arXiv:2610.23854v1 [cs.CV] 16 Oct 2026

RETHINKING THE DESIGN SPACE OF VANILLA PIXEL DIFFUSION TRANSFORMERS

Soumik Mukhopadhyay[1,2*] **Long Mai**[2] **Abhinav Shrivastava**[1] **Aniruddha Mahapatra**[2†]
[1] University of Maryland, College Park [2] Adobe Research



Figure 1: Sample 512px text-to-image generations with our MFLOW-XXL/16 model.

ABSTRACT

Pixel-space diffusion models have recently re-emerged as a compelling alternative to latent diffusion, avoiding the lossy latent bottleneck and separately trained autoencoder. JiT shows that a vanilla diffusion transformer can generate directly in pixel space by replacing the commonly used velocity prediction ($\boldsymbol{u}$) with clean image prediction ($\boldsymbol{x}_0$). In this work, we revisit this choice and introduce MFLOW, an adaptive parameterization that learns to interpolate between $\boldsymbol{x}_0$ and $\boldsymbol{u}$ prediction across principal directions of the data, which can be learned end-to-end with diffusion model training. We find that the preferred parameterization varies systematically with the data representation and model capacity. Notably, increasing model capacity shifts preference toward $\boldsymbol{u}$ prediction from $\boldsymbol{x}_0$ prediction. We evaluate MFLOW across ImageNet, UCF-101 video, and large-scale text-to-image generation, demonstrating competitive or improved generative performance across diverse settings. In addition, we find that vanilla pixel diffusion transformers, irrespective of parameterization, retain a persistent high-frequency deficit relative to real images. To address this, we introduce *ConvRefiner*, a lightweight convolutional module trained post-hoc with adversarial loss while keeping the diffusion backbone frozen. With only a small amount of additional training, *ConvRefiner* restores the high-frequency statistics of generated images while preserving FID. Project webpage: soumik-kanad.github.io/mflow-diffusion.

1 INTRODUCTION

Diffusion and flow-based models have become a leading paradigm for visual generation in recent years. Modern versions of these models predominantly operate in the latent space, obtained from compressing the pixel-space data with a separately-trained autoencoder. However, latent diffusion

*Work done while an intern at Adobe Research.
†Project lead.

models (LDMs) (Rombach et al. (2022); Peebles & Xie (2023); Ma et al. (2024); Esser et al. (2024); Yu et al. (2025); Singh et al. (2026)) introduce their own limitations – a lossy bottleneck leads to unfaithful reconstruction after compression, and there is a need for a separately trained compression autoencoder, which is typically unaware of the downstream generative objective.

Recent works (Li & He (2026); Chen et al. (2026b); Yu et al. (2026); Mukhopadhyay et al. (2026); Chen et al. (2025b); Ma et al. (2026b;c); Wang et al. (2026c)) have revived diffusion directly in pixel space for image generation, removing the aforementioned drawbacks in LDMs. In particular, JiT (Li & He (2026)), demonstrates that vanilla diffusion transformers can achieve competitive generative performance directly on raw pixels, eliminating the need for pretrained autoencoders or complicated architectural modifications. Its aggressive patchification keeps the sequence length manageable, *e.g.*, its 16×16 patchification produces equivalent number of tokens as a $16\times$ spatial compression VAE followed by 1×1 latent patchification. At the same time, such aggressive patchification poses a challenge in pixel-space diffusion as each token must model a high-dimensional block of raw pixels, raising two questions: *what prediction parameterization is best suited to such high-dimensional patches*, and *can a vanilla large-patch transformer faithfully model fine-scale image details?* In this work, we aim to improve the pixel-space diffusion paradigm along these axes.

JiT addresses the first question with clean-image ($\boldsymbol{x}_0$) prediction. Under the manifold assumption, clean data lies near a low-dimensional manifold, whereas noise and velocity span the ambient space, making $\boldsymbol{x}_0$ easier for an undercomplete network to predict. However, sampling requires converting $\boldsymbol{x}_0$ to $\boldsymbol{u}$ (velocity), which becomes ill-conditioned near the low-noise regime. AsymFlow (Chen et al. (2026a)) reduces this problem by predicting $\boldsymbol{u}$ in a selected rank-r PCA subspace and $\boldsymbol{x}_0$ in its complement, keeping the prediction target effectively in a low dimensionality. However, this hard binary cutoff rank must be chosen for each setting through expensive training sweeps, *e.g.*, AsymFlow needed 6 trainings to tune the rank-r. Where as Jin & Wang (2026) use a learnable but globally shared interpolation between $\boldsymbol{x}_0$ and $\boldsymbol{u}$, not reducing the effective noise dimensionality.

We introduce **MFLOW**, an adaptive parameterization with a continuous gain $m_i \in [0, 1]$ for every principal direction of data. Notably, this differentiable velocity-recovery formulation lets us jointly train the generator and gains with the standard velocity-space flow-matching loss, letting model learn the best combination of $\boldsymbol{u}$ and $\boldsymbol{x}_0$ per principal component. MFLOW therefore replaces the AsymFlow's manually selected rank and binary mask with a continuous parameterization that adapts to the data representation, objective, and model capacity. Our experiments show that MFLOW can be adapted not only to pixel-space image generation, but also to diverse paradigms, like pixel-space video generation, text-to-image generation at multiple model capacities, and high-dimensional latents like RAEv2 (Singh et al. (2026)). Further visualization of the learned direction-wise gains as $\boldsymbol{M}$-curves, a plot of m_i's in decreasing direction variance order, reveal interesting insights, such as increasing model capacity shifts more directions toward velocity-like prediction.

