

Understanding, Accelerating, and Improving MeanFlow Training

Jin-Young Kim* Hyojun Go^{1*} Lea Bogensperger² Julius Erbach^{1,3}
 Nikolai Kalischek⁴ Federico Tombari⁴ Konrad Schindler¹ Dominik Narnhofer¹

¹ETH Zurich ²University of Zurich ³Max Planck ETH Center for Learning Systems ⁴Google

Abstract

MeanFlow promises high-quality generative modeling in few steps, by jointly learning instantaneous and average velocity fields. Yet, the underlying training dynamics remain unclear. We analyze the interaction between the two velocities and find: (i) well-established instantaneous velocity is a prerequisite for learning average velocity; (ii) learning of instantaneous velocity benefits from average velocity when the temporal gap is small, but degrades as the gap increases; and (iii) task-affinity analysis indicates that smooth learning of large-gap average velocities, essential for one-step generation, depends on the prior formation of accurate instantaneous and small-gap average velocities. Guided by these observations, we design an effective training scheme that accelerates the formation of instantaneous velocity, then shifts emphasis from short- to long-interval average velocity. Our enhanced MeanFlow training yields faster convergence and significantly better few-step generation: With the same DiT-XL backbone, our method reaches an impressive FID of 2.87 on 1-NFE ImageNet 256×256, compared to 3.43 for the conventional MeanFlow baseline. Alternatively, our method matches the performance of the MeanFlow baseline with 2.5× shorter training time, or with a smaller DiT-L backbone. Our code is available at <https://github.com/seah0119/ImprovedMeanFlow>.

1. Introduction

Diffusion models [28, 61, 64] and Flow Matching [2, 39, 40] have achieved state-of-the-art results across image [4, 48, 75], video [1, 17, 70], and 3D generation [21, 22]. However, it remains a persistent weakness that the denoising relies on many, small iteration steps, making sampling computationally expensive [25]. Higher-order samplers [13, 32, 43, 44, 54, 62, 80] partially alleviate this, though achieving high fidelity with fewer than 10 steps remains a challenge. Consequently, recent work has focused on models that enable inference in a few steps, or even a

*Equal contribution




Figure 1. Our novel, enhanced training strategy reaches the performance of MeanFlow-XL in $\approx 2.5\times$ fewer training epochs, and converges to a final model with superior performance ($\approx 16\%$ lower FID).

single step.

Early approaches distill few-step generative models from pretrained multi-step diffusion models, using direct [45, 56, 81], adversarial [59, 60, 76], or score-based supervision [47, 77, 86]. The two-stage design increases complexity, requires two distinct training processes, and often depends on large-scale synthetic data generation [40, 45], or on propagation through teacher–student cascades [49, 56].

Consistency models [66] represent a step towards one-time, end-to-end training, by enforcing consistent outputs for all samples drawn along the same denoising trajectory. Despite various improvements [19, 26, 34, 42, 63, 71, 72], a substantial performance gap remains between few-step consistency models and multi-step diffusion models. More recent research [6, 7, 16, 34, 55, 74, 85] has proposed to char-

acterize diffusion/flow quantities along two distinct time indices. Among these attempts, MeanFlow [18] stands out as a stable end-to-end training scheme that markedly narrows the gap between one-step and multi-step generation.

The key to MeanFlow’s success is the idea of exploiting the intrinsic relationship between the *instantaneous velocity* (at a single time point) and the *average velocity* (integrated over a time interval), such that a single network learns both simultaneously. However, MeanFlow training is computationally expensive, and has only been analyzed rather superficially. In particular, it remains poorly understood how the two coupled velocity fields interact during learning and how their interplay can be coordinated to achieve high-quality one-step generation.

Here, we investigate these learning dynamics and develop a training strategy that greatly improves both generation quality and efficiency. Through controlled experiments, we determine that: (i) instantaneous velocity must be established early in the training process, because it provides the foundation for learning average velocity: if instantaneous velocities are poorly formed or corrupted, then learning average velocities fails altogether; (ii) in the opposite direction, the time interval over which average velocities are computed (the “temporal gap”) critically determines how they impact the learning of instantaneous velocities: small gaps facilitate instantaneous velocity formation and refinement, while large gaps destabilize it; (iii) task affinity analysis reveals that one should initially focus on small-gap average velocities, which lay the foundation for learning the large-gap average velocities that are required for one-step generation.

Standard MeanFlow training ignores these subtle, but impactful dynamics. Throughout the training process, it applies the same, fixed loss function and sampling scheme, disregarding the complex dependencies between the two velocity fields. This naive training objective interferes with the early formation of reasonable instantaneous velocities, which in turn delays the learning of average velocities. Ultimately, the current training practice significantly degrades overall performance compared to what would be achievable with a given model and dataset, and also slows down the training.

To remedy these issues, we propose a simple yet effective extension of MeanFlow training. To quickly establish reasonable instantaneous velocities, we adopt acceleration techniques from diffusion training [10, 20, 25, 35, 73]. To support the learning of correct average velocities, we design a progressive weighting scheme. In early training stages, the weighting prioritizes small gaps, which reinforces instantaneous velocity formation and prepares the ground for large-gap learning. As training progresses, the weighting gradually transitions to equal weighting across all gap lengths, ensuring accurate average velocities over

large gaps, which are the vital ingredient for few-step inference.

Empirically, the enhanced training protocol substantially improves the generation results and also accelerates convergence. On the standard 1-NFE ImageNet [11] 256×256 benchmark, we improve the FID of MeanFlow-XL from FID 3.43 to 2.87, see Fig. 1. To reach the performance of conventional training, our improved training scheme needs $2.5\times$ fewer iterations. Remarkably, it is even capable of matching that same performance with a smaller DiT-L backbone.

In a broader context, our work shows that there is still a lot of untapped potential to accelerate recent few-step generative models. With a better understanding of their internal dynamics and numerical properties, high-quality real-time generation may well be achievable.

2. Related Work

Acceleration of diffusion and flow matching training. Diffusion models [28, 61, 64] gradually perturb data with noise and train a network to reconstruct the clean signal. This noise-injection and denoising process can be described via stochastic differential equations (SDEs) or, equivalently, as a probability-flow ordinary differential equation (ODE) [32, 65]. Flow Matching [2, 39, 40] extends this with velocity fields, enabling the model to learn transport paths connecting data and reference distributions.

