embeddings cause instabilities. In addition, the EDM formulation suffers from numerical issues when $t \rightarrow \frac{\pi}{2}$, while TrigFlow (using positional embeddings) has stable partial derivatives for both \mathbf{x}_t and t .




Figure 5: **Comparing different training objectives for consistency distillation.** The diffusion models are EDM2 (Karras et al., 2024) pretrained on ImageNet 512×512. (a) 1-step and 2-step sampling of continuous-time CMs trained by using raw tangents $\frac{d\mathbf{f}_{\theta-}}{dt}$, clipped tangents $\text{clip}(\frac{d\mathbf{f}_{\theta-}}{dt}, -1, 1)$ and normalized tangents $(\frac{d\mathbf{f}_{\theta-}}{dt}) / (\|\frac{d\mathbf{f}_{\theta-}}{dt}\| + 0.1)$. (b) Quality of 1-step and 2-step samples from continuous-time CMs trained w/ and w/o adaptive weighting, both are w/ tangent normalization. (c) Quality of 1-step samples from continuous-time CMs vs. discrete-time CMs using varying number of time steps (N), trained using all techniques in Sec. 4.

has greater magnitudes and oscillates more vibrantly, causing worse instability. To avoid this, we use positional embeddings, which amounts to $s \approx 0.02$ in Fourier embeddings. This analysis provides a principled explanation for the observations in Song & Dhariwal (2023).

Adaptive Double Normalization. Song & Dhariwal (2023) found that the AdaGN layer (Dhariwal & Nichol, 2021), defined as $\mathbf{y} = \text{norm}(\mathbf{x}) \odot \mathbf{s}(t) + \mathbf{b}(t)$, negatively causes CM training to diverge. Our modification is *adaptive double normalization*, defined as $\mathbf{y} = \text{norm}(\mathbf{x}) \odot \text{pnorm}(\mathbf{s}(t)) + \text{pnorm}(\mathbf{b}(t))$, where $\text{pnorm}(\cdot)$ denotes pixel normalization (Karras, 2017). Empirically we find it retains the expressive power of AdaGN for diffusion training but removes its instability in CM training.

As shown in Figure 4, we visualize how our techniques stabilize the time-derivates for CMs trained on CIFAR-10. Empirically, we find that these improvements help stabilize the training dynamics of CMs without hurting diffusion model training (see Appendix G).


4.2 TRAINING OBJECTIVES

Using the TrigFlow formulation in Sec. 3 and techniques proposed in Sec. 4.1, the gradient of continuous-time CM training in Eq. (2) becomes

$$\nabla_{\theta} \mathbb{E}_{\mathbf{x}_t, t} \left[-w(t) \sigma_d \sin(t) \mathbf{F}_{\theta}^{\top} \left(\frac{\mathbf{x}_t}{\sigma_d}, t \right) \frac{d\mathbf{f}_{\theta-}(\mathbf{x}_t, t)}{dt} \right].$$

Below we propose additional techniques to explicitly control this gradient for improved stability.

Tangent Normalization. As discussed in Sec. 4.1, most gradient variance in CM training comes from the tangent function $\frac{d\mathbf{f}_{\theta-}}{dt}$. We propose to explicitly normalize the tangent function by replacing $\frac{d\mathbf{f}_{\theta-}}{dt}$ with $\frac{d\mathbf{f}_{\theta-}}{dt} / (\|\frac{d\mathbf{f}_{\theta-}}{dt}\| + c)$, where we empirically set $c = 0.1$. Alternatively, we can clip the



(a) FID (↓) as a function of single forward flops. (b) FID ratio (↓) as a function of single forward flops.

Figure 6: **sCD scales commensurately with teacher diffusion models.** We plot the (a) FID and (b) FID ratio against the teacher diffusion model (at the same model size) on ImageNet 64×64 and 512×512. sCD scales better than sCT, and has a *constant offset* in the FID ratio across all model sizes, implying that sCD has the same scaling property of the teacher diffusion model. Furthermore, the offset diminishes with more sampling steps. tangent within $[-1, 1]$, which also caps its variance. Our results in Figure 5(a) demonstrate that either normalization or clipping leads to substantial improvements for the training of continuous-time CMs.