Prediction parameterization alone does not resolve all limitations of vanilla pixel diffusion transformers. Irrespective of the parameterization, we find that JiT-family models exhibit a persistent high-frequency deficit relative to real images. This suggests a limitation of the minimalist large-patch design, where patch tokens are projected directly back to pixels using `MLP`-layers without an explicit local refinement pathway. We introduce ***ConvRefiner***, a lightweight residual convolutional module trained post hoc with perceptual and adversarial supervision while keeping the diffusion backbone frozen. With only 5% post-hoc training, *ConvRefiner* brings the generated high-frequency spectrum closer to that of real images while preserving FID. Overall, this work presents adaptive prediction parameterization and fine-scale signal modeling as complementary directions for improving vanilla pixel diffusion transformers.

To summarize, the key contributions are:

- We introduce MFLOW, a framework of continuous, direction-adaptive prediction parameterization whose gains are learned jointly with the generator under a standard velocity-space flow-matching objective. We show MFLOW is a general framework and can be used for diverse domains like pixel-space video generation, high-dimensional latents (like RAEv2) in addition to pixel-space image generation.
- We identify a high-frequency deficit in vanilla JiT-style models and introduce *ConvRefiner*, a lightweight post-hoc module that restores high-frequency details while preserving fidelity matching the frequency spectrum of real images.



Figure 2: **Additional Text-to-Image Results.** Sample 512px text-to-image generations with our MFLOW-XXL/16 model.

2 RELATED WORK

Pixel-space diffusion and prediction parameterization. Early diffusion models generated directly in pixel space (Ho et al. (2020); Nichol & Dhariwal (2021); Dhariwal & Nichol (2021); Karras et al. (2022)), while subsequent works improved the scalability of end-to-end pixel generation through architectural and formulation advances (Jabri et al. (2023); Hoogeboom et al. (2023; 2025); Crowson et al. (2024); Chen et al. (2025b); Mukhopadhyay et al. (2026)). Recent approaches have renewed interest in pixel-space generation using vanilla or specialized transformer architectures (Li & He (2026); Chen et al. (2026b); Yu et al. (2026); Guo et al. (2026)), with recent studies further exploring its scaling to text-to-image generation (Jiang et al. (2026); Wang et al. (2026c); Ma et al. (2026b); Yu et al. (2026); Ma et al. (2026c); Chen et al. (2025b)). In the context of our parameterization study, JiT (Li & He (2026)) advocates clean-image prediction for high-dimensional pixel modeling, an approach also adopted by recent large-scale pixel-space models such as Tuna-2 (Liu et al. (2026)). k-Diff Jin & Wang (2026) and AsymFlow Chen et al. (2026a) further explore adaptive combinations of clean-data and velocity prediction. Instead of deciding a predefined fixed formulation, which can vary across architectures and datasets, MFLOW lets the model learn the best fit between clean image and velocity prediction.

Details and frequency modeling. Several recent works study the difficulty of modeling fine-scale information in vanilla pixel-space diffusion transformers with large spatial patch sizes, such as 16×16–32×32. Existing approaches use perceptual supervision (Ma et al. (2026c)), frequency-aware objectives or flows (Degeorge et al. (2026); Pavasovic et al. (2026); Fan et al. (2026)), frequency-specialized architectures (Ma et al. (2026b;a)), or improved information pathways (Guo et al. (2026)). Concurrent work also shows that a frozen JiT model can recover additional high-frequency detail through self-guidance (Fu et al. (2026)). Although some of these approaches also improve training efficiency (Degeorge et al. (2026); Ma et al. (2026b)), they generally introduce specialized architectural components or modify the training objective, and their generated frequency statistics can still exhibit a high-frequency gap relative to real images (Degeorge et al. (2026); Ma et al. (2026b); Pavasovic et al. (2026)). In contrast, our lightweight *ConvRefiner* is trained post hoc while keeping the pretrained diffusion backbone frozen. It recovers the missing high-frequency statistics without modifying the backbone or its original training procedure, while preserving FID.

3 MODEL PREDICTION PARAMETERIZATION

3.1 MOTIVATION

Flow Matching. In the linear stochastic interpolant/rectified-flow formulation used by SiT (Ma et al. (2024)) and JiT (Li & He (2026)), for signal $\boldsymbol{x}_0 \sim p_{\text{data}}$ and noise $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$, both in $\mathbb{R}^D$, the forward interpolation at timestep t: $\boldsymbol{x}_t = t\,\boldsymbol{x}_0 + (1-t)\,\boldsymbol{\epsilon}$. The generative process is the probability-flow ODE that transports pure noise $\boldsymbol{\epsilon}$ ($t = 0$) to clean signal $\boldsymbol{x}_0$ ($t = 1$) along the conditional expectation of the interpolant's time derivative, the velocity field $\boldsymbol{u}(\boldsymbol{x}_t, t) = \mathbb{E}\left[\frac{\mathrm{d}\boldsymbol{x}_t}{\mathrm{d}t} \middle| \boldsymbol{x}_t\right] = \mathbb{E}[\boldsymbol{x}_0 - \boldsymbol{\epsilon} | \boldsymbol{x}_t]$. A network $f_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t)$ is trained by regressing a target derived from $(\boldsymbol{x}_0, \boldsymbol{\epsilon}, t)$.