Training these models is computationally expensive [52]. To accelerate it, prior work mostly focuses on critical timesteps using two strategies: (1) timestep-dependent loss weighting based on SNR [25], perceptual quality [10], or uncertainty [20, 33], and (2) modified sampling distributions [82–84]. Furthermore, there are hybrid approaches that combine the two [73], as well as adaptive schedulers [35].

We will leverage some of the existing acceleration techniques to speed up the formation of instantaneous velocity, which is important to speed up convergence, see Sec. 4.

Few-step generative models. Early work on few-step models revolves around distillation [45], progressive [5, 56], data-free approaches [23, 49, 87], and various forms of alternative supervision, *e.g.* adversarial [59, 60, 76], score-based [47, 77, 86], moment matching [57], operator learning [81], and physics-informed losses [68]. These are not end-to-end pipelines; they are two-stage (or tightly scheduled) procedures that require manually selected transition points—*i.e.*, deciding when to stop teacher training and begin distillation [16].

Consistency models [66] pioneered a major advance, enabling *end-to-end* training of few-step generators, and they have also been used in distillation setups [38, 41, 78]. The core principle is to enforce that model outputs from

any two points along the same trajectory are identical. This foundational concept has inspired a wide range of follow-ups, such as improving training stability and simplicity [19, 42, 63, 71], extending the framework to multi-step [26, 34, 72], adapting to latent space models [46], and incorporating adversarial loss [34, 36]. Despite these advances, a substantial gap remains between the few-step performance of end-to-end training and the performance of fully multi-step diffusion models.

More recently, several works have proposed to learn diffusion and flow quantities between two time points [6, 7, 16, 34, 55, 74, 85]. For example, Flow Maps [6] define the integral of the flow between time points and learn it via matching losses. Shortcut Models [16] augment flow matching with a regularization loss to learn shortcut paths, and Inductive Moment Matching (IMM) [85] enforces self-consistency of stochastic interpolants across time.

Among the methods developed so far, MeanFlow [18] stands out: it substantially narrows the performance gap between few-step and full multi-step diffusion models, while being trained end-to-end. The aim of the present paper is to understand, improve, and accelerate the training of MeanFlow. There are a few concurrent efforts: AlphaFlow [79] replaces the original MeanFlow objective with a softened one. In contrast, we retain the MeanFlow formulation and improve it based on a careful empirical analysis. CMT [30] splits learning into multiple stages, whereas our method preserves the simplicity of end-to-end MeanFlow training.

3. Background: MeanFlow

Flow matching. Flow matching [2, 39, 40] learns a time-dependent vector field $v_\theta(z_t, t)$ that transports a (typically Gaussian) source distribution $\epsilon \sim p_1(\epsilon)$ to a target data distribution $x \sim p_0(x)$. This transformation is defined as the solution to the ODE that characterizes the flow Φ

$$\frac{d}{dt} \Phi_t(z) = v_t(\Phi_t(z)). \quad (1)$$

A valid velocity field can be learned by optimizing the conditional flow matching objective, which utilizes a tractable conditional velocity $v_t(z_t|\epsilon)$ instead of the intractable marginal velocity $v_t(z_t)$ [39]. Different choices for the conditional flow paths are possible, arguably, the simplest and most popular ones are *optimal transport conditional* velocity fields. Given a pair (x, ϵ) and a time $t \in [0, 1]$, z_t is the linear interpolant $z_t = (1-t)x + t\epsilon$, whose conditional velocity [39] is the time derivative $v_t(z_t|\epsilon) = \dot{z}_t = \epsilon - x$. The neural network v_θ is by minimizing the conditional flow-matching loss $\mathcal{L}_{CFM} = \mathbb{E}_{x,\epsilon,t} [\|v_\theta(z_t, t) - v_t(z_t|\epsilon)\|_2^2]$. Data samples are generated by solving the probability-flow ODE in Eq. 1.

MeanFlow. The core idea of MeanFlow [18] is to interpret the flow-matching velocity as the **instantaneous ve-**

locity (v) at time t and to learn an **average velocity** (u) between time points r and t defined by the MeanFlow identity

$$u(z_t, r, t) \triangleq \frac{1}{t-r} \int_r^t v_t(z_\tau, \tau) d\tau. \quad (2)$$

By learning u with a neural network u_θ , MeanFlow approximates the finite-time ODE integral in Eq. 1, enabling a single update step $z_r = z_t - (t-r)u_\theta(z_t, r, t)$ that replaces multiple small solver steps. To train u_θ , MeanFlow exploits the relationship between u and v , which is defined as:

$$u(z_t, t, r) = v_t(z_t, t) - (t-r)(v_t(z_t, t)\partial_z u_\theta + \partial_t u_\theta). \quad (3)$$

The overall objective is then given by:

$$\mathcal{L}_{MF} = \mathbb{E}_{x,\epsilon,t,r} [\|u_\theta(z_t, r, t) - \text{sg}(u_{\text{tgt}})\|_2^2], \quad (4)$$

where $u_{\text{tgt}} = v_t(z_t, t) - (t-r)(v_t(z_t, t)\partial_z u_\theta + \partial_t u_\theta)$ and $\text{sg}(\cdot)$ denotes the stop-gradient operation.

The training strategy proposed by [18] samples a portion of the minibatch with $t=r$, in which case the MeanFlow objective in Eq. 4 reduces to flow matching by learning the instantaneous velocity v . Hence, their loss can be interpreted as the sum of two terms for the average velocity $\mathcal{L}_u(z_t, r, t)$ and the instantaneous velocity $\mathcal{L}_v(z_t, t)$:

$$\mathcal{L}_{MF} = \mathbb{E}_{x,\epsilon,t,r} [\mathcal{L}_u(z_t, r, t) \cdot \mathbb{I}(t \neq r) + \mathcal{L}_v(z_t, t) \cdot \mathbb{I}(t = r)]. \quad (5)$$

4. Observations

As described in Section 3, the MeanFlow objective decomposes into learning instantaneous velocity v and average velocity u . In the following, we study how these coupled quantities interact during training and how to optimize their interplay to maximize one-step generation performance. For these experiments, we use a DiT-B/4 [51] architecture and the ImageNet 256 \times 256 dataset [11]. Further details are provided in Appendix A.