Adaptive Weighting. Previous works (Song & Dhariwal, 2023; Geng et al., 2024) design weighting functions $w(t)$ manually for CM training, which can be suboptimal for different data distributions and network architectures. Following EDM2 (Karras et al., 2024), we propose to train an adaptive weighting function alongside the CM, which not only eases the burden of hyperparameter tuning but also outperforms manually designed weighting functions with better empirical performance and negligible training overhead. Key to our approach is the observation that $\nabla_{\theta} \mathbb{E}[\mathbf{F}_{\theta}^{\top} \mathbf{y}] = \frac{1}{2} \nabla_{\theta} \mathbb{E}[\|\mathbf{F}_{\theta} - \mathbf{F}_{\theta-} + \mathbf{y}\|_2^2]$, where \mathbf{y} is an arbitrary vector independent of θ . When training continuous-time CMs using Eq. (2), we have $\mathbf{y} = -w(t)\sigma_d \sin(t) \frac{d\mathbf{f}_{\theta-}}{dt}$. This observation allows us to convert Eq. (2) into the gradient of an MSE objective. We can therefore use the same approach in Karras et al. (2024) to train an adaptive weighting function that minimizes the variance of MSE losses across time steps (details in Appendix D). In practice, we find that integrating a prior weighting $w(t) = \frac{1}{\sigma_d \tan(t)}$ further reduces training variance. By incorporating the prior weighting, we train both the network \mathbf{F}_{θ} and the adaptive weighting function $w_{\phi}(t)$ by minimizing

$$\mathcal{L}_{\text{sCM}}(\theta, \phi) := \mathbb{E}_{\mathbf{x}_t, t} \left[\frac{e^{w_{\phi}(t)}}{D} \left\| \mathbf{F}_{\theta} \left(\frac{\mathbf{x}_t}{\sigma_d}, t \right) - \mathbf{F}_{\theta-} \left(\frac{\mathbf{x}_t}{\sigma_d}, t \right) - \cos(t) \frac{d\mathbf{f}_{\theta-}(\mathbf{x}_t, t)}{dt} \right\|_2^2 - w_{\phi}(t) \right], \quad (8)$$

where D is the dimensionality of \mathbf{x}_0 , and we sample $\tan(t)$ from a log-Normal proposal distribution (Karras et al., 2022), that is, $e^{\sigma_d \tan(t)} \sim \mathcal{N}(P_{\text{mean}}, P_{\text{std}}^2)$ (details in Appendix G).

Diffusion Finetuning and Tangent Warmup. For consistency distillation, we find that finetuning the CM from a pretrained diffusion model can speed up convergence, which is consistent with Song et al. (2023); Geng et al. (2024). Recall that in Eq. (6), the tangent $\frac{d\mathbf{f}_{\theta-}}{dt}$ can be decomposed into two parts: the first term $\cos(t)(\sigma_d \mathbf{F}_{\theta-} - \frac{d\mathbf{x}_t}{dt})$ is relatively stable, whereas the second term $\sin(t)(\mathbf{x}_t + \sigma_d \frac{d\mathbf{f}_{\theta-}}{dt})$ may cause instability. We introduce an optional technique named as *tangent warmup* by replacing the coefficient $\sin(t)$ with $r \cdot \sin(t)$, where r linearly increases from 0 to 1 over the first 10k training iterations. We find that the tangent normalization does not affect sample quality but may reduce some gradient spikes during training.

With all techniques in place, the stability of both discrete-time and continuous-time CM training substantially improves. We provide detailed algorithms for discrete-time CMs in Appendix E, and train continuous-time CMs and discrete-time CMs with the same setting. As demonstrated in Figure 5(c), increasing the number of discretization steps N in discrete-time CMs improves sample quality by reducing discretization errors, but degrades once N becomes too large (after $N > 1024$) to suffer from numerical precision issues. By contrast, continuous-time CMs significantly outperform discrete-time CMs across all N 's which provides strong justification for choosing continuous-time CMs over discrete-time counterparts. We call our model **sCM** (s for *simple, stable, and scalable*), and provide detailed pseudo-code for sCM training in Appendix A.

5 SCALING UP CONTINUOUS-TIME CONSISTENCY MODELS

Below we test all the improvements proposed in previous sections by training large-scale sCMs on a variety of challenging datasets.




Figure 7: **sCD has higher diversity compared to VSD**: Sample quality comparison of the EDM2 (Karras et al., 2024) diffusion model, VSD (Wang et al., 2024; Yin et al., 2024b), sCD, and the combination of VSD and sCD, across varying guidance scales. All models are of EDM2-M size and trained on ImageNet 512×512 .

Table 1: Sample quality on unconditional CIFAR-10 and class-conditional ImageNet 64×64 .