Adaptive parameterization. Common non-adaptive parameterization/regression-targets include: noise $\boldsymbol{\epsilon}$-prediction, velocity $\boldsymbol{u}$-prediction , and clean signal $\boldsymbol{x}_0$-prediction. These are algebraically equivalent but differ in conditioning. EDM (Karras et al. (2022)) introduces an input/output preconditioning, where the network effectively predicts a noise-level dependent combination of clean data and noise. However, JiT (Li & He (2026)) demonstrates that such a preconditioning substantially underperforms $\boldsymbol{x}_0$-prediction for high-dimensional pixel patches using a vanilla transformer backbone, motivating the use of standard clean image prediction. More recently, k-Diff (Jin & Wang (2026)) introduces a learnable scalar interpolation between prediction targets, showing that the preferred parameterization varies across pixel and latent representations. However, its globally shared coefficient cannot account for directional differences within a representation.

AsymFlow (Chen et al. (2026a)) addresses this limitation by assigning velocity $\boldsymbol{u}$-prediction to a selected low-rank subspace while clean image $\boldsymbol{x}_0$-prediction in its complement. More concretely, let $\boldsymbol{U} = [\boldsymbol{v}_1, \ldots, \boldsymbol{v}_D]$ be the orthonormal patch-PCA eigenvectors of the data (variances $\lambda_1 \geq \cdots \geq \lambda_D$), and let $\boldsymbol{A} = \boldsymbol{U}_{:,1:r}$ be the top-r directions with $\boldsymbol{A}^{\mathrm{T}}\boldsymbol{A} = \boldsymbol{I}_r$, then the orthogonal projector $\boldsymbol{P} = \boldsymbol{A}\boldsymbol{A}^{\mathrm{T}}$ has the properties $\boldsymbol{P}^2 = \boldsymbol{P}$ and $(\boldsymbol{I} - \boldsymbol{P})\boldsymbol{P} = \mathbf{0}$. Their model predicts the asymmetric

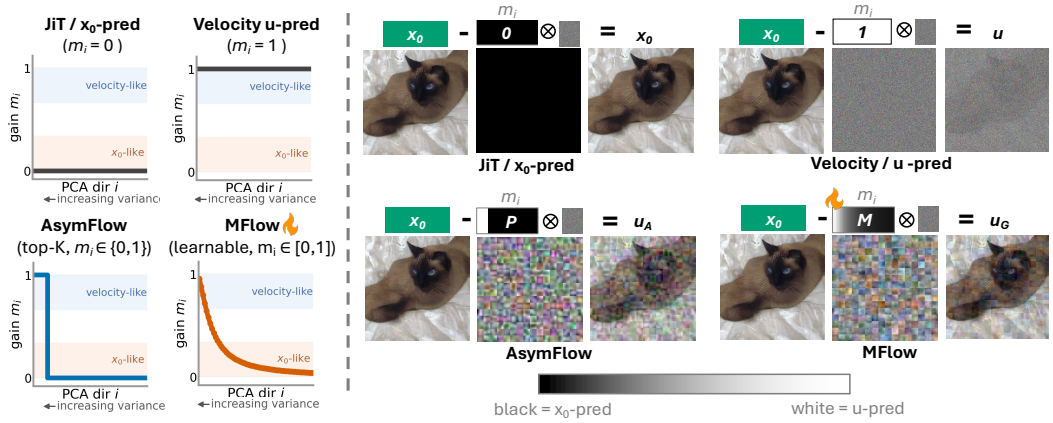


Figure 3: **Overview of MFLOW Formulation.** We propose a principal data direction-wise learnable parameterization as a generalized framework in contrast with previous fixed parameterizations.

target $\boldsymbol{u}_{\mathrm{A}} = \boldsymbol{x}_0 - \boldsymbol{P}\boldsymbol{\epsilon}$, *i.e.*, predict the full clean signal, but only the low-rank projection of the noise. Because $\boldsymbol{P}$ is a projector, the target splits cleanly into the velocity inside the subspace $\boldsymbol{P}\,\boldsymbol{u}_{\mathrm{A}} = \boldsymbol{P}(\boldsymbol{x}_0 - \boldsymbol{\epsilon}) = \boldsymbol{P}\,\boldsymbol{u}$, and the clean data on the complement $(\boldsymbol{I} - \boldsymbol{P})\,\boldsymbol{u}_{\mathrm{A}} = (\boldsymbol{I} - \boldsymbol{P})\boldsymbol{x}_0$. So each PCA direction is treated either as velocity ($i \le r$) or as $\boldsymbol{x}_0$ ($i > r$). The full velocity is recovered from $\boldsymbol{u}_{\mathrm{A}}$ and $\boldsymbol{x}_t$ by: $\boldsymbol{u} = \boldsymbol{P}\,\boldsymbol{u}_{\mathrm{A}} + (\boldsymbol{I} - \boldsymbol{P})\,\frac{\boldsymbol{x}_t + \boldsymbol{u}_{\mathrm{A}}}{1-t}$.

