

Deeply Supervised Flow-Based Generative Models

Inkyu Shin Chenglin Yang Liang-Chieh Chen

ByteDance Seed

<https://deepflow-project.github.io>

Abstract

Flow-based generative models have charted an impressive path across multiple visual generation tasks by adhering to a simple principle: learning velocity representations of a linear interpolant. However, we observe that training velocity solely from the final layer’s output under-utilizes the rich inter-layer representations, potentially impeding model convergence. To address this limitation, we introduce **DeepFlow**, a novel framework that enhances velocity representation through inter-layer communication. DeepFlow partitions transformer layers into balanced branches with deep supervision and inserts a lightweight Velocity Refiner with Acceleration (VeRA) block between adjacent branches, which aligns the intermediate velocity features within transformer blocks. Powered by the improved deep supervision via the internal velocity alignment, DeepFlow converges **8× faster** on ImageNet-256×256 with equivalent performance and further reduces FID by **2.6** while halving training time compared to previous flow-based models without a classifier-free guidance. DeepFlow also outperforms baselines in text-to-image generation tasks, as evidenced by evaluations on MS-COCO and zero-shot GenEval.

1. Introduction

In the era of generative AI, it is indisputable that the strategy “denoising from noise” has significantly propelled the advancement of visual generation. The processes of introducing and removing noise have given rise to two prominent families of generative models: diffusion-based models [14, 38, 41, 40] and flow-based models [22, 1, 25, 23]. Diffusion-based models utilize a curved trajectory of diffusion forward process and denoise it back using a noise prediction. In contrast, flow-based models simply adopt the linear interpolation between noise and target signals, learning to predict velocity of interpolated noisy image under the principles of normalizing flows [4, 37]. Owing to these

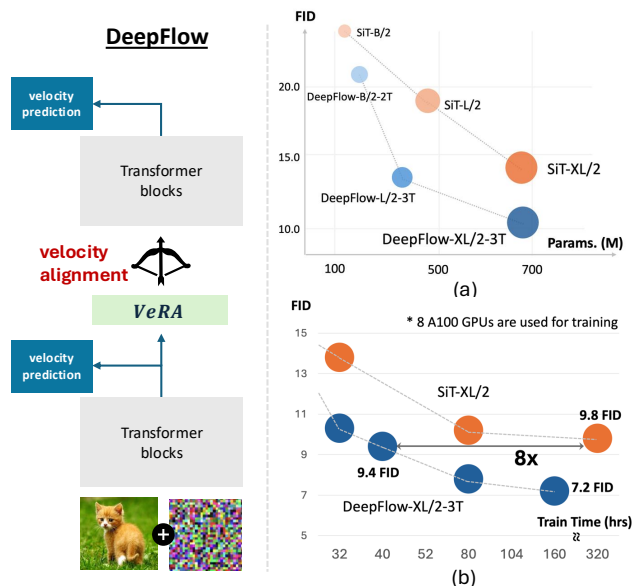


Figure 1. **Overview of DeepFlow.** *Left:* DeepFlow incorporates deep supervision by evenly adding velocity prediction within transformer blocks, further enhanced by the proposed Velocity Alignment block (VeRA). *Right:* (a) On the ImageNet-256 benchmark, DeepFlow consistently outperforms SiT [28] in FID scores across various model sizes. (b) DeepFlow-XL achieves an 8× training efficiency improvement over SiT-XL. See Table 8 for details.

straightforward yet effective noising and denoising mechanisms, flow-based models have achieved state-of-the-art performance across numerous visual generation benchmarks [28, 52, 26, 8, 17, 9].

Despite recent advancements, current fundamental flow-based models [28, 8] have largely overlooked the potential for enhancing their internal velocity representations. As illustrated in Figure 2(a), SiT [28], a representative flow-based model, relies on sequentially stacked multi-layered transformer [45] blocks to learn the velocity exclusively from the final layer. This approach under-utilizes the significance of intermediate velocity representations, leading to challenges such as slow train-




Figure 2. **Importance of Internal Feature Alignment for Flow-Based Models.** Our DeepFlow enhances the baseline flow-based model (a) by explicitly aligning intermediate velocity features with final layer features. As shown in (b), simply applying deep supervision reduces the feature distance between intermediate velocity v_1^* (from middle 6th layer) and final v_2^* (from 12th layer), improving FID scores (light blue bars in (d, e)). To further minimize this distance, we introduce the VeRA block, which refines deeply-supervised intermediate features into $v_{1 \rightarrow 2}^*$, more closely aligned with v_2^* (dark blue bar in (d)). This leads to even better image generation quality (dark blue bars in (e)).

ing convergence and low performance [33, 52]. Recently, to address this limitation, REPA [52] aligns the internal velocity representation with external features from pre-trained self-supervised models (e.g., DINO [3, 30]), resulting in better generative models with fewer training time needed. However, relying solely on external self-supervised models overlooks the opportunity to internally rectify feature misalignment within transformer layers and fails to fully leverage the properties of flow-based models. A natural question thus arises: *Can flow-based models be improved by internally aligning velocity representations across transformer layers instead of relying on external models?*

We begin with a straightforward approach by incorporating deep supervision [20] within the transformer layers of flow-based models to enhance alignment. As depicted in Figure 2(b), flow-based models can employ deep supervision across multiple velocity layers by partitioning transformer blocks into equal-sized branches, each trained to predict the same ground-truth velocity. This approach can align intermediate and final velocity features, as demonstrated by the reduced feature distance¹ in Figure 2(d); SiT-B/2: **7.7** vs. SiT-B/2 with deep supervision: **7.2**. This alignment, in turn, positively affects image generation performance in Figure 2(e); SiT-B/2: **34.4 FID** vs. SiT-B/2 with deep supervision: **33.0 FID**.