4.1. Impact of v -Learning on u -Learning

We first investigate how learning the *instantaneous velocity* v affects the quality of the learned *average velocity* u . Through controlled experiments, we demonstrate that establishing v is a prerequisite for u -learning.

Learning v facilitates learning u . We examine whether and how instantaneous velocity v affects the learning of the average velocity u . We conduct a two-stage training: we first train the model with v -loss ($\mathcal{L}_v(z_t, t)$ in Eq. 5) and then finetune it with u -loss ($\mathcal{L}_u(z_t, r, t)$ in Eq. 5). This experimental setup makes it possible to analyze what effect v -pretraining has on learning u . We use 1-NFE FID as an evaluation metric for the quality of u .

Figure 2 shows the results of two complementary settings. In the first setting (top), we fix the budget for finetuning u at 60 epochs while varying the duration of v -pretraining between $\{0, 5, 10, 15, 20\}$ epochs. In the second setting (bottom), we fix the total training budget at




Figure 2. ***v*-learning facilitates *u*-learning.** (Top) 1-NFE FID during *u*-finetuning according to *v*-pretraining epochs. (Bottom) 1-NFE FID under a fixed 80-epoch budget with varying allocation between *v*-pretraining and *u*-finetuning. Both settings show that investing in *v*-learning improves *u*-learning quality.




Figure 3. **Corruption in *v*-learning disrupts *u*-learning.** 1-NFE FID when training with \mathcal{L}_{MF} while injecting Gaussian noise scaled by $k \cdot \|v_t(z_t|\epsilon)\|$ into the target velocity of \mathcal{L}_v . Even small noise ($k = 0.03$) disrupts *v*-learning and severely degrades *u*-learning performance compared to clean training ($k = 0$).

80 epochs and allocate $\{0, 5, 10, 15, 20\}$ epochs to *v*-pretraining, with the remainder dedicated to *u*-finetuning. Clearly, investing more into *v*-pretraining yields more stable and accurate *u*-learning. In the second setting, even under a fixed compute budget, prioritizing *v* earlier is more effective and accelerates convergence. Overall, the results suggest that a well-formed *v* is a necessary prerequisite for subsequent *u*-learning.

Corruption in *v*-learning disrupts *u*-learning. To examine the opposite case, we investigate whether *u* can be accurately learned when *v*-learning is deliberately corrupted. During MeanFlow training, we intentionally inject Gaussian noise (scaled by $k \cdot \|v_t(z_t|\epsilon)\|$) into the target conditional velocity of the *v*-loss, thereby degrading *v*-learning while leaving the *u*-loss intact. We again use 1-NFE FID as an evaluation metric for *u*-learning quality and illustrate the results across noise scales k in Fig. 3. Even with small noise ($k = 0.03$), *u*-learning severely degrades. In other




Figure 4. **Impact of Δt of *u*-learning on *v*-learning.** 32-NFE FID after 40 epochs of *u* finetuning across different Δt ranges, starting from either random initialization (blue) or *v*-pretrained model (orange, 40 epochs). Small Δt enables constructing and improving *v*, while large Δt degrades pretrained *v*. The green line denotes the performance of the *v*-pretrained model.




Figure 5. **Task affinity between *v*- and *u*-learning across Δt ranges.** Small- Δt *u*-pretraining achieves higher affinity for large Δt compared to *v*-pretraining, providing a better regime for learning large-gap average velocity with instantaneous velocity.

words, a corrupted instantaneous velocity makes learning of the average velocity a lot harder.

Implication. The two experiments above reveal symmetric dependencies: *u*-learning benefits from well-formed *v*, while failing when *v* is corrupted. This aligns with the mathematical structure of MeanFlow—*u* is defined as the temporal integral of *v*, so the latter must be reasonably well established to learn the former. These findings suggest that instantaneous velocity should be prioritized early in training, to lay the foundation for subsequent learning of the average velocity.

4.2. Impact of *u*-Learning on *v*-Learning

Complementary to Sec. 4.1, we analyze how the temporal gap $\Delta t = t - r$ in average velocity supervision influences the instantaneous velocity. We consider two initializations: (1) a model pretrained with *v*-loss for 40 epochs, and (2) random initialization. Both models are then finetuned with the *u*-loss for 40 epochs while restricting Δt to one of four ranges: $[0.1, 0.3]$, $[0.3, 0.5]$, $[0.5, 0.7]$, or $[0.7, 0.9]$. The experimental design exposes how *u*-supervision with different Δt modifies a pretrained instantaneous velocity, and how effectively it forms one from scratch. To measure the quality of *v*, we set $v(z_t, t) = u_\theta(z_t, t, t)$ and evaluate 32-NFE FID, see results in Fig. 4.

Small Δt supervision forms v . Restricting u -learning to small temporal gaps ($\Delta t \in [0.1, 0.3]$) reveals two notable effects (Fig. 4). First, a model trained from scratch with this u -loss achieves a 32-NFE FID, comparable to 40 epochs of v -pretraining (green line), demonstrating that small- Δt supervision is a viable proxy for v -learning. Second, when finetuning the v -pretrained model, the same u -loss yields additional FID gains, indicating that it also improves the pretrained instantaneous velocity. Together, these results demonstrate that small- Δt supervision provides an effective learning signal for the instantaneous velocity v .

Large Δt supervision deteriorates v . In contrast, u -learning with larger temporal gaps ($\Delta t \in [0.3, 0.5]$, $[0.5, 0.7]$, $[0.7, 0.9]$) yields poor 32-NFE FID, with both initializations. *i.e.*, large- Δt supervision is of limited use to construct v from scratch and, even worse, seriously degrades an already pretrained instantaneous velocity.

Implication. When training is not properly managed, there is a self-destructive dynamic: u -learning requires a stable v foundation, yet large- Δt supervision destabilizes it. Together with the earlier finding that v should be established early (Sec. 4.1), this translates into the following guideline: suppress large- Δt supervision in early training stages, to avoid disrupting v -learning.

4.3. Task Affinity Analysis

From Sec. 4.1, we observe that v should be established early as a foundation for u -learning. However, as shown in Sec. 4.2, supervision with large temporal gaps destabilizes v , while supervision with small temporal gaps can actually benefit v learning. Therefore, we should exclude large- Δt supervision from the early training phase, which should be dedicated to forming v . In this regard, two viable strategies emerge: (1) pure v -learning via v -loss, or (2) combining u -loss with small temporal gap supervision, which serves as an effective signal for v -learning.