Limitations of previous adaptive parameterizations. Chen et al. (2026a) note that the default $\boldsymbol{x}_0$-prediction and k-Diff parameterization are less expressive. On the other hand, AsymFlow makes hard-cutoff decision at a fixed rank r, and this chosen subspace is dataset/modality dependent, requiring an expensive hyperparameter search over multiple parameterization setting. Furthermore, restricting the noise term to the top-r (highest-variance) directions is arbitrary – *e.g.*, we found that $\mathrm{PCA}[9:16]$ can beat $\mathrm{PCA}[1:8]$ (FID 3.54 vs 3.59 on the B-level model trained on ImageNet). Finally, the binary nature of the cutoff, while useful for interpretation, is not necessary: it extends naturally to a continuous, per-direction weighting that softly mixes the clean-signal ($\boldsymbol{x}_0$) and velocity ($\boldsymbol{v}$) prediction. Building on these insights, we introduce MFLOW, which learns a continuous, direction-dependent parameterization, allowing the prediction target to adapt across directions without requiring a hard subspace partition or costly search for parameterization.

3.2 MFLOW: DIRECTION-ADAPTIVE LEARNABLE REPARAMETERIZATION

Instead of a hard cutoff, we assign each spectral direction a soft, learnable weight in $[0, 1]$. Replace the binary projector $\boldsymbol{P}$ by a symmetric graded operator:

$$\boldsymbol{M} = \boldsymbol{U}\,\mathrm{diag}(m_1, \ldots, m_D)\,\boldsymbol{U}^{\mathrm{T}}, \qquad m_i \in [0,1], \qquad \boldsymbol{M}\boldsymbol{v}_i = m_i \boldsymbol{v}_i, \tag{1}$$

with the graded target (as shown in Fig. 3):

$$\boldsymbol{u}_{\mathrm{G}} = \boldsymbol{x}_0 - \boldsymbol{M}\boldsymbol{\epsilon}\,. \tag{2}$$

$\boldsymbol{M}$ is generally not a projector ($\boldsymbol{M}^2 \neq \boldsymbol{M}$ unless every $m_i \in \{0,1\}$), so it does not split $\mathbb{R}^D$ into the low-rank subspace $\mathrm{Im}(\boldsymbol{P})$ and its orthogonal complement $\mathrm{Im}(\boldsymbol{I} - \boldsymbol{P})$. Instead every direction carries its own continuous velocity-data mixing weight.

Component-wise (PCA basis). Let $\tilde{\boldsymbol{x}} := \boldsymbol{U}^{\mathrm{T}}\boldsymbol{x}$ denote the transformation of $\boldsymbol{x}$ in the PCA basis, then we have the i^{th} component $\tilde{x}_i = \boldsymbol{v}_i^{\mathrm{T}}\boldsymbol{x}$. Because $\boldsymbol{M}$ is diagonal in this basis, $(\widetilde{\boldsymbol{M}\boldsymbol{x}})_i = m_i \tilde{x}_i$, and the graded target decouples per direction:

$$\tilde{u}_{\mathrm{G},i} = \tilde{x}_{0,i} - m_i\,\tilde{\epsilon}_i \tag{3}$$

where the directional velocity is $\tilde{u}_i = \tilde{x}_{0,i} - \tilde{\epsilon}_i$. Here m_i denotes how velocity-like the target is in this direction: $m_i = 1 \Rightarrow \boldsymbol{u}$-pred while $m_i = 0 \Rightarrow \boldsymbol{x}_0$-pred. MFLOW therefore strictly generalizes both the endpoints as well as AsymFlow ($m_i = \mathbf{1}[i \le r]$).

Choosing the gains. We initialize the gains with $m_i = \frac{\hat{\lambda}_i}{\hat{\lambda}_i + c}$ where $\hat{\lambda}_i = \lambda_i / \lambda_{\max}$ is the normalized variance, and c is a small constant. By default these gains are learnable $m_i = \mathrm{sigmoid}(\phi_i)$ with a learnable per-direction logit ϕ_i.

Exact velocity recovery. The forward $\boldsymbol{x}_t = t\,\boldsymbol{x}_0 + (1-t)\,\boldsymbol{\epsilon}$ and target Eq. (3) give two linear equations in the two unknowns$(\tilde{x}_{0,i}, \tilde{\epsilon}_i)$ in each eigendirection:

$$\begin{bmatrix} \tilde{x}_{t,i} \\ \tilde{u}_{\mathrm{G},i} \end{bmatrix} = \begin{bmatrix} t & 1-t \\ 1 & -m_i \end{bmatrix} \begin{bmatrix} \tilde{x}_{0,i} \\ \tilde{\epsilon}_i \end{bmatrix}, \tag{4}$$

The system matrix has determinant $-\Delta_i$ with $\Delta_i = (1-t) + m_i t \geq 1 - t > 0$; since it is nonzero (strictly negative) for every $m_i \in [0,1]$, the solve is always exact:

$$\tilde{x}_{0,i} = \frac{(1-t)\,\tilde{u}_{\mathrm{G},i} + m_i\,\tilde{x}_{t,i}}{\Delta_i}, \qquad \tilde{\epsilon}_i = \frac{\tilde{x}_{t,i} - t\,\tilde{u}_{\mathrm{G},i}}{\Delta_i}, \tag{5}$$

Equivalently, in the operator form the vx_0 and vu can be recovered from the target $\boldsymbol{u}_{\mathrm{G}}$ using:

$$\begin{aligned} \boldsymbol{x}_0 &= \big[t\,\boldsymbol{M} + (1-t)\,\boldsymbol{I}\big]^{-1}\big[(1-t)\,\boldsymbol{u}_{\mathrm{G}} + \boldsymbol{M}\,\boldsymbol{x}_t\big], \\ \boldsymbol{u} &= \frac{\boldsymbol{x}_0 - \boldsymbol{x}_t}{1-t} = \big[t\,\boldsymbol{M} + (1-t)\,\boldsymbol{I}\big]^{-1}\big[\boldsymbol{u}_{\mathrm{G}} - (\boldsymbol{I} - \boldsymbol{M})\,\boldsymbol{x}_t\big]. \end{aligned} \tag{6}$$

Setting $m_i \in \{0,1\}$ reduces Eq (6) exactly to the AsymFlow recovery rule.