However, deep supervision alone is insufficient for achieving optimal alignment between intermediate and final layers, as intermediate layers exhibit a limited capacity for velocity prediction compared to the fi-

¹We compute euclidean distance between v_1^* (from middle 6th layer) and v_2^* (from final 12th layer) across all 250 timesteps using SDE while generating 50k samples on the ImageNet-256×256 dataset. This metric quantifies the alignment between intermediate and final layer velocity features.

nal layer. Motivated by this, we propose a redesigned flow-based model that explicitly aligns internal velocity representations while effectively integrating deep supervision—hereafter referred to as **DeepFlow**. It aims to refine deeply-supervised intermediate velocity features to be aligned for following branch as depicted in Figure 2(c). To achieve this, DeepFlow consists of a lightweight block between adjacent branches, explicitly tailored to learn the mapping of velocity features from preceding branch to subsequent one. This block, termed the **Velocity Refiner with Acceleration (VeRA)** block, is specifically designed to model acceleration. It refines velocity features by conditioning on adjacent branches across different time steps. This process is guided by principles of second-order dynamics. To achieve this, we implement a simple MLP that takes previous velocity feature and is trained to generate an acceleration feature using the second-order ODE as visualized in Figure 3. Afterwards, we concatenate the previous velocity features with the computed acceleration features and apply a time-gap-conditioned adaptive layer normalization, ensuring aligned velocity features for the subsequent branch. Finally, we further refine these features by incorporating spatial information through cross-space attention, facilitating interaction between the refined velocity and spatial feature spaces. We can observe that the velocity feature refined from VeRA block is significantly closer to the final output velocity feature in Figure 2(d); SiT-B/2 with deep supervision: **7.2** vs. DeepFlow-B/2-2T: **2.9**. It successfully leads to enhanced image quality as shown in Figure 2(e); SiT-B/2 with deep supervision: **33.0 FID** vs. DeepFlow-B/2-2T: **28.1 FID**.

Driven by feature alignment strategy for enhanced deep supervision, DeepFlow significantly improves both training efficiency and final image generation quality,

all without dependence on external models. Figure 1(a) shows that DeepFlow-L/2-3T model with smaller number of parameters outperforms the SiT-XL/2 model after 80 epochs of training on ImageNet-256 benchmark [6]. Moreover, our DeepFlow-XL/2-3T model delivers performance comparable to the SiT-XL/2 model while reducing training time by eightfold, as shown in Figure 1(b). It further improves generation quality by outperforming SiT-XL/2 only using half of training time. For optimal image generation, we can seamlessly integrate feature alignment using external self-supervised model (*e.g.*, DINOv2 [30]) and classifier-free guidance [15], yielding better results on both ImageNet-256 and ImageNet-512 while requiring fewer training epochs needed. Additionally, we performed extensive comparisons with a conventional flow-based model on the text-to-image generation benchmark using the MSCOCO dataset [21] and GenEval benchmark [10].

2. Related Works

Generative Models with Denoising Transformers. Recent studies have advanced the field of visual generation by leveraging transformer [45] architecture as a denoising model [41, 14, 29, 38, 24, 47, 34, 35, 12, 36]. Specifically, U-ViT [2] and DiffiT [11] integrate skip connections [39] into transformer-based backbones, whereas DiT [32] demonstrates that a simple transformer-based diffusion network without skip connections can serve as a scalable and effective backbone for diffusion models. Based on this simple architecture from DiT, SiT [28] employs the principle of flow matching [22] and normalizing flows [37], resulting in better image generation quality. Although both studies recognize that scaling laws hold as the number of transformer blocks increases, they overlook the role of internal feature representations across transformer layers.

Feature Enhancement in Generative Models. Several approaches have sought to enhance internal representations in denoising transformers. For example, REPA [52] improves generative modeling by aligning internal features in both diffusion and flow-based models with external representations from pre-trained self-supervised models like MAE [13] and DINO [3, 30]. Similarly, VA-VAE [48] refines tokenizer [19] representations by incorporating external foundational models. However, relying solely on external models may overlook the self-correcting potential in addressing feature misalignment across intermediate layers. Deep supervision [20] addresses this in classification tasks by providing multi-layer supervision that refines internal features for better discriminative performance. Inspired by this, we extend deep supervision to flow-based generative models, where the discriminative quality of velocity

representations is critical.

3. Method

In this section, we first introduce the preliminaries on flow matching in Section 3.1, followed by a detailed presentation of the proposed method, DeepFlow, in Section 3.2.

3.1. Preliminaries

Flow Matching. Normalizing flows [4] conceptualized time-dependent velocity field, $v : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, which can provide flow map, $\phi : [0, 1] \times \mathbb{R}^d$. This flow map aims to push-forward simple pure noise z to target distribution, x_0 . Flow matching [22] applied this finding to generative model by designing a neural network, $v_\theta(\mathbf{x}_t)$ that predicts the velocity of \mathbf{x}_t at time-step t with parameter θ . Thus, it can generate target samples from pure gaussian noise using progressive denoising step with predicted velocity. To this end, during training, forward noising step is conducted with simple linear interpolation between prior noise ($\mathbf{x}_1 \sim \mathcal{N}(0, 1)$) and target distribution (\mathbf{x}_0) as below:

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

where $t \in [0, 1]$ denotes time-step used for interpolation coefficient. Then, the flow-based method is learned to transform $v_\theta(\mathbf{x}_t)$ to be similar to corresponding ground-truth velocity $V = \mathbf{x}_1 - \mathbf{x}_0$ with following objective function:

$$\mathcal{L}(\theta) = \mathbb{E} \|v_\theta(\mathbf{x}_t) - V\|^2 \quad (2)$$

Consistent with a conventional flow-based model [28], we use DiT Transformer [32] as $v_\theta(\cdot)$. It comprises multiple transformer blocks that apply self-attention to the input tokens, followed by AdaLN-Zero modulation conditioned on the time step and class (or text), and a final velocity layer. Equation (2) is thus expressed as below:

$$\mathcal{L}(\theta) = \mathbb{E} \|v_\theta(\mathbf{x}_t, t, c) - V\|^2, \quad (3)$$

where t and c are input time-step and class features. Similarly, SiT [28] adopts DiT Transformer while leveraging flow matching-based noising and denoising process. Depending on the number of transformer blocks and channel dimension used, SiT has four variants, SiT-{S,B,L,XL}.

3.2. DeepFlow

Our DeepFlow aims to exploit the potential of internal features across transformer blocks by enhancing feature alignment. It can be achieved by incorporating Deep Supervision (Section 3.2.1) into flow-based models and designing VeRA Block (Section 3.2.2) for explicit alignment between internal features.