Unconditional CIFAR-10			Class-Conditional ImageNet 64×64		
METHOD	NFE (↓)	FID (↓)	METHOD	NFE (↓)	FID (↓)
Diffusion models & Fast Samplers			Diffusion models & Fast Samplers		
Score SDE (deep) (Song et al., 2021b)	2000	2.20	ADM (Dhariwal & Nichol, 2021)	250	2.07
EDM (Karras et al., 2022)	35	2.01	RIN (Jabri et al., 2022)	1000	1.23
Flow Matching (Lipman et al., 2022)	142	6.35	DPM-Solver (Lu et al., 2022a)	20	3.42
OT-CFM (Tong et al., 2023)	1000	3.57	EDM (Heun) (Karras et al., 2022)	79	2.44
DPM-Solver (Lu et al., 2022a)	10	4.70	EDM2 (Heun) (Karras et al., 2024)	63	1.33
DPM-Solver++ (Lu et al., 2022b)	10	2.91	Joint Training		
DPM-Solver-v3 (Zheng et al., 2023c)	10	2.51	StyleGAN-XL (Sauer et al., 2022)	1	1.52
Joint Training			Diff-Instruct (Luo et al., 2024)	1	5.57
Diffusion GAN (Xiao et al., 2022)	4	3.75	EMD (Xie et al., 2024b)	1	2.20
Diffusion StyleGAN (Wang et al., 2022)	1	3.19	DMD (Yin et al., 2024b)	1	2.62
StyleGAN-XL (Sauer et al., 2022)	1	1.52	DMD2 (Yin et al., 2024a)	1	1.28
CTM (Kim et al., 2023)	1	1.87	SiD (Zhou et al., 2024)	1	1.52
Diff-Instruct (Luo et al., 2024)	1	4.53	CTM (Kim et al., 2023)	1	1.92
DMD (Yin et al., 2024b)	1	3.77	Moment Matching (Salimans et al., 2024)	1	3.00
SiD (Zhou et al., 2024)	1	1.92		2	3.86
Diffusion Distillation			Diffusion Distillation		
DFNO (LPIPS) (Zheng et al., 2023b)	1	3.78	DFNO (LPIPS) (Zheng et al., 2023b)	1	7.83
2-Rectified Flow (Liu et al., 2022)	1	4.85	PID (LPIPS) (Tee et al., 2024)	1	9.49
PID (LPIPS) (Tee et al., 2024)	1	3.92	TRACT (Berthelot et al., 2023)	1	7.43
Consistency-FM (Yang et al., 2024)	2	5.34		2	4.97
PD (Salimans & Ho, 2022)	1	8.34	PD (Salimans & Ho, 2022)	1	10.70
	2	5.58	(reimpl. from Heek et al. (2024))	2	4.70
TRACT (Berthelot et al., 2023)	1	3.78	CD (LPIPS) (Song et al., 2023)	1	6.20
	2	3.32		2	4.70
CD (LPIPS) (Song et al., 2023)	1	3.55	MultiStep-CD (Heek et al., 2024)	1	3.20
	2	2.93		2	1.90
sCD (ours)	1	3.66	sCD (ours)	1	2.44
	2	2.52		2	1.66
Consistency Training			Consistency Training		
iCT (Song & Dhariwal, 2023)	1	2.83	iCT (Song & Dhariwal, 2023)	1	4.02
	2	2.46		2	3.20
iCT-deep (Song & Dhariwal, 2023)	1	2.51	iCT-deep (Song & Dhariwal, 2023)	1	3.25
	2	2.24		2	2.77
ECT (Geng et al., 2024)	1	3.60	ECT (Geng et al., 2024)	1	2.49
	2	2.11		2	1.67
sCT (ours)	1	2.85	sCT (ours)	1	2.04
	2	2.06		2	1.48

5.1 TANGENT COMPUTATION IN LARGE-SCALE MODELS

The common setting for training large-scale diffusion models includes using half-precision (FP16) and Flash Attention (Dao et al., 2022; Dao, 2023). As training continuous-time CMs requires computing the tangent $\frac{d\mathbf{f}_{\theta-}}{dt}$ accurately, we need to improve numerical precision and also support memory-efficient attention computation, as detailed below.

JVP Rearrangement. Computing $\frac{d\mathbf{f}_{\theta-}}{dt}$ involves calculating $\frac{d\mathbf{f}_{\theta-}}{dt} = \nabla_{\mathbf{x}_t} \mathbf{F}_{\theta-} \cdot \frac{d\mathbf{x}_t}{dt} + \partial_t \mathbf{F}_{\theta-}$, which can be efficiently obtained via the Jacobian-vector product (JVP) for $\mathbf{F}_{\theta-}(\cdot, \frac{\cdot}{\sigma_d}, \cdot)$ with the input

Table 2: Sample quality on class-conditional ImageNet 512×512 . [†]Our reimplemented teacher diffusion model based on EDM2 (Karras et al., 2024) but with modifications in Sec. 4.1.