Training and sampling. $\boldsymbol{M}$ changes neither the forward process nor the transport target $\boldsymbol{u}$, only the coordinate-wise conditioning of the network's regression target. The network outputs $\hat{\boldsymbol{u}}_G$ as the learned approximation for target $\boldsymbol{u}_{\mathrm{G}}$ which we map it to the full velocity $\hat{\boldsymbol{u}}$ via Eq (6) and train a velocity-space flow-matching loss,

$$\mathcal{L}(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{x}_0, \boldsymbol{\epsilon}, t}\big[\, w(t)\, \|\hat{\boldsymbol{u}}(\hat{\boldsymbol{u}}_G; \boldsymbol{x}_t, t) - (\boldsymbol{x}_0 - \boldsymbol{\epsilon})\|_2^2 \,\big]. \tag{7}$$

Additionally, we can add an LPIPS loss on the predicted clean signal when required. At sampling time we run standard Euler/Heun on $\dot{\boldsymbol{x}} = \boldsymbol{u}$ with $\boldsymbol{u}$ from Eq. (6).

3.3 Toy Example

We revisit the toy manifold experiment of JiT (Li & He (2026)) to test two properties of MFlow: *(a)* whether continuous, non-binary *m*-values remain tractable in high-dimensional observation spaces, and *(b)* whether they can be learned jointly with the generator. We sample a 2D spiral and embed it into a D-dimensional space using a fixed random column-orthogonal projection. We fix the intrinsic dimension to $d = 2$, vary $D \in \{2, 8, 16, 512\}$, and train a five-layer ReLU MLP with 256 hidden units. We compare $\boldsymbol{x}_0$-prediction, full $\boldsymbol{u}$-prediction, AsymFlow with $r = d$, MFlow with fixed normalized-variance gains, and MFlow with gains learned from the same initialization. All variants use the same architecture and velocity-space training objective. As shown in Fig. 4, full $\boldsymbol{u}$-prediction deteriorates as D increases, whereas $\boldsymbol{x}_0$-prediction and rank-matched AsymFlow remain stable. Both fixed and learnable MFlow also recover the spiral, including at $D = 512$. Thus, high-dimensional tractability does not require binary gains $m_i \in \{0,1\}$: continuous gains remain effective. The learnable variant further shows that this allocation can be optimized end-to-end rather than specified in a heuristic manner.

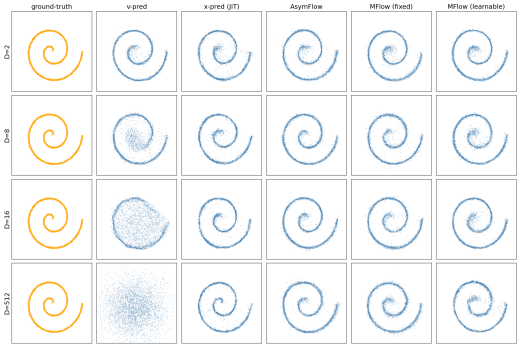


Figure 4: **Toy manifold experiment.** A 2-D spiral is embedded in $D \in \{2, 8, 16, 512\}$ dimensions and projected back for visualization. Full $\boldsymbol{u}$-prediction degrades with D, while $\boldsymbol{x}_0$-prediction, rank-matched AsymFlow, and fixed or learned MFlow recover the spiral, including in the undercomplete $D = 512$ setting.

4 SIGNAL MODELING

4.1 MOTIVATION

JiT lacks high-frequency but some other pixel diffusion models do not. JiT's minimalist architecture lacks an explicit pathway dedicated to fine-grained signal details, like skip connections (Bao et al. (2023); Hoogeboom et al. (2023; 2025)) or parallel refinement networks (Chen et al. (2026b); Yu et al. (2026)), motivating an investigation into whether this simplicity introduces limitations in its generated frequency distribution. We study the frequency spectrum of the images generated by models we trained which share the JiT architecture. We use radially averaged power spectrum (RAPS) for this analysis. For context, images on a pixel grid can have at most 0.5 cycles/pixel horizontal or vertical frequencies (f_x,f_y). However, in the case of radial averaging, the radial frequencies f_r can reach till the diagonal frequency of 0.707 ($= \sqrt{0.5^2 + 0.5^2}$). Interestingly, as seen in Fig. 5, we observe that all these models lack radial frequencies in the range of $0.5 - 0.7$. Recent concurrent works (Pavasovic et al. (2026); Degeorge et al. (2026); Fu et al. (2026)) have also observed this. As seen in Fig. 6, we tried investigating other LDMs and pixel diffusion backbones. We find that neither VAE-reconstruction/generations of DiT (Peebles & Xie (2023)), REPA-E (Leng et al. (2025)), nor the generations of PixelDiT (Yu et al. (2026)) (which as a pixel level pathway) suffer from this issue, indicating that it might be a JiT backbone specific problem, more explicitely the use of high patchidication (and unpatchification) of $16 - 32\times$ without any intra-patch local operation.