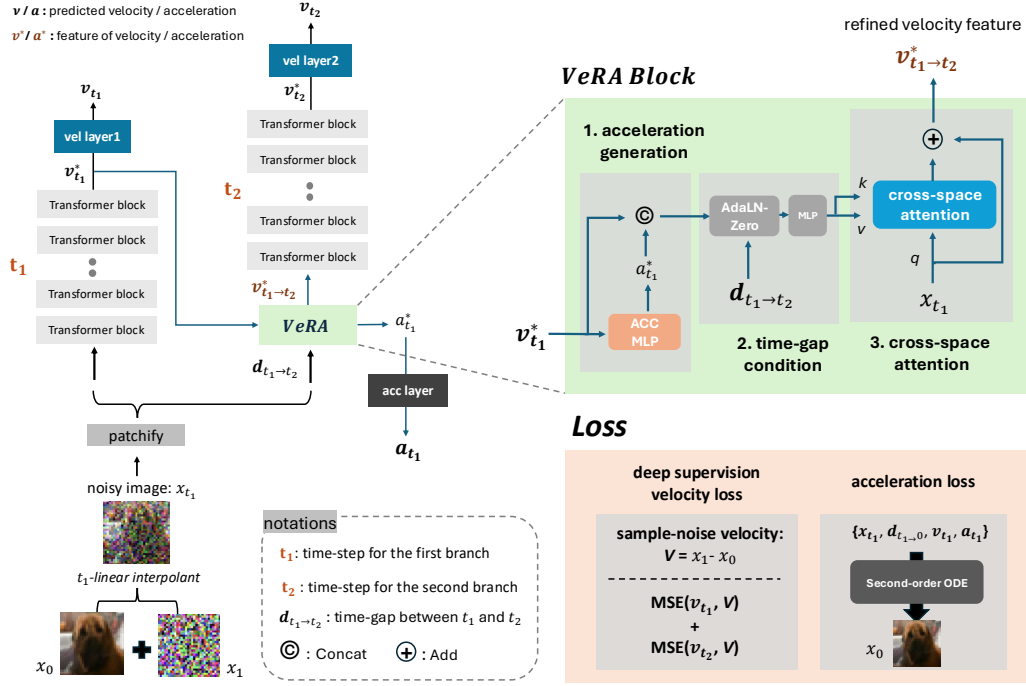


Figure 3. **DeepFlow Architecture.** We introduce advanced deep supervision by partitioning transformer blocks into equal-sized branches and employing multiple velocity layers (dark blue boxes), enabling each branch to predict velocity at a distinct time-step. Then, **VeRA Block** is inserted between adjacent branches for explicit feature refinement. It consists of three sub-blocks: 1. **acceleration generation**: we design a simple MLP (**ACC MLP**) to generate acceleration feature. It is trained with acceleration loss using second-order ODE function (Equation (7)). Meanwhile, we concatenate previous velocity feature and computed acceleration feature for following sub-block. 2. **time-gap condition**: we modulate concatenated features using **AdaLN-Zero** layer conditioned by time-gap. 3. **cross-space attention**: we design a novel cross-attention that integrates two features from different spaces, modulated velocity features from temporal dynamics and spatial features from original patchified image. By leveraging time-gap conditioning between branches with **VeRA block**, **DeepFlow** enhances feature alignment and ultimately improves image generation quality.

3.2.1 Deep Supervision

DeepFlow employs deep supervision [20] by inserting auxiliary velocity layers after selected intermediate transformer blocks. The corresponding deep supervision loss at these key transformer layers is defined as follows:

$$\mathcal{L}_{\text{deep}}(\theta) = \mathbb{E} \left[\sum_{i=1}^k \beta^i (\|v_{\theta}^i(\mathbf{x}_t^i, t, c) - V\|^2) \right] \quad (4)$$

Here, k denotes the number of key layers, v^i indicates the velocity prediction from the i^{th} velocity layer following the corresponding transformer branch, and \mathbf{x}_t^i denotes input features for the i^{th} branch (defined as \mathbf{x}_t for $i=1$, and velocity features from the previous branch otherwise). β^i represents the deep supervision coefficient. This loss encourages each velocity layer to produce outputs that closely match the target V . This design enables our DeepFlow to support various configurations. For example, DeepFlow- $\{k\}$ T denotes a

variant in which the transformer blocks are divided into k equal-sized Transformer branches, with each branch concluding with a velocity layer to facilitate deep supervision. For simplicity, we set k as 2 for following explanation about VeRA block.

3.2.2 VeRA Block

To enhance cross-layer deep supervision and improve feature alignment across different branches, we introduce a novel lightweight module called Velocity Refiner with Acceleration (VeRA) block. This module explicitly aligns the deeply supervised velocity features between consecutive branches. We provide a step-by-step explanation of its overall architecture, as illustrated in Figure 3.

Branch Conditioned with Different Time-step.

During training, we deliberately differentiate time-steps conditioned on adjacent branches. Two branches correspond to two different time-steps as illustrated in Fig-

ure 3 with \mathbf{t}_1 and \mathbf{t}_2 , which effectively enables inserted VeRA block to be trained with second-order dynamics using time-gap. We first transform Equation (4) for deep supervision as below to train each branch with its corresponding time-step.

$$\mathcal{L}_{\text{deep}^*} = \mathbb{E} \left[\sum_{i=1}^2 \beta^i (\|v_{\theta}^i(\mathbf{x}_t^i, \mathbf{t}_i, c) - V\|^2) \right], \quad (5)$$

$$\mathbf{v}_{t_1} = v_{\theta}^1(\mathbf{x}_{t_1}^1, \mathbf{t}_1, c), \quad \mathbf{v}_{t_2} = v_{\theta}^2(\mathbf{x}_{t_2}^2, \mathbf{t}_2, c),$$

where \mathbf{x}_t^1 and \mathbf{x}_t^2 indicate initial noisy image \mathbf{x}_{t_1} and previous velocity feature $\mathbf{v}_{t_1}^*$, respectively. It highlights that different time-steps are used for conditioning different branches to generate distinctive velocities.