METHOD	NFE (\downarrow)	FID (\downarrow)	#Params	METHOD	NFE (\downarrow)	FID (\downarrow)	#Params
Diffusion models				[†]Teacher Diffusion Model			
ADM-G (Dhariwal & Nichol, 2021)	250×2	7.72	559M	EDM2-S (Karras et al., 2024)	63×2	2.29	280M
RIN (Jabri et al., 2022)	1000	3.95	320M	EDM2-M (Karras et al., 2024)	63×2	2.00	498M
U-ViT-H/4 (Bao et al., 2023)	250×2	4.05	501M	EDM2-L (Karras et al., 2024)	63×2	1.87	778M
DiT-XL/2 (Peebles & Xie, 2023)	250×2	3.04	675M	EDM2-XL (Karras et al., 2024)	63×2	1.80	1.1B
SimDiff (Hoogeboom et al., 2023)	512×2	3.02	2B	EDM2-XXL (Karras et al., 2024)	63×2	1.73	1.5B
VDM++ (Kingma & Gao, 2024)	512×2	2.65	2B	Consistency Training (sCT, ours)			
DiffiT (Hatamizadeh et al., 2023)	250×2	2.67	561M	sCT-S (ours)	1	10.13	280M
DiMR-XL/3R (Liu et al., 2024)	250×2	2.89	525M		2	9.86	280M
DiffuSSM-XL (Yan et al., 2024)	250×2	3.41	673M	sCT-M (ours)	1	5.84	498M
DiM-H (Teng et al., 2024)	250×2	3.78	860M		2	5.53	498M
U-DiT (Tian et al., 2024b)	250	15.39	204M	sCT-L (ours)	1	5.15	778M
SiT-XL (Ma et al., 2024)	250×2	2.62	675M		2	4.65	778M
Large-DiT (Alpha-VLLM, 2024)	250×2	2.52	3B	sCT-XL (ours)	1	4.33	1.1B
MaskDiT (Zheng et al., 2023a)	79×2	2.50	736M		2	3.73	1.1B
DiS-H/2 (Fei et al., 2024a)	250×2	2.88	900M	sCT-XXL (ours)	1	4.29	1.5B
DRWKV-H/2 (Fei et al., 2024b)	250×2	2.95	779M		2	3.76	1.5B
EDM2-S (Karras et al., 2024)	63×2	2.23	280M	Consistency Distillation (sCD, ours)			
EDM2-M (Karras et al., 2024)	63×2	2.01	498M	sCD-S	1	3.07	280M
EDM2-L (Karras et al., 2024)	63×2	1.88	778M		2	2.50	280M
EDM2-XL (Karras et al., 2024)	63×2	1.85	1.1B	sCD-M	1	2.75	498M
EDM2-XXL (Karras et al., 2024)	63×2	1.81	1.5B		2	2.26	498M
GANs & Masked Models				sCD-L	1	2.55	778M
BigGAN (Brock, 2018)	1	8.43	160M		2	2.04	778M
StyleGAN-XL (Sauer et al., 2022)	1×2	2.41	168M	sCD-XL	1	2.40	1.1B
VQGAN (Esser et al., 2021)	1024	26.52	227M		2	1.93	1.1B
MaskGIT (Chang et al., 2022)	12	7.32	227M	sCD-XXL	1	2.28	1.5B
MAGViT-v2 (Yu et al., 2023)	64×2	1.91	307M		2	1.88	1.5B
MAR (Li et al., 2024)	64×2	1.73	481M				
VAR-d36-s (Tian et al., 2024a)	10×2	2.63	2.3B				

vector (\mathbf{x}_t, t) and the tangent vector $(\frac{d\mathbf{x}_t}{dt}, 1)$. However, we empirically find that the tangent may overflow in intermediate layers when t is near 0 or $\frac{\pi}{2}$. To improve numerical precision, we propose to rearrange the computation of the tangent. Specifically, since the objective in Eq. (8) contains $\cos(t)\frac{d\mathbf{f}_{\theta-}}{dt}$ and $\frac{d\mathbf{f}_{\theta-}}{dt}$ is proportional to $\sin(t)\frac{d\mathbf{F}_{\theta-}}{dt}$, we can compute the JVP as:

$$\cos(t)\sin(t)\frac{d\mathbf{F}_{\theta-}}{dt} = \left(\nabla_{\frac{\mathbf{x}_t}{\sigma_d}}\mathbf{F}_{\theta-}\right) \cdot \left(\cos(t)\sin(t)\frac{d\mathbf{x}_t}{dt}\right) + \partial_t\mathbf{F}_{\theta-} \cdot (\cos(t)\sin(t)\sigma_d),$$

which is the JVP for $\mathbf{F}_{\theta-}(\cdot, \cdot)$ with the input $(\frac{\mathbf{x}_t}{\sigma_d}, t)$ and the tangent $(\cos(t)\sin(t)\frac{d\mathbf{x}_t}{dt}, \cos(t)\sin(t)\sigma_d)$. This rearrangement greatly alleviates the overflow issues in the intermediate layers, resulting in more stable training in FP16.

JVP of Flash Attention. Flash Attention (Dao et al., 2022; Dao, 2023) is widely used for attention computation in large-scale model training, providing both GPU memory savings and faster training. However, Flash Attention does not compute the Jacobian-vector product (JVP). To fill this gap, we propose a similar algorithm (detailed in Appendix F) that efficiently computes both softmax self-attention and its JVP in a single forward pass in the style of Flash Attention, significantly reducing GPU memory usage for JVP computation in attention layers.

5.2 EXPERIMENTS

To test our improvements, we employ both consistency training (referred to as **sCT**) and consistency distillation (referred to as **sCD**) to train and scale continuous-time CMs on CIFAR-10 (Krizhevsky, 2009), ImageNet 64×64 and ImageNet 512×512 (Deng et al., 2009). We benchmark the sample quality using FID (Heusel et al., 2017). We follow the settings of Score SDE (Song et al., 2021b) on CIFAR10 and EDM2 (Karras et al., 2024) on both ImageNet 64×64 and ImageNet 512×512 , while changing the parameterization and architecture according to Section 4.1. We adopt the method proposed by Song et al. (2023) for two-step sampling of both sCT and sCD, using a fixed intermediate time step $t = 1.1$. For sCD models on ImageNet 512×512 , since the teacher diffusion model relies on classifier-free guidance (CFG) (Ho & Salimans, 2021), we incorporate an additional input s into the model \mathbf{F}_{θ} to represent the guidance scale (Meng et al., 2023). We train the model with sCD

by uniformly sampling $s \in [1, 2]$ and applying the corresponding CFG to the teacher model during distillation (more details are provided in Appendix G). For sCT models, we do not test CFG since it is incompatible with consistency training.

Training compute of sCM. We use the same batch size as the teacher diffusion model across all datasets. The effective compute per training iteration of sCD is approximately twice that of the teacher model. We observe that the quality of two-step samples from sCD converges rapidly, achieving results comparable to the teacher diffusion model using less than 20% of the teacher training compute. In practice, we can obtain high-quality samples after only 20k finetuning iterations with sCD.

Benchmarks. In Tables 1 and 2, we compare our results with previous methods by benchmarking the FIDs and the number of function evaluations (NFEs). First, sCM outperforms all previous few-step methods that do not rely on joint training with another network and is on par with, or even exceeds, the best results achieved with adversarial training. Notably, the 1-step FID of sCD-XXL on ImageNet 512×512 surpasses that of StyleGAN-XL (Sauer et al., 2022) and VAR (Tian et al., 2024a). Furthermore, the two-step FID of sCD-XXL outperforms all generative models except diffusion and is comparable with the best diffusion models that require 63 sequential steps. Second, the two-step sCM model significantly narrows the FID gap with the teacher diffusion model to within 10%, achieving FIDs of 2.06 on CIFAR-10 (compared to the teacher FID of 2.01), 1.48 on ImageNet 64×64 (teacher FID of 1.33), and 1.88 on ImageNet 512×512 (teacher FID of 1.73). Additionally, we observe that sCT is more effective at smaller scales but suffers from increased variance at larger scales, while sCD shows consistent performance across both small and large scales.

Scaling study. Based on our improved training techniques, we successfully scale continuous-time CMs without training instability. We train various sizes of sCMs using EDM2 configurations (S, M, L, XL, XXL) on ImageNet 64×64 and 512×512 , and evaluate FID under optimal guidance scales, as shown in Fig. 6. First, as model FLOPs increase, both sCT and sCD show improved sample quality, showing that both methods benefit from scaling. Second, compared to sCD, sCT is more compute efficient at smaller resolutions but less efficient at larger resolutions. Third, sCD scales predictably for a given dataset, maintaining a consistent relative difference in FIDs across model sizes. This suggests that the FID of sCD decreases at the same rate as the teacher diffusion model, and therefore *sCD is as scalable as the teacher diffusion model*. As the FID of the teacher diffusion model decreases with scaling, the *absolute* difference in FID between sCD and the teacher model also diminishes. Finally, the relative difference in FIDs decreases with more sampling steps, and the sample quality of the two-step sCD becomes on par with that of the teacher diffusion model.