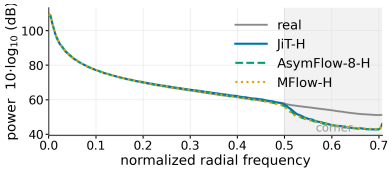


Figure 5: JiT-family generations lacks high frequency irrespective of parameterization.

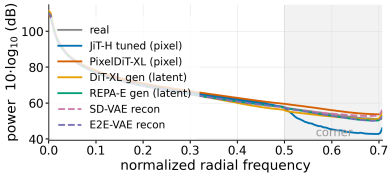


Figure 6: Latent model generations, VAE reconstructions, and other pixel diffusion models like PixelDiT do not lack high frequency.

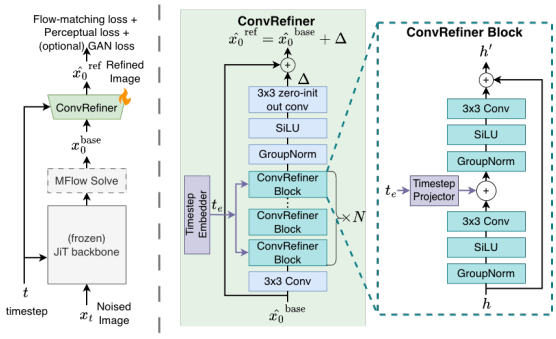


Figure 7: **Overview of *ConvRefiner* Formulation.** We propose a residual convolutional refiner post hoc training for a few epochs, with a fixed backbone, to regain the high-frequency similar to real images spectrum in vanilla vision transformers.

4.2 CONVREFINER: A SIMPLE POST-HOC FIX

Based on our analysis we suspect that this high-frequency deficit might be specific to the JiT-based pixel diffusion architecture. However, as seen in the previous section, convolutional VAE decoders are able to recover high-frequency. Chen et al. (2025a) note that only finetuning the last few layers of a pretrained VAE with an adversarial loss can improve local details at a reduced training cost.

Motivated by this, we add a few lightweight convolutional blocks and post-hoc train only these blocks using an additional GAN loss apart from flow-matchting losses (Fig 7). We keep the JiT backbone frozen to keep the refinement light. The *ConvRefiner* predicts a residual over $\boldsymbol{x}$-prediction of the network on which the perceptual and adversarial losses are then applied. For other parameterizations, the prediction is first converted to $\boldsymbol{x}$ (clean image) and then the *ConvRefiner* residual is applied on top of it.

Table 1: **Summary of Results.** Our MFLOW + *ConvRefiner* provides a flexible parameterization, saving expensive hyperparameter sweeps while having better FID, as well as solving the high-frequency deficit (measured in mean radial power deficit compared to real images) on ImageNet-1K 256×256.

Method	FID↓	HF-deficit(dB)
JiT-H/16	1.82	-6.69
MFLOW-H/16 MFLOW	**1.81**	-6.63
MFLOW-H/16 + *ConvRefiner* (ours)	**1.81**	**-1.29**

5 EXPERIMENTS

We extensively evaluate our proposed method on class-conditional ImageNet image generation to validate its effectiveness, as summarized in Table 1. Then we delve into individual components – MFLOW and *ConvRefiner*. We further show that MFLOW can be extended to video generation and scaled to text-to-image generation. Finally, we ablate the design decision of each component.

Datasets and Evaluation. For class-conditioned image generation experiments we train on ImageNet-1K (Russakovsky et al. (2015)). By default we show results on the standard 256×256 resolution images with the training and evaluation similar to Li & He (2026). For evaluation we use FID@50K (Heusel et al. (2017)). We conduct class-conditioned video generation experiments on UCF-101 (Soomro et al. (2012)) dataset. We use 16 frames of 128×128 resolution videos for training and evaluation following Wang et al. (2025). Similar to images, FVD@10K (Unterthiner et al. (2018)) serves as our video generation evaluation metric. For text-to-image generation, we train on internal dataset of around 900M images (pretraining) and 33.5M (SFT), and evaluate on standard benchmarks of GenEval (Ghosh et al. (2023)), DPG-Bench (Hu et al. (2024)), and PRISM-Bench (Fang et al. (2026)).

Implementation Details. We build on top of the JiT repository and use their architecture for conditional image and video generation experiments. Hence, we use JiT both to refer to the vanilla transformer backbone and $\boldsymbol{x}_0$-prediction parameterization (especially when comparing against MFLOW parameterization). We refer the base, large, and huge sizes as B-/L-/H-level. For these experiments, for ImageNet we used patch size of 16×16 leading to 256 tokens (with additional 32 in-context tokens like JiT), while for video we use $4 \times 8 \times 8$ patch size (with 4 as the temporal patch size) leading to 1024 tokens. For text-to-image, we train with a patch size of 16×16 on mixed-resolution 256p and 512p images with LPIPS loss (0.1 loss weighting). For calculating patch-wise PCA-basis we follow AsymFlow. In the case of *ConvRefiner*, we use a PatchGAN (Isola et al. (2017)) to apply the GAN loss only on the near-clean images (see appendix for details).