The VeRA Block Architecture. The VeRA block is strategically placed between consecutive branches, refining deeply-supervised velocity features for use in the subsequent branch. The key operations include:

Acceleration Learning via Second-Order ODE Training:

The primary goal of the VeRA block is to refine previous velocity features by incorporating acceleration information. To achieve this, we introduce a simple MLP block, termed ACC_MLP, which projects the previous velocity feature ($\mathbf{v}_{t_1}^*$) to a higher dimension and then back to the original dimension, producing acceleration feature $\mathbf{a}_{t_1}^*$.

$$\mathbf{a}_{t_1}^* = \text{ACC_MLP}(\mathbf{v}_{t_1}^*) \quad (6)$$

Then, we can endow $\mathbf{a}_{t_1}^*$ with acceleration property using a second-order ordinary differential equation (2^{nd}-ODE) as following equation:

$$L_{\text{acc}} = \mathbb{E} \|2^{\text{nd}}\text{-ODE}(\mathbf{x}_{t_1}, \mathbf{v}_{t_1}, \mathbf{a}_{t_1}, d_{t_1 \rightarrow 0}) - \mathbf{x}_0\|^2, \quad (7)$$

$$2^{\text{nd}}\text{-ODE} = \mathbf{x}_{t_1} + \mathbf{v}_{t_1} \odot d_{t_1 \rightarrow 0} + \frac{1}{2} \mathbf{a}_{t_1} \odot (d_{t_1 \rightarrow 0})^2,$$

where \mathbf{v}_{t_1} and \mathbf{a}_{t_1} are outputs of velocity and acceleration layers from $\mathbf{v}_{t_1}^*$ and $\mathbf{a}_{t_1}^*$, reducing their dimension to image space. $d_{t_1 \rightarrow 0}$ is time gap between time-steps of t_1 and 0. \odot indicates Hadamard product for element-wise matrix multiplication. In this setup, the acceleration is learned in such a way that it aligns closely with the clean image representation (\mathbf{x}_0).

Feature Concatenation and Time-gap Conditioning:

After computing the acceleration features ($\mathbf{a}_{t_1}^*$), we concatenate these with the original velocity features ($\mathbf{v}_{t_1}^*$). To enable this concatenated feature to be aware of time-gap, we apply a time-gap-conditioned adaptive layer normalization [32] with a following MLP as below:

$$\text{modulate}(\mathbf{v}_{t_1}^*) = \text{MLP}(\text{AdaLN-Zero}(\text{concat}(\mathbf{v}_{t_1}^*, \mathbf{a}_{t_1}^*), T(d_{t_1 \rightarrow t_2}))) \quad (8)$$

$d_{t_1 \rightarrow t_2}$ denotes the time gap between t_1 and t_2 and passes through time embedder T . This dynamically modulates the concatenated feature statistics based on the time difference, which steps forward refined velocity feature.

Spatial Information Integration via Cross-Attention:

Beyond feature alignment with temporal property using different time-steps, the VeRA block also integrates spatial context by employing a cross-attention (CA) mechanism. This mechanism facilitates interaction between two spaces: modulated velocity feature space from previous step and spatial feature space from an original patchified image as noted in following equation.

$$\mathbf{v}_{t_1 \rightarrow t_2}^* = \text{CA}(\text{modulate}(\mathbf{v}_{t_1}^*), \mathbf{x}_{t_1}), \quad (9)$$

where $\text{modulate}(\mathbf{v}_{t_1}^*)$ is used for key and value, while \mathbf{x}_{t_1} is used for query. This approach using cross-space attention effectively highlights pertinent spatial features that may have been underrepresented in the pure temporal transformation, leading to final refined velocity feature.

In summary, DeepFlow incorporates deep supervision and VeRA block, which together enable training to align the internal velocity representation with Equation (5) and Equation (7), as described below:

$$L_{\text{total}} = L_{\text{deep}^*} + \lambda L_{\text{acc}}, \quad (10)$$

where λ is a hyperparameter that balances the deep-supervised velocity loss and the acceleration loss.

4. Experiments

In this section, we demonstrate the effectiveness of DeepFlow through extensive experiments. Section 4.1 details the implementation of DeepFlow. Section 4.2 presents the main results of DeepFlow on class-conditional image generation and text-to-image generation. Finally, comprehensive ablation studies are provided in Section 4.3.

4.1. Implementation Details

model	depth	hidden dim	key layer indices
DeepFlow-B/2-2T	12	768	{6, 12}
DeepFlow-L/2-3T	18	1024	{6, 12, 18}
DeepFlow-XL/2-3T	24	1152	{8, 16, 24}

Table 1. We propose three variants of DeepFlow (B, L, and XL), each differing in depth (*i.e.*, the number of blocks) and channel dimensions. DeepFlow- $\{k\}$ T denotes a model with k key layers, where deep supervision is applied. These key layers are evenly distributed across the Transformer blocks.

Our framework is implemented in PyTorch [31] and closely follows the flow matching and transformer setup

model	epoch	SSL align	sampling	FID↓	sFID↓	IS↑
SiT-B/2 [28]	80	✗	uniform	34.4	6.5	43.7
	80	✗	lognormal	29.7	6.2	51.0
	80	DINOv1	uniform	28.3	6.6	53.8
	80	DINOv1	lognormal	24.4	6.3	62.1
	80	DINOv2	uniform	23.0	6.4	59.9
	80	DINOv2	lognormal	20.4	6.3	72.7
DeepFlow-B/2-2T	80	✗	uniform	28.1 (-6.3)	5.8 (-0.7)	49.8 (+6.1)
	80	✗	lognormal	23.1 (-6.6)	5.6 (-0.6)	60.3 (+9.3)
	80	DINOv1	uniform	25.6 (-2.7)	6.4 (-0.2)	56.6 (+2.8)
	80	DINOv1	lognormal	21.7 (-2.7)	6.1 (-0.2)	66.0 (+3.9)
	80	DINOv2	uniform	20.0 (-3.0)	6.2 (-0.2)	70.2 (+10.3)
	80	DINOv2	lognormal	17.2 (-3.2)	6.0 (-0.3)	77.8 (+5.1)

Table 2. Quantitative comparisons of flow-based generative models under the **Base** configuration for class-conditional image generation on ImageNet-256×256 without classifier-free guidance [15]. DeepFlow-B/2-2T consistently outperforms SiT-B/2 across various settings, whether using SSL alignment [52] or not, and under both uniform and lognormal sampling strategies.

introduced in SiT [28]. Specifically, we extract patchified features from raw images using the VAE [19] encoder pretrained on Stable Diffusion [38]. For the transformer blocks, we divide them into equal-sized branches to enable deep supervision, with VeRA blocks inserted at key layers (see Table 1 for model configuration of DeepFlow). We explain more about used hyperparameters and implementation details in Appendix A.3.