To determine whether pure v -loss or small- Δt u -loss better prepares the ground for large- Δt learning, we resort to the Task Affinity Score (TAS) [15, 20, 67]. TAS measures how smoothly two tasks can be jointly trained, quantifying their lack of conflict through training iterations. We compute this between v -loss and u -loss (across different Δt -ranges) under three distinct initialization schemes: *baseline*: random initialization; *strategy 1*: pretrained with v -loss for 40 epochs; *strategy 2*: pretrained with u -loss ($\Delta t \in [0.1, 0.3]$) for 40 epochs. The results, shown in Fig. 5, clarify how to best initialize large- Δt learning.

The models pretrained with strategies 1 and 2 both have higher TAS than random initialization across all temporal gap ranges, showing that both pretraining strategies create a more favorable regime for joint MeanFlow learning. Among the two, pretraining with small- Δt supervision exhibits a stronger affinity for large- Δt regimes compared to

pure v -loss pretraining, indicating that small- Δt supervision provides a more favorable initialization for the later learning stages that extend u to large temporal gaps.

Implication. Small- Δt supervision (Strategy 2) creates a better initialization for the challenging large- Δt regime than pure v -loss (Strategy 1). Hence, small- Δt supervision should be included early in training, preparing for the subsequent stages that learn u over large temporal gaps—the mode ultimately required for one-step generation.

5. Method

In Sec. 4, we made three key observations: **(O1)** instantaneous velocity (v) must be established early, as it provides the foundation for learning average velocity (u)—when v is poorly formed or corrupted, u -learning fails; **(O2)** the temporal gap determines how u -learning affects v : small temporal gaps facilitate formation and refinement of v , while large temporal gaps degrade it; and **(O3)** task affinity analysis reveals that including small-gap supervision creates a more favorable initialization for learning the large-gap average velocity required for one-step generation.

The original MeanFlow objective does not consider these properties and thus trains with standard v -loss and the whole range of Δt from the start. Thus, the training suffers from the observed inefficiencies: it fails to rapidly form v , which in turn delays learning of u , since it depends on well-formed instantaneous velocities; resulting in slow training and suboptimal model performance.

We translate our insights into a unified training strategy that directly addresses these limitations, thereby accelerating convergence and improving overall generation quality. Our strategy has two components: (1) the use of well-established training acceleration techniques for diffusion and flow models to rapidly form v , and (2) progressive weighting of the u -loss ($\mathcal{L}_u(z_t, r, t)$), such that it prioritizes small- Δt early in training and gradually transitions to uniform weighting across all Δt .

Accelerating v -learning. To rapidly establish v , we adopt two standard acceleration techniques: specialized timestep sampling and time-dependent loss weighting. For timestep sampling, we replace the base sampling in Eq. 4 with a custom distribution $p_{\text{acc}}(t)$. For loss weighting, we apply a time-dependent weight $\alpha(t)$ to the v -loss, modifying the corresponding term in Eq. 4 to $\alpha(t) \cdot \mathcal{L}_v(z_t, t)$. The specific forms of $p_{\text{acc}}(t)$ and $\alpha(t)$ are adopted from established acceleration methods [10, 20, 25, 33, 73, 82–84] and are designed to focus the model training on more critical timesteps, thereby accelerating convergence.

Progressive \mathcal{L}_u weighting. We weight $\mathcal{L}_u(z_t, r, t)$ to emphasize small temporal gaps early in training, then gradually transition to uniform weighting. Specifically, we use

$$\beta(\Delta t, s) = 1 - s + \lambda s(1 - \Delta t),$$

Method	#Params	Epoch	1-NFE FID↓	2-NFE FID↓	Method	#Params	NFE	FID↓
<i>Comparison to MeanFlow across model sizes</i>					<i>GANs / Normalizing Flows</i>			
MeanFlow-B/4	131M	240	11.58	7.85	BigGAN [8]	112M	1	6.95
+ Ours w MinSNR	131M	240	9.87	7.08	StyleGAN-XL [58]	166M	1	2.30
+ Ours w DTD	131M	240	<u>10.20</u>	<u>7.31</u>	GigaGAN [31]	569M	1	3.45
MeanFlow-M/2	308M	240	5.01	4.61	STARFlow [24]	1.4B	1	2.40
+ Ours w MinSNR	308M	240	<u>4.61</u>	<u>4.30</u>	<i>Autoregressive / Masking models</i>			
+ Ours w DTD	308M	240	4.43	4.10	VQ-GAN [14]	227M	1024	26.52
MeanFlow-L/2	459M	240	3.84	3.35	MaskGIT [9]	227M	8	6.18
+ Ours w MinSNR	459M	240	<u>3.79</u>	<u>3.31</u>	VAR [69]	2B	10×2	1.92
+ Ours w DTD	459M	240	3.47	3.24	MAR-H [37]	943M	256×2	1.55
<i>Few-step diffusion/flow models from end-to-end training</i>					<i>Diffusion / Flow models</i>			
iCT-XL/2 [‡] [63]	675M	-	34.24	20.30	ADM [12]	554M	250×2	10.94
Shortcut-XL/2 [16]	675M	160	10.60	-	LDM [53]	400M	250×2	3.60
iMM-XL/2 [†] [85]	675M	3840	-	7.77	U-ViT-H/2 [3]	501M	50×2	2.29
MeanFlow-XL/2+ [18]	676M	1000	-	2.20	SimDiff [29]	2B	512×2	2.77
MeanFlow-XL/2	676M	240	3.43	2.93	DTR-L/2 [50]	458M	250×2	2.33
+ Ours w DTD	676M	240	2.87	2.64	DiT-XL/2 [51]	675M	250×2	2.27
					SiT-XL/2 [48]	675M	250×2	2.06

Table 1. **Results for class-conditional generation on ImageNet 256×256.** (Left) Comparison of few-step diffusion/flow models. (Right) Other generative model families as reference. “†” in left table and “×2” in right table indicate that Classifier-Free Guidance (CFG) doubles NFE per sampling step. ‡: iCT results from [85].

where $s \in [0, 1]$ denotes the training progress. At initialization $s = 1$ and $\beta(\Delta t, 1) = \lambda(1 - \Delta t)$ prioritizes small Δt ; at convergence $\beta(\Delta t, 0) = 1$ weights all gaps equally. To maintain a uniform expectation at initialization, we set $\lambda = 1/\mathbb{E}_{\Delta t}[1 - \Delta t]$. We use a simple linear schedule $s = 1 - i/T$, where i and T denote the current iteration and the total number of iterations, respectively. A slower initial transition can be achieved by setting $s = 1 - (i/T)^k$ with $k > 1$; conversely, $k < 1$ yields a faster transition.