Table 2: **Effect of LPIPS on ImageNet 256×256 generation.** We see that LPIPS helps all the models in FID (evaluated on base JiT-B evaluation hyperparameters). However, it improves MFLOW performance the most, beating the baselines.

Method	FID↓	
	w/o LPIPS	w/ LPIPS
JiT-B/16 (Li & He (2026))	3.70	3.42
AsymFlow-B/16 (Chen et al. (2026a))	3.59	3.40
MFLOW-B/16 (ours)	3.74	**3.33**

ImageNet. *(a)* MFLOW*:* On class-conditioned image generation on ImageNet 256×256 experiment at B-level, as shown in Table 2, without LPIPS loss and at the base JiT-B parameters, we observe that AsymFlow works better than both JiT and MFLOW. Based on PixelGen (Ma et al. (2026c)) we tried to add LPIPS loss to our baselines and MFLOW model, and found that it provides performance gain across the board. Surprisingly, we see that this addition makes MFLOW the best at this configuration. Subsequently, we train the same models with LPIPS at H-scale as shown in Table 3. We hypothesize the very slight lead of AsymFlow as a result of it being heavily tuned for ImageNet. Even so, we find that MFLOW is very competitive to all the prior pixel diffusion transformer models without any parameterization search. *(b) ConvRefiner:* We experiment with training the JiT and MFLOW models on ImageNet 256×256 using our *ConvRefiner* for an additional 30 epochs on top of the frozen 600 epoch trained base model. As shown in Fig. 8 (left), at both B- and H-levels and across both JiT and MFLOW we see that the high radial frequency deficit is fixed. At the same time, Fig. 8 (right) shows that this training maintains the FID. Hence, *ConvRefiner* is effective at solving the high-frequency problem without affecting the FID at all. We provide some visual improvement results in Fig. 22.

Videos. We show the effectiveness of MFLOW in the domain of videos apart from images. At the B-level class-conditioned UCF-101 video generation as shown in Table 4, we see that MFLOW outperforms the WAN2.1 (Team (2025)) latent diffusion model, pixel-based JiT, and AsymFlow models we trained. We expand to L-level, where MFLOW beats JiT and AsymFlow. With a small post-hoc Framewise-LPIPS training, we achieve best performance. Notably, we find training with Framewise-LPIPS from start (like ImageNet) hurts video generation.

Text-to-Image. *(a)* MFLOW*:* We study whether MFLOW scales beyond class-conditioned generation to large-scale text-to-image (T2I). We instantiate JiT, AsymFlow ($r = 8$), and MFLOW using the same 1.1B-parameter MMDiT-L/16 backbone and compare them under the same SFT setting. This controlled comparison isolates the effect of prediction parameterization while keeping the backbone and reported model size fixed. As shown in Table 6, the three parameterizations perform similarly

Table 3: **ImageNet 256×256 comparisons**. We trained JiT, AsymFlow, and MFLOW from scratch with additional LPIPS (Zhang et al. (2018)) loss (but without REPA (Yu et al. (2025))) and evaluated under the JiT protocol with individual hyperparameter tuning (denoted by *). We use the best hyperparameter for each after doing a sweep. The AsymFlow number remains in the range of the non-REPA performance noted in Chen et al. (2026a). We use AsymFlow's FLOP convention as reported in Yu et al. (2026).

Method	Pred (±)	Params	GFLOPs	FID↓
Hierarchical CNNs (skip connections / U-Net-like (Ronneberger et al. (2015))				
ADM-G (Dhariwal & Nichol (2021))	ϵ	554M	2240	4.59
Hierarchical transformers (skip connections / U-ViT-like (Bao et al. (2023)))				
RIN (Jabri et al. (2023))	ϵ	320M	668	3.42
SiD, U-ViT/2 (Hoogeboom et al. (2023))	ϵ	2B	1110	2.44
VDM++, U-ViT/2 (Kingma & Gao (2023))	ϵ	2B	1110	2.12
SiD2, U-ViT/1 (Hoogeboom et al. (2025))	ϵ	-	1306	**1.38**
Hierarchical transformers (decoder head / DDT-like (Wang et al. (2026b)))				
DiP-XL/16 (Chen et al. (2026b))	$x_0 - \epsilon$	631M	-	1.79
DeCo-XL/16 (Ma et al. (2026b))	$x_0 - \epsilon$	682M	245	1.62
PixelDiT-XL/16 (Yu et al. (2026))	$x_0 - \epsilon$	797M	311	**1.61**
Vanilla transformers (DiT-like (Peebles & Xie (2023)))				
PixelFlow-XL/4 (Chen et al. (2025b))	$x_0 - \epsilon$	677M	5818	1.98
PixelGen-XL/16 (Ma et al. (2026c))	x_0	676M	260	1.83
PixelREPA-H/16 (Shin et al. (2026))	x_0	953M	363	1.81
JiT-H/16 (Li & He (2026))	x_0	953M	363	1.82*
AsymFlow-H/16 (Chen et al. (2026a))	$x_0 - P\epsilon$	953M	363	**1.79***
MFLOW-H/16 (Ours)	$x_0 - M\epsilon$	953M	363	1.81*

Table 4: **UCF-101 video comparisons (B-level).** We show results on models we trained with best-tuned hyperparameters for each.