DeepFlow Training & Evaluation. For training VeRA block in DeepFlow- $\{k\}$ T, we employ k distinct time-steps while constraining the maximum gap between consecutive branches to α . We empirically observed that assigning a low β value (*e.g.*, 0.2) to intermediate velocity predictions while maintaining a β of 1.0 for final layer during training leads to improved generation quality. During evaluation, we condition all transformer layers on a single time-step, which still enables the refinement of velocity features across adjacent branches. We adopt basic training recipes and sampling strategies from SiT [28] and REPA [52] for fair comparisons.

4.2. Main Results

4.2.1 Class-conditional Image Generation

ImageNet-1k [6] is widely used for class-conditional image generation benchmark. We provide the detail of its usage for training and sampling in Appendix A.7.

Comparison with Flow-based Models. Using ImageNet-256×256, we compare our DeepFlow with the representative flow-based model SiT [28] in Table 2 and Table 3, evaluated under the Base and XLarge configurations, respectively. In Table 2, we consider two key comparison criteria: (i) the use of external self-supervised alignment (*SSL align*), which leverages DINO_{v1} or DINO_{v2} as introduced in REPA [52], versus no SSL alignment; (ii) the strategy for sampling time-steps during training, comparing *uniform* sampling with

model	epoch	SSL align	sampling	FID↓	sFID↓	IS↑
SiT-XL/2 [28]	80	✗	lognormal	13.8	5.0	91.0
	200	✗	lognormal	10.2	5.3	113.6
	800	✗	lognormal	9.8	7.3	128.2
	80	DINOv2	lognormal	7.2	5.1	134.3
	200	DINOv2	lognormal	6.2	5.3	149.6
	800	DINOv2	lognormal	5.7	6.4	171.0
DeepFlow-XL/2-3T	80	✗	lognormal	10.3 (-3.4)	4.8 (-0.2)	105.2 (+14.2)
	200	✗	lognormal	7.8 (-2.4)	4.8 (-0.5)	127.3 (+13.7)
	400	✗	lognormal	7.2	5.1	138.5
	80	DINOv2	lognormal	6.5 (-0.7)	4.9 (-0.2)	134.9 (+0.6)
	200	DINOv2	lognormal	5.4 (-0.8)	5.0 (-0.3)	151.6 (+2.0)
	400	DINOv2	lognormal	5.0	5.2	162.0

Table 3. Quantitative comparisons of flow-based generative models under the **XLarge** configuration for class-conditional image generation on ImageNet-256×256 without classifier-free guidance [15]. Our DeepFlow-XL/2-3T not only consistently delivers superior image generation quality compared to SiT-XL/2 when trained for an equivalent number of epochs, but it also converges significantly faster, requiring only half the training epochs (800 epochs → 400 epochs).




Figure 4. **Qualitative Comparison in Different Epochs.** Images are generated from models trained in different epochs. DeepFlow-XL/2-3T converges faster than SiT-XL/2 and produces high-quality samples even with fewer training epochs.

lognormal sampling, as also evaluated in SD3 [8]. Without SSL align, DeepFlow-B/2-2T consistently outperforms SiT under both uniform and lognormal sampling strategies with a healthy margin, lowering 6.3 FID and 6.6 FID, respectively. When paired with SSL alignment using either DINO_{v1} or DINO_{v2}, DeepFlow-B/2-2T further reduces the FID by an average of 3.0 points compared to SiT-B/2, achieving a remarkable generation quality with FID as low as 17.2. More surprisingly, DeepFlow-B/2-2T without SSL align can achieve comparable or better performance than SiT with SSL align from DINO_{v1}² — for example, 28.3 FID *vs.* 28.1 FID in uniform sampling, 24.4 FID *vs.* 23.3 FID in lognormal sampling), demonstrating the effectiveness of DeepFlow in obviating the need for external feature alignment. Experiments using the XLarge configuration (see Table 3)

²DINO_{v1} is used for this comparison since it was pretrained with ImageNet-1k, while DINO_{v2} was pretrained on the additional dataset, LVD-142M [30].

ImageNet-256×256	epoch	params	FID↓	sFID↓	IS↑
<i>Autoregressive Models</i>					
LlamaGen-3B [43]	50	3.1B	3.05	-	222.3
Open-MAGViT2-XL [27]	350	1.5B	2.33	-	271.8
<i>Masked Generative Models</i>					
TiTok-S-128 [51]	200	287M	1.97	-	281.8
MAGViT-v2 [49]	270	307M	1.78	-	319.4
<i>Diffusion-based Models</i>					
ADM-U [7]	400	554M	3.94	6.14	186.7
Simple diffusion [16]	800	-	2.77	-	211.8
DiT-XL/2 [32]	1400	675M	2.27	4.60	278.2
<i>Flow-based Models</i>					
SiT-XL/2 [28]	1400	675M	2.06	4.50	270.3
DeepFlow-XL/2-3T	400	681M	1.97	4.39	264.7
DeepFlow-XL/2-3T	600	681M	1.89	4.40	263.3
SiT-XL/2 + SSL align [52]	800	675M	1.80	4.50	284.0
DeepFlow-XL/2-3T + SSL align	400	681M	1.77	4.44	271.3

Table 4. Comparison with state-of-the-art models in class-conditional image generation on ImageNet 256×256.

reveal two key observations. First, DeepFlow-XL/2-3T consistently outperforms SiT-XL/2—whether using SSL alignment or not—under lognormal sampling at the same training epochs (*e.g.*, 80 and 200 epochs). Second, DeepFlow-XL/2-3T converges significantly faster than SiT-XL/2. For example, training DeepFlow-XL/2-3T for just 100 epochs without SSL alignment yields a 9.8 FID, matching the performance of SiT-XL/2 trained for 800 epochs. Furthermore, with 400 epochs, DeepFlow-XL/2-3T surpasses SiT-XL/2’s 800-epoch results. As illustrated in Figure 4, our model also produces higher-quality images with fewer training epochs.