Comparison with VSD. Variational score distillation (VSD) (Wang et al., 2024; Yin et al., 2024b) and its multi-step generalization (Xie et al., 2024b; Salimans et al., 2024) represent another diffusion distillation technique that has demonstrated scalability on high-resolution images (Yin et al., 2024a). We apply one-step VSD from time T to 0 to finetune a teacher diffusion model using the EDM2-M configuration and tune both the weighting functions and proposal distributions for fair comparisons. As shown in Figure 7, we compare sCD, VSD, a combination of sCD and VSD (by simply adding the two losses), and the teacher diffusion model by sweeping over the guidance scale. We observe that VSD has artifacts similar to those from applying large guidance scales in diffusion models: it increases fidelity (as evidenced by higher precision scores) while decreasing diversity (as shown by lower recall scores). This effect becomes more pronounced with increased guidance scales, ultimately causing severe mode collapse. In contrast, the precision and recall scores from two-step sCD are comparable with those of the teacher diffusion model, resulting in better FID scores than VSD.

6 CONCLUSION

Our improved formulations, architectures, and training objectives have simplified and stabilized the training of continuous-time consistency models, enabling smooth scaling up to 1.5 billion parameters on ImageNet 512×512 . We ablated the impact of TrigFlow formulation, tangent normalization, and adaptive weighting, confirming their effectiveness. Combining these improvements, our method demonstrated predictable scalability across datasets and model sizes, outperforming other few-step sampling approaches at large scales. Notably, we narrowed the FID gap with the teacher model to within 10% using two-step generation, compared to state-of-the-art diffusion models that require significantly more sampling steps.

REFERENCES

- Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.
- Michael Samuel Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=li7qeBbCRlt>.
- Alpha-VLLM. Large-DiT-ImageNet. <https://github.com/Alpha-VLLM/LLaMA2-Accessory/tree/f7fe19834b23e38f333403b91bb0330afe19f79e/Large-DiT-ImageNet>, 2024. Commit f7fe198.
- Fan Bao, Shen Nie, Kaiwen Xue, Yue Cao, Chongxuan Li, Hang Su, and Jun Zhu. All are worth words: A vit backbone for diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 22669–22679, 2023.
- David Berthelot, Arnaud Autef, Jierui Lin, Dian Ang Yap, Shuangfei Zhai, Siyuan Hu, Daniel Zheng, Walter Talbott, and Eric Gu. Tract: Denoising diffusion models with transitive closure time-distillation. *arXiv preprint arXiv:2303.04248*, 2023.
- Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.
- Andrew Brock. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. Video generation models as world simulators. 2024. URL <https://openai.com/research/video-generation-models-as-world-simulators>.
- Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11315–11325, 2022.
- Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning (2023). *arXiv preprint arXiv:2307.08691*, 2023.
- Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems*, 35:16344–16359, 2022.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255. IEEE, 2009.
- Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems*, volume 34, pp. 8780–8794, 2021.
- Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12873–12883, 2021.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024.
- Zach Evans, CJ Carr, Josiah Taylor, Scott H Hawley, and Jordi Pons. Fast timing-conditioned latent audio diffusion. *arXiv preprint arXiv:2402.04825*, 2024.