Integration with MeanFlow components. MeanFlow training employs a number of stabilization techniques, including specialized loss metrics and sampling strategies. In Appendix B we provide detailed instructions on how to integrate our proposed adaptations with these components.

6. Experiments

In the following, we show that our method accelerates convergence and attains higher final performance. Together, these results show that our observations translate into practical training improvements.

6.1. Experimental Setup

We follow the original experimental setup of MeanFlow [18] and conduct experiments on ImageNet [11] generation at 256×256 resolution with DiT architectures [51]. To measure few-step generation performance, we utilize the

FID [27] score on 50K samples from either 1-NFE or 2-NFE generation.

Implementation details. We test several ways of accelerating velocity training, one method from each category. For their simplicity and good empirical performance, we choose MinSNR [25] as loss weighting approach, and DTD [35] as timestep sampling approach. More sophisticated methods exist [20, 73, 84], but we do not expect large differences in the context of our progressive weighting scheme.

6.2. Comparative Evaluation

Comparison to MeanFlow. As shown in Table 1 (top left), we first compare our method against MeanFlow using the DiT-B/4, DiT-M/2, and DiT-L/2 models. Our method consistently outperforms MeanFlow, regardless of the velocity learning acceleration technique employed. When comparing the two acceleration methods, MinSNR and DTD, a clear pattern emerges: MinSNR surpasses DTD on DiT-B/4, but this advantage diminishes at larger scales (L/2, M/2). We attribute this discrepancy to MeanFlow’s adaptive loss weighting, which normalizes loss values by their norm to balance their influence. Loss-weighting strategies like MinSNR interfere with this adaptive mechanism, consequently reducing robustness across different model scales. In contrast, timestep-sampling methods like DTD only modify the sampling distribution and leave the loss weighting scheme intact, thus preserving compatibility with




Figure 6. Convergence speed comparison between MeanFlow and our methods across model sizes.

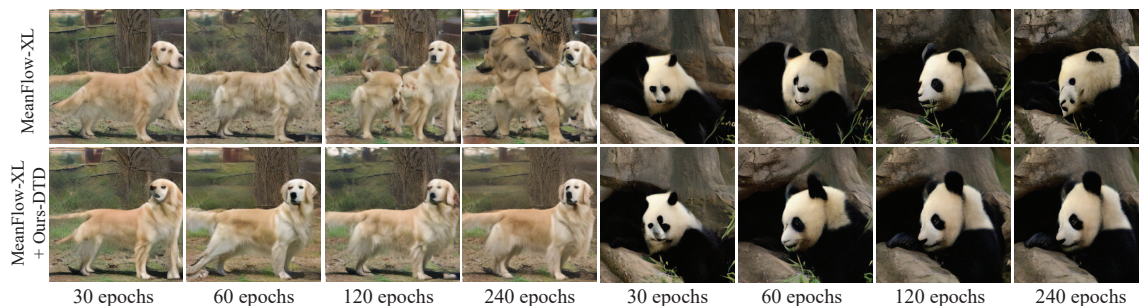


Figure 7. Qualitative comparison of generated samples across training epochs on DiT-XL/2.

MeanFlow’s adaptive design. Therefore, due to its consistent performance across model sizes, we select DTD as our primary acceleration method.

ImageNet 256×256 benchmark. We scale up our method with DTD to the DiT-XL model and compare it with previous one- and few-step diffusion/flow models in Table 1 (bottom left). Our method also outperforms MeanFlow in the DiT-XL setup (240 epochs), improving MeanFlow’s 1-NFE FID from 3.43 to 2.87 and its 2-NFE FID from 2.93 to 2.64. We highlight that this improvement on DiT-XL (3.43→2.87), a notable 16% reduction in FID, substantially narrows the gap between one-step generation and multi-step diffusion models, as shown in Table 1 (right).

Convergence speed. Next, we compare the convergence behavior of our method against vanilla MeanFlow training, both quantitatively and qualitatively. As shown in Fig. 1 and 6, our approach achieves substantially faster convergence across all model sizes. Notably, our method with DTD demonstrates superior convergence on DiT-XL/2, -L/2, and -M/2, achieving approximately 2.5×, 2.3×, and 2.1× speedup, respectively. Qualitatively, Fig. 7 shows samples from DiT-XL/2 at different training epochs. At equivalent epochs, our method generates noticeably sharper and more detailed images than vanilla MeanFlow. Remarkably, our samples at 120 epochs exhibit quality compa-

CFG Setup	Method	FID (1-NFE)↓	FID (2-NFE)↓
$\kappa = 0.5$ $\omega = 2.0$	MeanFlow-L/2	4.60	4.58
	+ Ours w DTD	4.08	3.97
	MeanFlow-XL/2	4.72	4.64
	+ Ours w DTD	4.31	4.29
$\kappa = 0.92$ $\omega = 2.5$	MeanFlow-L/2	3.84	3.35
	+ Ours w DTD	3.47	3.24
$\kappa = 0.92$ $\omega = 2.5$	MeanFlow-XL/2	3.43	2.93
	+ Ours w DTD	2.87	2.64

Table 2. **Robustness to CFG configurations.** Performance of DiT-L and DiT-XL with different CFG setups. Our method consistently improves across all configurations.

table to or exceeding vanilla MeanFlow’s 240-epoch results, demonstrating $\approx 2\times$ faster convergence. Together, these quantitative and qualitative results confirm that our approach not only substantially accelerates training but also improves final generation quality.