Comparison with state-of-the-art Models. In addition to flow-based comparisons, we compare against state-of-the-art image generators—including autoregressive [43, 27, 50], masked generative [49, 51, 46, 18], diffusion-based [7, 16, 32], and flow-based models [28, 52]—using classifier-free guidance [15]. As shown in Table 4 and Table 5, DeepFlow-XL/2-3T not only significantly outperforms existing flow-based models but also achieves competitive or superior performance compared to other state-of-the-art generators, all while training efficiently. For instance, DeepFlow-XL/2-3T delivers superior results on ImageNet-256 in only 400 epochs, and on ImageNet-512, it outperforms the corresponding SiT model with fewer epochs.

4.2.2 Text-to-Image Generation

Following REPA [52], we modify the architecture of MMDiT [8] by incorporating a flow matching objective, seamlessly integrating DeepFlow into MMDiT’s architecture. Both our model and SiT (MMDiT + flow matching) are trained on the MS-COCO [21, 5] training set with a hidden dimension of 768 and 24 transformer layers while sampling time-steps from uniform

ImageNet-512×512	epoch	params	FID↓	sFID↓	IS↑
<i>Autoregressive Models</i>					
GIVT-Causal-L+A [44]	500	1.7B	2.92	-	-
<i>Masked Generative Models</i>					
TiTok-B-128 [51]	200	177M	2.13	-	-
MAGViT-v2 [49]	270	307M	1.91	-	324.3
<i>Diffusion-based Models</i>					
LDM-4-G [38]	167	400M	3.60	-	247.7
DiT-XL/2 [32]	600	676M	3.04	-	-
<i>Flow-based Models</i>					
SiT-XL/2 [28]	600	676M	2.62	4.18	252.2
DeepFlow-XL/2-3T	200	682M	2.59	4.23	231.1
DeepFlow-XL/2-3T	400	682M	2.22	4.27	249.4
SiT-XL/2 + SSL align [52]	200	676M	2.08	4.19	274.6
DeepFlow-XL-3T + SSL align	160	682M	2.03	4.29	248.8
DeepFlow-XL-3T + SSL align	200	682M	1.96	4.28	260.3

Table 5. Comparison with state-of-the-art models in class-conditional image generation on ImageNet 512×512.

model	SSL align	MS-COCO Eval				GenEval Overall↑
		FID↓	FD _{DINOv2} ↓	IS↑	CLIP↑	
SiT [28]	✗	5.38	408.27	31.1	0.2954	0.2672
	DINOv2	4.24	274.19	34.4	0.2979	0.3166
DeepFlow-3T	✗	5.32	343.60	33.0	0.2974	0.2957
	DINOv2	4.18	216.20	35.9	0.3010	0.3458

Table 6. Quantitative comparison on text-to-image generation. To ensure a fair comparison, both SiT [28] and DeepFlow utilize 24 layers of transformer blocks and use same MS-COCO [21] training set. Then, they are both compared for MS-COCO evaluation and GenEval [10] benchmark (please refer to category-level performance of GenEval in Appendix). We reproduced MS-COCO evaluation performance of SiT variants.

distribution. We extensively conducted evaluation on MS-COCO validation set and GenEval [10] benchmark. For MS-COCO evaluation, we observe that DeepFlow outperforms in all of the metrics (FID, FD_{DINOv2} [42], IS, CLIP score). Furthermore, we show that DeepFlow significantly increase overall score of GenEval compared to SiT (check Appendix A.5. for category-level performance in GenEval benchmark).

4.3. Ablation Studies

4.3.1 DeepFlow Components

Table 7a summarizes the incremental improvements of DeepFlow over the baseline SiT-B/2 [28] (FID 34.4) using uniform sampling. Deep supervision on each branch reduces FID to 33.0 by aligning intermediate velocity features with the ground-truth. Adding a time-gap mechanism further lowers FID to 31.1 by enhancing robustness to time-step variations. An inter-layer acceleration pipeline within the VeRA block decreases FID to 29.9, and incorporating a cross-space attention module refines spatial fusion to achieve a final FID of 28.1. Overall, these enhancements yield a total FID improvement of 6.3, demonstrating the complementary

model	FID↓
SiT-B/2 [28]	34.4
+ deep supervision	33.0
+ time-gap	31.1
+ inter-layer acceleration	29.9
+ cross-space attention	28.1

model	α	FID↓
DeepFlow-B/2-2T	0.1	28.9
	0.05	28.6
	0.01	28.1
	0.005	28.3
	0.001	28.5

model	λ	ACC_MLP	FID↓
DeepFlow-B/2-2T	1.0	{2048, 768}	30.4
	0.5		27.8
	0.75		28.1
	1.0	{2048, 4096, 2048, 768}	28.1
	1.25		28.5
DeepFlow-B/2-2T*	1.5		28.8
	0.5	{2048, 4096, 2048, 768}	23.5
	1.0		23.1

(a) **Component Ablation:** study on effectiveness of components for DeepFlow.

(b) **Time-gap:** ablation on maximum time-gap (α) between adjacent branches used during training.

(c) **Acceleration Design:** ablation on λ and number of channels for ACC_MLP. * indicates the model trained with lognormal sampling of time-steps.