- Zhengcong Fei, Mingyuan Fan, Changqian Yu, and Junshi Huang. Scalable diffusion models with state space backbone. *arXiv preprint arXiv:2402.05608*, 2024a.
- Zhengcong Fei, Mingyuan Fan, Changqian Yu, Debang Li, and Junshi Huang. Diffusion-rwkv: Scaling rwkv-like architectures for diffusion models. *arXiv preprint arXiv:2404.04478*, 2024b.
- Zhengyang Geng, Ashwini Pokle, William Luo, Justin Lin, and J Zico Kolter. Consistency models made easy. *arXiv preprint arXiv:2406.14548*, 2024.
- Ali Hatamizadeh, Jiaming Song, Guilin Liu, Jan Kautz, and Arash Vahdat. Diffit: Diffusion vision transformers for image generation. *arXiv preprint arXiv:2312.02139*, 2023.
- Jonathan Heek, Emiel Hooeboom, and Tim Salimans. Multistep consistency models. *arXiv preprint arXiv:2403.06807*, 2024.
- Eric Heitz, Laurent Belcour, and Thomas Chambon. Iterative α -(de) blending: A minimalist deterministic diffusion model. In *ACM SIGGRAPH 2023 Conference Proceedings*, pp. 1–8, 2023.
- 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 Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30, pp. 6626–6637, 2017.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851, 2020.
- Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022.
- Emiel Hooeboom, Jonathan Heek, and Tim Salimans. simple diffusion: End-to-end diffusion for high resolution images. In *International Conference on Machine Learning*, pp. 13213–13232. PMLR, 2023.
- Allan Jabri, David Fleet, and Ting Chen. Scalable adaptive computation for iterative generation. *arXiv preprint arXiv:2212.11972*, 2022.
- Michael Janner, Yilun Du, Joshua Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. In *International Conference on Machine Learning*, 2022.
- Tero Karras. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *arXiv preprint arXiv:2206.00364*, 2022.
- Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. Analyzing and improving the training dynamics of diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 24174–24184, 2024.
- Dongjun Kim, Chieh-Hsin Lai, Wei-Hsiang Liao, Naoki Murata, Yuhta Takida, Toshimitsu Ue-saka, Yutong He, Yuki Mitsufuji, and Stefano Ermon. Consistency trajectory models: Learning probability flow ode trajectory of diffusion. *arXiv preprint arXiv:2310.02279*, 2023.
- Diederik Kingma and Ruiqi Gao. Understanding diffusion objectives as the elbo with simple data augmentation. *Advances in Neural Information Processing Systems*, 36, 2024.
- Diederik P Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. In *Advances in Neural Information Processing Systems*, 2021.

- Nikita Kornilov, Petr Mokrov, Alexander Gasnikov, and Alexander Korotin. Optimal flow matching: Learning straight trajectories in just one step. *arXiv preprint arXiv:2403.13117*, 2024.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- Tianhong Li, Yonglong Tian, He Li, Mingyang Deng, and Kaiming He. Autoregressive image generation without vector quantization. *arXiv preprint arXiv:2406.11838*, 2024.
- Shanchuan Lin, Bingchen Liu, Jiashi Li, and Xiao Yang. Common diffusion noise schedules and sample steps are flawed. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 5404–5411, 2024.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- 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 preprint arXiv:2301.12503*, 2023a.
- Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265*, 2019.
- Qihao Liu, Zhanpeng Zeng, Ju He, Qihang Yu, Xiaohui Shen, and Liang-Chieh Chen. Alleviating distortion in image generation via multi-resolution diffusion models. *arXiv preprint arXiv:2406.09416*, 2024.
- Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. 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, 2023b.
- Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35:5775–5787, 2022a.
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*, 2022b.
- Eric Luhman and Troy Luhman. Knowledge distillation in iterative generative models for improved sampling speed. *arXiv preprint arXiv:2101.02388*, 2021.
- Simian Luo, Yiqin Tan, Longbo Huang, Jian Li, and Hang Zhao. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*, 2023.
- Weijian Luo, Tianyang Hu, Shifeng Zhang, Jiacheng Sun, Zhenguo Li, and Zhihua Zhang. Diff-instruct: A universal approach for transferring knowledge from pre-trained diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Nanye Ma, Mark Goldstein, Michael S Albergo, Nicholas M Boffi, Eric Vanden-Eijnden, and Saining Xie. Sit: Exploring flow and diffusion-based generative models with scalable interpolant transformers. *arXiv preprint arXiv:2401.08740*, 2024.
- Chenlin Meng, Robin Rombach, Ruiqi Gao, Diederik Kingma, Stefano Ermon, Jonathan Ho, and Tim Salimans. On distillation of guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14297–14306, 2023.
- Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pp. 8162–8171. PMLR, 2021.

- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023.
- Stefano Peluchetti. Non-denoising forward-time diffusions, 2022. URL <https://openreview.net/forum?id=oVfIKuhqfC>.
- Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with CLIP latents. *arXiv preprint arXiv:2204.06125*, 2022.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10684–10695, 2022.
- Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2022.
- Tim Salimans, Thomas Mensink, Jonathan Heek, and Emiel Hoogeboom. Multistep distillation of diffusion models via moment matching. *arXiv preprint arXiv:2406.04103*, 2024.
- Axel Sauer, Katja Schwarz, and Andreas Geiger. Stylegan-xl: Scaling stylegan to large diverse datasets. In *ACM SIGGRAPH 2022 conference proceedings*, pp. 1–10, 2022.
- Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion distillation. *arXiv preprint arXiv:2311.17042*, 2023.
- Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. *Advances in neural information processing systems*, 31, 2018.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pp. 2256–2265. PMLR, 2015.
- Yang Song and Prafulla Dhariwal. Improved techniques for training consistency models. *arXiv preprint arXiv:2310.14189*, 2023.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems*, volume 32, pp. 11895–11907, 2019.
- Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. In *Advances in Neural Information Processing Systems*, volume 34, pp. 1415–1428, 2021a.
- Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021b.
- Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *International Conference on Machine Learning*, pp. 32211–32252. PMLR, 2023.
- Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in neural information processing systems*, 33:7537–7547, 2020.
- Joshua Tian Jin Tee, Kang Zhang, Hee Suk Yoon, Dhananjaya Nagaraja Gowda, Chanwoo Kim, and Chang D Yoo. Physics informed distillation for diffusion models. *Transactions on Machine Learning Research*, 2024.

- Yao Teng, Yue Wu, Han Shi, Xuefei Ning, Guohao Dai, Yu Wang, Zhenguo Li, and Xihui Liu. Dim: Diffusion mamba for efficient high-resolution image synthesis. *arXiv preprint arXiv:2405.14224*, 2024.
- Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *arXiv preprint arXiv:2404.02905*, 2024a.
- Yuchuan Tian, Zhijun Tu, Hanting Chen, Jie Hu, Chao Xu, and Yunhe Wang. U-dits: Downsample tokens in u-shaped diffusion transformers. *arXiv preprint arXiv:2405.02730*, 2024b.
- Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Hugué, Yanlei Zhang, Jarrod Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023.
- A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- Zhendong Wang, Huangjie Zheng, Pengcheng He, Weizhu Chen, and Mingyuan Zhou. Diffusion-gan: Training gans with diffusion. *arXiv preprint arXiv:2206.02262*, 2022.
- Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Pro-lificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *Advances in Neural Information Processing Systems*, 36, 2024.
- Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with denoising diffusion GANs. In *International Conference on Learning Representations*, 2022.
- Qingsong Xie, Zhenyi Liao, Zhijie Deng, Shixiang Tang, Haonan Lu, et al. Mlcm: Multistep consistency distillation of latent diffusion model. *arXiv preprint arXiv:2406.05768*, 2024a.
- Sirui Xie, Zhisheng Xiao, Diederik P Kingma, Tingbo Hou, Ying Nian Wu, Kevin Patrick Murphy, Tim Salimans, Ben Poole, and Ruiqi Gao. EM distillation for one-step diffusion models. *arXiv preprint arXiv:2405.16852*, 2024b.
- Jing Nathan Yan, Jiatao Gu, and Alexander M Rush. Diffusion models without attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8239–8249, 2024.
- Ling Yang, Zixiang Zhang, Zhilong Zhang, Xingchao Liu, Minkai Xu, Wentao Zhang, Chenlin Meng, Stefano Ermon, and Bin Cui. Consistency flow matching: Defining straight flows with velocity consistency. *arXiv preprint arXiv:2407.02398*, 2024.
- Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and William T Freeman. Improved distribution matching distillation for fast image synthesis. *arXiv preprint arXiv:2405.14867*, 2024a.
- Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6613–6623, 2024b.
- Lijun Yu, José Lezama, Nitesh B Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G Hauptmann, et al. Language model beats diffusion-tokenizer is key to visual generation. *arXiv preprint arXiv:2310.05737*, 2023.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 586–595, 2018.
- Hongkai Zheng, Weili Nie, Arash Vahdat, and Anima Anandkumar. Fast training of diffusion models with masked transformers. *arXiv preprint arXiv:2306.09305*, 2023a.
- Hongkai Zheng, Weili Nie, Arash Vahdat, Kamyar Azizzadenesheli, and Anima Anandkumar. Fast sampling of diffusion models via operator learning. In *International conference on machine learning*, pp. 42390–42402. PMLR, 2023b.

Kaiwen Zheng, Cheng Lu, Jianfei Chen, and Jun Zhu. Dpm-solver-v3: Improved diffusion ode solver with empirical model statistics. *Advances in Neural Information Processing Systems*, 36: 55502–55542, 2023c.

Mingyuan Zhou, Huangjie Zheng, Zhendong Wang, Mingzhang Yin, and Hai Huang. Score identity distillation: Exponentially fast distillation of pretrained diffusion models for one-step generation. In *Forty-first International Conference on Machine Learning*, 2024.