6.3. Ablation and Analysis

Effectiveness across different CFG configurations. MeanFlow integrates Classifier-Free Guidance (CFG) by mixing three components in the target u_{tgt} : conditional velocity, conditional velocity prediction, and unconditional velocity prediction, using coefficients ω and κ (see Ap-

Method	FID (1-NFE)↓	FID (2-NFE)↓
MeanFlow-B/4	11.58	7.85
+ MinSNR	10.57	7.38
+ DTD	10.96	7.55
+ \mathcal{L}_u weighting.	10.98	7.58
+ MinSNR + \mathcal{L}_u weighting.	9.87	7.08
+ DTD + \mathcal{L}_u weighting.	<u>10.20</u>	<u>7.31</u>

Table 3. **Ablation of method components.** Velocity acceleration methods and \mathcal{L}_u weighting each improve upon vanilla MeanFlow training, with their combination achieving the best performance.

	MeanFlow-B/4	MeanFlow-B/4 + Ours w DTD			
		$k = 0.5$	$k = 1$	$k = 2$	$k = 3$
FID↓ (1-NFE)	11.58	<u>11.16</u>	10.20	11.44	11.99

Table 4. **1-NFE FID with varying settings of the schedule parameter k .** The linear schedule ($k = 1$) achieves the lowest FID.

pendix A in [18]). For smaller models (DiT-B and DiT-M), $\kappa = 0.5$ and $\omega = 2.0$ are used, while larger models (DiT-L and DiT-XL) use $\kappa = 0.92$ and $\omega = 2.5$. Additionally, for DiT-L and DiT-XL, CFG mixing is only applied when t is sampled within specific ranges ([0.0, 0.8] for DiT-L, [0.0, 0.75] for DiT-XL). To verify our method’s robustness to these configurations, we train DiT-L and DiT-XL using the CFG settings of both the small and large models.

Table 2 shows that our method consistently improves performance across all CFG configurations. This demonstrates the robustness of our method to different guidance scenarios, underscoring its role as a simple but highly effective technique for achieving consistent performance gains.

Effect of individual components. To validate the contribution of each component, we train DiT-B/4 with three configurations: (1) applying only velocity acceleration methods (MinSNR or DTD), (2) applying only progressive \mathcal{L}_u weighting, and (3) the full method combining both components. All models are trained for 240 epochs, matching the setup in Sec. 6.2. The results are illustrated in Table 3. As shown in the results, only applying velocity acceleration methods improves vanilla MeanFlow training (11.58 \rightarrow 10.57 with MinSNR, 10.96 with DTD at 1-NFE), demonstrating that rapid v formulation benefits training. Moreover, applying only progressive weighting on \mathcal{L}_u improves the performance to 10.98 at 1-NFE, showing the benefit of proper temporal gap scheduling. The combined approach yields the strongest results, with a combination of MinSNR and weighting yielding an FID of 9.87 at 1-NFE and 10.20 when combining DTD and weighting. This demonstrates that the two components are complementary: velocity acceleration rapidly establishes the instantaneous velocity foundation, while progressive weighting drives effective average velocity learning.

Method	FID↓		
	32-NFE	64-NFE	128-NFE
MeanFlow	7.61	7.26	7.16
+ Ours w MinSNR	7.09	6.91	6.86
+ Ours w DTD	<u>7.25</u>	<u>7.01</u>	<u>6.93</u>

Table 5. **Comparison of multi-step generation with estimated instantaneous velocity on DiT-B/4 models.**

Effect of schedule parameter k . While the pace of our progressive weighting $s = 1 - (i/T)^k$ can be modulated by adjusting k , we use $k = 1$ (linear schedule) in all experiments for simplicity. To verify the impact of this choice, we compare different k values on DiT-B/4 in Table 4. The results show that $k = 1$ achieves the best performance (10.20 FID), validating our choice of a simple linear schedule. We observe that slower transitions ($k < 1$) and faster transitions ($k > 1$) both degrade performance, indicating that a linear schedule provides a good balance.

Enhanced instantaneous velocity. We argue that accelerating instantaneous velocity formation is crucial for effective training (Sec. 4). To verify that the acceleration methods improve v quality, we evaluate multi-step generation performance using $u_\theta(z_t, t, t)$ as the velocity estimate (which equals the model’s instantaneous velocity prediction when $r = t$). As shown in Table 5, our methods consistently outperform MeanFlow across all sampling steps (32, 64, 128 NFE) regardless of the applied velocity acceleration methods. This demonstrates that our training strategy actually improves the underlying instantaneous velocity estimated by the model, confirming that our approach yields higher-quality velocity fields.

7. Conclusion

In this work, we revisited MeanFlow through a detailed analysis of its instantaneous and average velocity components and their interaction during training. We identified key learning dynamics: instantaneous velocity should be established early to enable effective average velocity learning, and supervision on small temporal gaps creates a more favorable foundation for subsequently learning large-gap average velocity. Building on these insights, we propose an improved training strategy that accelerates instantaneous velocity formation and prioritizes small temporal gaps in the early training phase before gradually transitioning to larger gaps. Our approach achieves faster convergence, improved one-step generation, and reduced training cost. We believe these findings offer both a deeper understanding of MeanFlow’s learning behavior and a practical foundation for designing efficient few-step generative models.

Acknowledgements. This work was supported by ETHAR (ETH Zürich Augmented Reality Research Lab), funded through a Google donation; and by a grant from the Swiss National Supercomputing Center within the Swiss AI Initiative (project a144). Julius Erbach is supported by the Max Planck ETH Center for Learning Systems.