Table 7. DeepFlow Ablations.

model	depth	params	GFLOPS	FID↓
SiT-B/2 [28]	12	130M	24	29.7
SiT-L/2 [28]	24	458M	80	16.1
SiT-XL/2 [28]	28	675M	114	13.7
DeepFlow-B/2-2T	10	144M	24	27.3
	12	166M	28	23.1
DeepFlow-L/2-2T	20	431M	72	13.9
	22	469M	78	13.3
	24	507M	84	12.8
DeepFlow-XL/2-2T	22	588M	97	11.9
	24	636M	106	11.7
	26	683M	114	11.1
	28	731M	122	11.1
DeepFlow-L/2-3T	18	433M	72	13.3
	21	490M	82	12.0
	24	547M	92	11.9
DeepFlow-XL/2-3T	18	538M	89	12.4
	21	609M	101	10.9
	24	681M	113	10.3
	27	753M	125	10.0
DeepFlow-XL/2-4T	28	822M	137	9.7

Table 8. Tradeoff between Efficiency and Performance of DeepFlow. All of the experiments, including the reproduced, better baseline SiT variants, utilize lognormal sampling with 80 epochs of training in ImageNet-256×256 and are evaluated using SDE 250 steps without a classifier-free guidance [15]. Best viewed in color.

benefits of each component in our framework.

4.3.2 Time-gap

Table 7b presents the impact of varying the maximum time-gap α between adjacent branches during training for DeepFlow. When α is set to 0.01, the model achieves its best FID of 28.1 under the base configuration. Additionally, we observe that the model’s performance remains relatively stable across different values of α , indicating that DeepFlow is not highly sensitive to variations in the time-gap parameter.

4.3.3 Acceleration Design

In Table 7c, we explore how varying the balancing coefficient λ used in Equation (10) and the architecture of the acceleration MLP (ACC_MLP) affect performance of DeepFlow. Considering the optimal λ that works both

in uniform sampling and lognormal sampling (marked with an asterisk), we choose λ as 1.0. We can also find out that expanding the ACC_MLP capacity with deeper layers (the number of channels for ACC_MLP: {2048, 4096, 2048, 768}) can increase the capability of acceleration, yielding further performance gains.

4.3.4 Efficiency vs. Performance

Table 8 compares various DeepFlow configurations (DeepFlow-B/2, DeepFlow-L/2, DeepFlow-XL/2) with baseline SiT models in terms of transformer depth, parameter count, GFLOPS, and FID. All models are trained for 80 epochs with lognormal sampling and evaluated using SDE (250 steps) on the ImageNet-256×256 benchmark without a classifier-free guidance. Two key observations emerge: (i) **Efficient Performance Scaling:** Even smaller DeepFlow variants match larger SiT models. For example, DeepFlow-L/2-3T (18 layers, 433M parameters, 72 GFLOPS) achieves an FID of 13.3, comparable to SiT-XL/2 (675M parameters, 114 GFLOPS, 13.7 FID). (ii) **Superior Performance:** DeepFlow variants consistently outperform comparable SiT models. DeepFlow-XL/2-3T reduces FID by 3.4 points relative to SiT-XL/2, while maintaining a similar computational cost. Moreover, DeepFlow scales effectively, as DeepFlow-XL/2-4T further improves performance down to 9.7 FID.

5. Conclusion

We introduced DeepFlow, a novel flow-based generative model that enhances internal velocity representations via deep supervision and explicit feature alignment using proposed VeRA block. Our extensive experiments show that DeepFlow not only dramatically improves training efficiency of flow-based models while achieving competitive performance on various image generation benchmarks. We believe this work lays a strong foundation for efficient and high performing flow-based models.

References

- [1] 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. **1**
- [2] 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 *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2023. **3**
- [3] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proc. of Int’l Conf. on Computer Vision (ICCV)*, 2021. **2, 3**
- [4] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Proc. of Neural Information Processing Systems (NeurIPS)*, 2018. **1, 3**
- [5] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. **7**
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2009. **3, 6**
- [7] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *Proc. of Neural Information Processing Systems (NeurIPS)*, 2021. **7**
- [8] 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 *Proc. of Int’l Conf. on Machine Learning (ICML)*, 2024. **1, 6, 7**
- [9] Peng Gao, Le Zhuo, Ziyi Lin, Chris Liu, Junsong Chen, Ruoyi Du, Enze Xie, Xu Luo, Longtian Qiu, Yuhang Zhang, et al. Lumina-t2x: Scalable flow-based large diffusion transformer for flexible resolution generation. In *Proc. of Int’l Conf. on Learning Representations (ICLR)*, 2025. **1**
- [10] Dhruba Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. In *Proc. of Neural Information Processing Systems (NeurIPS)*, 2024. **3, 7**
- [11] Ali Hatamizadeh, Jiaming Song, Guilin Liu, Jan Kautz, and Arash Vahdat. Diffit: Diffusion vision transformers for image generation. In *Proc. of European Conf. on Computer Vision (ECCV)*, 2024. **3**
- [12] Ju He, Qihang Yu, Qihao Liu, and Liang-Chieh Chen. Flowtok: Flowing seamlessly across text and image tokens. *arXiv preprint arXiv:2503.10772*, 2025. **3**
- [13] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2022. **3**
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proc. of Neural Information Processing Systems (NeurIPS)*, 2020. **1, 3**
- [15] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. **3, 6, 7, 8**
- [16] Emiel Hoogeboom, Jonathan Heek, and Tim Salimans. simple diffusion: End-to-end diffusion for high resolution images. In *Proc. of Int’l Conf. on Machine Learning (ICML)*, 2023. **7**
- [17] Yang Jin, Zhicheng Sun, Ningyuan Li, Kun Xu, Hao Jiang, Nan Zhuang, Quzhe Huang, Yang Song, Yadong Mu, and Zhouchen Lin. Pyramidal flow matching for efficient video generative modeling. *arXiv preprint arXiv:2410.05954*, 2024. **1**
- [18] Dongwon Kim, Ju He, Qihang Yu, Chenglin Yang, Xiaohui Shen, Suha Kwak, and Liang-Chieh Chen. Democratizing text-to-image masked generative models with compact text-aware one-dimensional tokens. *arXiv preprint arXiv:2501.07730*, 2025. **7**
- [19] Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. **3, 6**
- [20] Chen-Yu Lee, Saining Xie, Patrick Gallagher, Zhengyou Zhang, and Zhuowen Tu. Deeply-supervised nets. In *Proc. of Artificial Intelligence and Statistics (AISTATS)*, 2015. **2, 3, 4**
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Proc. of European Conf. on Computer Vision (ECCV)*, 2014. **3, 7**
- [22] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *Proc. of Int’l Conf. on Learning Representations (ICLR)*, 2023. **1, 3**
- [23] Yaron Lipman, Marton Havasi, Peter Holderrieth, Neta Shaul, Matt Le, Brian Karrer, Ricky TQ Chen, David Lopez-Paz, Heli Ben-Hamu, and Itai Gat. Flow matching guide and code. *arXiv preprint arXiv:2412.06264*, 2024. **1**
- [24] Qihao Liu, Zhanpeng Zeng, Ju He, Qihang Yu, Xiaohui Shen, and Liang-Chieh Chen. Alleviating distortion in image generation via multi-resolution diffusion models. In *Proc. of Neural Information Processing Systems (NeurIPS)*, 2024. **3**
- [25] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *Proc. of Int’l Conf. on Learning Representations (ICLR)*, 2023. **1**
- [26] Xingchao Liu, Xiwen Zhang, Jianzhu Ma, Jian Peng, et al. InstafLOW: One step is enough for high-quality diffusion-based text-to-image generation. In *Proc. of Int’l Conf. on Learning Representations (ICLR)*, 2023. **1**
- [27] Zhuoyan Luo, Fengyuan Shi, Yixiao Ge, Yujiu Yang, Limin Wang, and Ying Shan. Open-magvit2: An open-source project toward democratizing auto-regressive