References

- [1] Krea AI. Krea Realtime 14B: Real-time video generation, 2025. 1
- [2] Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *International Conference on Learning Representations (ICLR)*, 2023. 1, 2, 3
- [3] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 6
- [4] Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, Sumith Kulal, et al. FLUX.1 Kontext: Flow matching for in-context image generation and editing in latent space. *preprint arXiv:2506.15742*, 2025. 1
- [5] 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. *preprint arXiv:2303.04248*, 2023. 2
- [6] Nicholas M Boffi, Michael S Albergo, and Eric Vanden-Eijnden. Flow map matching. *preprint arXiv:2406.07507*, 2024. 1, 3
- [7] Nicholas M Boffi, Michael S Albergo, and Eric Vanden-Eijnden. How to build a consistency model: Learning flow maps via self-distillation. *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. 1, 3
- [8] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations (ICLR)*, 2019. 6
- [9] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. MaskGIT: Masked generative image transformer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 6
- [10] Jooyoung Choi, Jungbeom Lee, Chaehun Shin, Sungwon Kim, Hyunwoo Kim, and Sungroh Yoon. Perception prioritized training of diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 5
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 2, 3, 6
- [12] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 6
- [13] Tim Dockhorn, Arash Vahdat, and Karsten Kreis. GENIE: Higher-order denoising diffusion solvers. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 1
- [14] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 6
- [15] Chris Fifty, Ehsan Amid, Zhe Zhao, Tianhe Yu, Rohan Anil, and Chelsea Finn. Efficiently identifying task groupings for multi-task learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 5
- [16] Kevin Frans, Danijar Hafner, Sergey Levine, and Pieter Abbeel. One step diffusion via shortcut models. In *International Conference on Learning Representations (ICLR)*, 2025. 1, 2, 3, 6
- [17] Xin Gao, Li Hu, Siqi Hu, Mingyang Huang, Chaonan Ji, Dechao Meng, Jinwei Qi, Penchong Qiao, Zhen Shen, Yafei Song, et al. Wan-S2V: Audio-driven cinematic video generation. *preprint arXiv:2508.18621*, 2025. 1
- [18] Zhengyang Geng, Mingyang Deng, Xingjian Bai, J Zico Kolter, and Kaiming He. Mean flows for one-step generative modeling. *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. 2, 3, 6, 8, 1
- [19] Zhengyang Geng, Ashwini Pople, William Luo, Justin Lin, and J Zico Kolter. Consistency models made easy. *International Conference on Learning Representations (ICLR)*, 2025. 1, 3
- [20] Hyojun Go, Yunsung Lee, Seunghyun Lee, Shinhyeok Oh, Hyeongdon Moon, and Seungtaek Choi. Addressing negative transfer in diffusion models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 2, 5, 6
- [21] Hyojun Go, Dominik Narnhofer, Goutam Bhat, Prune Truong, Federico Tombari, and Konrad Schindler. VIST3A: Text-to-3d by stitching a multi-view reconstruction network to a video generator. *preprint arXiv:2510.13454*, 2025. 1
- [22] Hyojun Go, Byeongjun Park, Jiho Jang, Jin-Young Kim, Soonwoo Kwon, and Changick Kim. SplatFlow: Multi-view rectified flow model for 3d gaussian splatting synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 1
- [23] Jiatao Gu, Shuangfei Zhai, Yizhe Zhang, Lingjie Liu, and Joshua M Susskind. BOOT: Data-free distillation of denoising diffusion models with bootstrapping. In *ICML Workshop on Structured Probabilistic Inference & Generative Modeling*, 2023. 2
- [24] Jiatao Gu, Tianrong Chen, David Berthelot, Huangjie Zheng, Yuyang Wang, Ruixiang Zhang, Laurent Dinh, Miguel Angel Bautista, Josh Susskind, and Shuangfei Zhai. STARFlow: Scaling latent normalizing flows for high-resolution image synthesis. *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. 6
- [25] Tiankai Hang, Shuyang Gu, Chen Li, Jianmin Bao, Dong Chen, Han Hu, Xin Geng, and Baining Guo. Efficient diffusion training via min-SNR weighting strategy. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 1, 2, 5, 6

- [26] Jonathan Heek, Emiel Hoogeboom, and Tim Salimans. Multistep consistency models. *preprint arXiv:2403.06807*, 2024. 1, 3
- [27] 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. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 6
- [28] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 1, 2
- [29] Emiel Hoogeboom, Jonathan Heek, and Tim Salimans. Simple diffusion: End-to-end diffusion for high resolution images. In *International Conference on Machine Learning (ICML)*, 2023. 6
- [30] Zheyuan Hu, Chieh-Hsin Lai, Yuki Mitsufuji, and Stefano Ermon. CMT: Mid-training for efficient learning of consistency, mean flow, and flow map models. *preprint arXiv:2509.24526*, 2025. 3
- [31] Minguk Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. Scaling up GANs for text-to-image synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 6
- [32] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 1, 2
- [33] Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. Analyzing and improving the training dynamics of diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 5
- [34] Dongjun Kim, Chieh-Hsin Lai, Wei-Hsiang Liao, Naoki Murata, Yuhta Takida, Toshimitsu Uesaka, Yutong He, Yuki Mitsufuji, and Stefano Ermon. Consistency trajectory models: Learning probability flow ode trajectory of diffusion. In *International Conference on Learning Representations (ICLR)*, 2024. 1, 3
- [35] Jin-Young Kim, Hyojun Go, Soonwoo Kwon, and Hyun-Gyoon Kim. Denoising task difficulty-based curriculum for training diffusion models. In *International Conference on Learning Representations (ICLR)*, 2025. 2, 6, 1
- [36] Fei Kong, Jinhao Duan, Lichao Sun, Hao Cheng, Renjing Xu, Hengtao Shen, Xiaofeng Zhu, Xiaoshuang Shi, and Kaidi Xu. ACT-Diffusion: Efficient adversarial consistency training for one-step diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [37] Tianhong Li, Yonglong Tian, He Li, Mingyang Deng, and Kaiming He. Autoregressive image generation without vector quantization. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 6
- [38] Zongrui Li, Minghui Hu, Qian Zheng, and Xudong Jiang. Connecting consistency distillation to score distillation for text-to-3d generation. In *European Conference on Computer Vision (ECCV)*, 2024. 2
- [39] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations (ICLR)*, 2023. 1, 2, 3
- [40] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *International Conference on Learning Representations (ICLR)*, 2023. 1, 2, 3
- [41] Yunpeng Liu, Boxiao Liu, Yi Zhang, Xingzhong Hou, Guanglu Song, Yu Liu, and Haihang You. See further when clear: Curriculum consistency model. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 2
- [42] Cheng Lu and Yang Song. Simplifying, stabilizing and scaling continuous-time consistency models. *International Conference on Learning Representations (ICLR)*, 2025. 1, 3
- [43] 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 (NeurIPS)*, 2022. 1
- [44] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. PM-Solver++: Fast solver for guided sampling of diffusion probabilistic models. *Machine Intelligence Research*, 22:730–751, 2025. 1
- [45] Eric Luhman and Troy Luhman. Knowledge distillation in iterative generative models for improved sampling speed. *preprint arXiv:2101.02388*, 2021. 1, 2
- [46] Simian Luo, Yiqin Tan, Longbo Huang, Jian Li, and Hang Zhao. Latent consistency models: Synthesizing high-resolution images with few-step inference. *preprint arXiv:2310.04378*, 2023. 3
- [47] 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 (NeurIPS)*, 2023. 1, 2
- [48] 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. In *European Conference on Computer Vision (ECCV)*, 2024. 1, 6
- [49] Chenlin Meng, Robin Rombach, Ruiqi Gao, Diederik Kingma, Stefano Ermon, Jonathan Ho, and Tim Salimans. On distillation of guided diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 2
- [50] Byeongjun Park, Sangmin Woo, Hyojun Go, Jin-Young Kim, and Changick Kim. Denoising task routing for diffusion models. In *International Conference on Learning Representations (ICLR)*, 2024. 6
- [51] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 3, 6
- [52] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *preprint arXiv:2204.06125*, 2022. 2

- [53] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 6
- [54] Amirmojtaba Sabour, Sanja Fidler, and Karsten Kreis. Align your steps: Optimizing sampling schedules in diffusion models. In *International Conference on Machine Learning (ICML)*, 2024. 1
- [55] Amirmojtaba Sabour, Sanja Fidler, and Karsten Kreis. Align your flow: Scaling continuous-time flow map distillation. *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. 1, 3
- [56] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations (ICLR)*, 2022. 1, 2
- [57] Tim Salimans, Thomas Mensink, Jonathan Heek, and Emiel Hoogeboom. Multistep distillation of diffusion models via moment matching. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 2
- [58] Axel Sauer, Katja Schwarz, and Andreas Geiger. StyleGAN-XL: Scaling stylegan to large diverse datasets. In *ACM SigGraph*, 2022. 6
- [59] Axel Sauer, Frederic Boesel, Tim Dockhorn, Andreas Blattmann, Patrick Esser, and Robin Rombach. Fast high-resolution image synthesis with latent adversarial diffusion distillation. In *ACM SigGraph Asia*, 2024. 1, 2
- [60] Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion distillation. In *European Conference on Computer Vision (ECCV)*, 2024. 1, 2
- [61] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning (ICML)*, 2015. 1, 2
- [62] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations (ICLR)*, 2021. 1
- [63] Yang Song and Prafulla Dhariwal. Improved techniques for training consistency models. *International Conference on Learning Representations (ICLR)*, 2024. 1, 3, 6
- [64] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 1, 2
- [65] 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 (ICLR)*, 2021. 2
- [66] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *International Conference on Machine Learning (ICML)*, 2023. 1, 2
- [67] Trevor Standley, Amir Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Which tasks should be learned together in multi-task learning? In *International Conference on Machine Learning (ICML)*, 2020. 5
- [68] 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 (TMLR)*, 2024. 2
- [69] Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 6
- [70] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *preprint arXiv:2503.20314*, 2025. 1
- [71] Fu-Yun Wang, Zhengyang Geng, and Hongsheng Li. Stable consistency tuning: Understanding and improving consistency models. *preprint arXiv:2410.18958*, 2024. 1, 3
- [72] Fu-Yun Wang, Zhaoyang Huang, Alexander Bergman, Dazhong Shen, Peng Gao, Michael Lingelbach, Keqiang Sun, Weikang Bian, Guanglu Song, Yu Liu, et al. Phased consistency models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 1, 3
- [73] Kai Wang, Mingjia Shi, Yukun Zhou, Zekai Li, Zhihang Yuan, Yuzhang Shang, Xiaojiang Peng, Hanwang Zhang, and Yang You. A closer look at time steps is worthy of triple speed-up for diffusion model training. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 2, 5, 6
- [74] Zidong Wang, Yiyuan Zhang, Xiaoyu Yue, Xiangyu Yue, Yangguang Li, Wanli Ouyang, and Lei Bai. Transition models: Rethinking the generative learning objective. *preprint arXiv:2509.04394*, 2025. 1, 3
- [75] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, et al. Qwen-image technical report. *preprint arXiv:2508.02324*, 2025. 1
- [76] Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and Bill Freeman. Improved distribution matching distillation for fast image synthesis. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 1, 2
- [77] Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1, 2
- [78] Han Zhang and Fan Cheng. Improving consistency distillation with rectified trajectories. In *International Conference on Artificial Neural Networks (ICANN)*, 2025. 2
- [79] Huijie Zhang, Aliaksandr Siarohin, Willi Menapace, Michael Vasilkovsky, Sergey Tulyakov, Qing Qu, and Ivan Skorokhodov. AlphaFlow: Understanding and improving MeanFlow models. *preprint arXiv:2510.20771*, 2025. 3
- [80] Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator. In *International Conference on Learning Representations (ICLR)*, 2023. 1
- [81] Hongkai Zheng, Weili Nie, Arash Vahdat, Kamyar Azizadenesheli, and Anima Anandkumar. Fast sampling of diffusion models via operator learning. In *International Conference on Machine Learning (ICML)*, 2023. 1, 2
- [82] Tianyi Zheng, Cong Geng, Peng-Tao Jiang, Ben Wan, Hao Zhang, Jinwei Chen, Jia Wang, and Bo Li. Non-uniform

- timestep sampling: Towards faster diffusion model training. In *ACM International Conference on Multimedia*, 2024. 2, 5
- [83] Tianyi Zheng, Peng-Tao Jiang, Ben Wan, Hao Zhang, Jinwei Chen, Jia Wang, and Bo Li. Beta-tuned timestep diffusion model. In *European Conference on Computer Vision (ECCV)*, 2024.
- [84] Tianyi Zheng, Jiayang Zou, Peng-Tao Jiang, Hao Zhang, Jinwei Chen, Jia Wang, and Bo Li. Bidirectional beta-tuned diffusion model. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2025. 2, 5, 6
- [85] Linqi Zhou, Stefano Ermon, and Jiaming Song. Inductive moment matching. In *International Conference on Machine Learning (ICML)*, 2025. 1, 3, 6
- [86] 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 *International Conference on Machine Learning (ICML)*, 2024. 1, 2
- [87] Mingyuan Zhou, Zhendong Wang, Huangjie Zheng, and Hai Huang. Guided score identity distillation for data-free one-step text-to-image generation. In *International Conference on Learning Representations (ICLR)*, 2025. 2