- visual generation. *arXiv preprint arXiv:2409.04410*, 2024. 7
- [28] 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 *Proc. of European Conf. on Computer Vision (ECCV)*, 2024. 1, 3, 6, 7, 8
- [29] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *Proc. of Int'l Conf. on Machine Learning (ICML)*, 2021. 3
- [30] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 2, 3, 6
- [31] 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. In *Proc. of Neural Information Processing Systems (NeurIPS)*, 2019. 5
- [32] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proc. of Int'l Conf. on Computer Vision (ICCV)*, 2023. 3, 5, 7
- [33] Pablo Pernias, Dominic Rampas, Mats Leon Richter, Christopher Pal, and Marc Aubreville. Würstchen: An efficient architecture for large-scale text-to-image diffusion models. In *Proc. of Int'l Conf. on Learning Representations (ICLR)*, 2024. 2
- [34] Sucheng Ren, Qihang Yu, Ju He, Xiaohui Shen, Alan Yuille, and Liang-Chieh Chen. Flowar: Scale-wise autoregressive image generation meets flow matching. *arXiv preprint arXiv:2412.15205*, 2024. 3
- [35] Sucheng Ren, Qihang Yu, Ju He, Xiaohui Shen, Alan Yuille, and Liang-Chieh Chen. Beyond next-token: Next-x prediction for autoregressive visual generation. *arXiv preprint arXiv:2502.20388*, 2025. 3
- [36] Sucheng Ren, Qihang Yu, Ju He, Alan Yuille, and Liang-Chieh Chen. Grouping first, attending smartly: Training-free acceleration for diffusion transformers. *arXiv preprint arXiv:2505.14687*, 2025. 3
- [37] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *Proc. of Int'l Conf. on Machine Learning (ICML)*, 2015. 1, 3
- [38] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 3, 6, 7
- [39] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Proc. of International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2015. 3
- [40] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. In *Proc. of Neural Information Processing Systems (NeurIPS)*, 2020. 1
- [41] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *Proc. of Int'l Conf. on Learning Representations (ICLR)*, 2021. 1, 3
- [42] George Stein, Jesse Cresswell, Rasa Hosseinzadeh, Yi Sui, Brendan Ross, Valentin Vilecroze, Zhaoyan Liu, Anthony L Caterini, Eric Taylor, and Gabriel Loaizaganem. Exposing flaws of generative model evaluation metrics and their unfair treatment of diffusion models. In *Proc. of Neural Information Processing Systems (NeurIPS)*, 2023. 7
- [43] Peize Sun, Yi Jiang, Shoufa Chen, Shilong Zhang, Bingyue Peng, Ping Luo, and Zehuan Yuan. Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv preprint arXiv:2406.06525*, 2024. 7
- [44] Michael Tschannen, Cian Eastwood, and Fabian Mentzer. Givit: Generative infinite-vocabulary transformers. In *Proc. of European Conf. on Computer Vision (ECCV)*, 2025. 7
- [45] A Vaswani. Attention is all you need. In *Proc. of Neural Information Processing Systems (NeurIPS)*, 2017. 1, 3
- [46] Mark Weber, Lijun Yu, Qihang Yu, Xueqing Deng, Xiaohui Shen, Daniel Cremers, and Liang-Chieh Chen. Maskbit: Embedding-free image generation via bit tokens. *arXiv preprint arXiv:2409.16211*, 2024. 7
- [47] Chenglin Yang, Celong Liu, Xueqing Deng, Dongwon Kim, Xing Mei, Xiaohui Shen, and Liang-Chieh Chen. 1.58-bit flux. *arXiv preprint arXiv:2412.18653*, 2024. 3
- [48] Jingfeng Yao and Xinggang Wang. Reconstruction vs. generation: Taming optimization dilemma in latent diffusion models. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2025. 3
- [49] Lijun Yu, José Lezama, Nitesh B Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Vighnesh Birodkar, Agrim Gupta, Xiuye Gu, et al. Language model beats diffusion—tokenizer is key to visual generation. In *Proc. of Int'l Conf. on Learning Representations (ICLR)*, 2024. 7
- [50] Qihang Yu, Ju He, Xueqing Deng, Xiaohui Shen, and Liang-Chieh Chen. Randomized autoregressive visual generation. *arXiv preprint arXiv:2411.00776*, 2024. 7
- [51] Qihang Yu, Mark Weber, Xueqing Deng, Xiaohui Shen, Daniel Cremers, and Liang-Chieh Chen. An image is worth 32 tokens for reconstruction and generation. In *Proc. of Neural Information Processing Systems (NeurIPS)*, 2024. 7
- [52] Sihyun Yu, Sangkyung Kwak, Huiwon Jang, Jongheon Jeong, Jonathan Huang, Jinwoo Shin, and Saining Xie. Representation alignment for generation: Training diffusion transformers is easier than you think. In *Proc. of Int'l Conf. on Learning Representations (ICLR)*, 2025. 1, 2, 3, 6, 7