Method	FVD↓
Latent diffusion models	
WAN-VAE2.1-1.3B(4×8×8) + latent-DiT-B/1×1×1	169
Pixel diffusion models	
Video JiT-B/4×8×8 (Li & He (2026))	152
Video AsymFlow-B/4×8×8 (Li & He (2026))	126
Video MFLOW-B/4×8×8 (ours)	**122**

Table 5: **UCF-101 video comparisons (L-level).** The rows in gray were trained by us and evaluated with individual hyperparameter tuning. †Note that we use the mapping network as generator and synthesis network as tokenizer convention from Ren et al. (2026).

Method	Params		#Tokens	FVD↓
	Tokenizer	Generator		
Autoregressive Models				
CogVideo (Hong et al. (2022))	–	9.4B	2065	625
TATS (Ge et al. (2022))	32M	321M	1024	332
MAGVIT-AR (Yu et al. (2023))	158M	306M	1024	265
OmniTokenizer (Wang et al. (2024))	82.2M	650M	1280	191
MAGVITv2-AR (Yu et al. (2024))	–	840M	1280	109
LARP (Wang et al. (2025))	173M	343M	1024	102
Implicit weight-space generation				
DIGAN (Yu et al. (2022))	58M†	5.5M†	–	465
NeRV-Diffusion-L (Ren et al. (2026))	∼229M	467M	128	97
Pixel diffusion models				
VideoFusion (Luo et al. (2023))	–	2B	–	173
Video JiT-L/4×8×8 (Li & He (2026))	–	458M	1024	111
Video AsymFlow-L/4×8×8 (Chen et al. (2026a))	–	458M	1024	98
Video MFLOW-L/4×8×8 (ours)	–	458M	1024	97
+ post-hoc Framewise-LPIPS (ours)	–	458M	1024	**94**

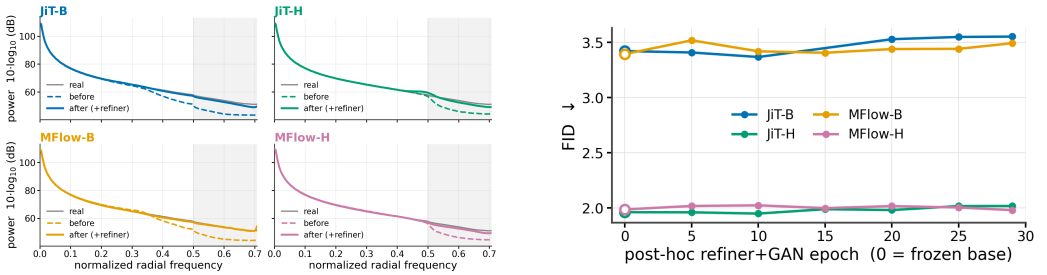


Figure 8: (Left) *ConvRefiner* recovers high frequency across JiT/MFLOW and model sizes B/H. (Right) The generative performance remains almost the same in this post-training.

at this scale, but MFLOW is best on every benchmark, even when compared against previous pixel-space T2I models. Figure 10 presents the benchmark progression over training iterations. We further scale MFLOW to an MMDiT-XXL/16 mixture-of-experts (MoE) model with 13B total parameters and 2.9B active parameters that achieves the best performance among all models in the comparison. Interestingly, model scaling also changes the parameterization learned by MFLOW as seen in Sec. 6, and MFLOW can leverage this capacity automatically without any manual tuning. *(b) ConvRefiner:* We further apply *ConvRefiner* as a post-training stage for text-to-image generation with MFLOW-XXL. Starting from the SFT checkpoint, we train the *ConvRefiner* for an additional 5k steps on top of the frozen base model. As shown in Fig. 9 (right), without *ConvRefiner* we observe the same high-frequency deficit in the 0.5–0.7 band, even though it seems to be somewhat alleviated with increased model capacity compared to Fig. 8 (left). *ConvRefiner* completely closes in this gap. This also shows up visually: Fig. 9 (left) shows sharper fine-scale texture, such as the veins of the leaf. Hence, *ConvRefiner* also fixes the high-frequency problem at scale for text-to-image generation. Additionally, as seen in 6, *ConvRefiner* does not have noticable effect on text-to-image benchmark performance.

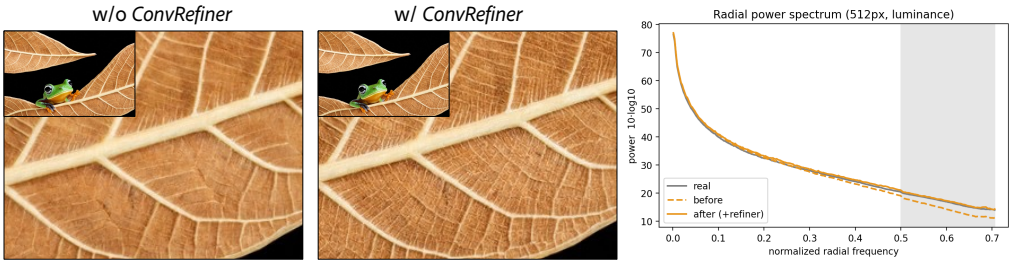


Figure 9: **Visualizing effect of *ConvRefiner* for text-to-image generation.** Here we visualize the effect of adding our post-training to MFLOW-XXL on text-to-image generations. **(left)** We see improved details in the leaf veins in the generated image. **(right)** We also note the average power gain in the high-frequency (0.5–0.7) region. We see overall improvement in texture details comparable to real image spectrum.

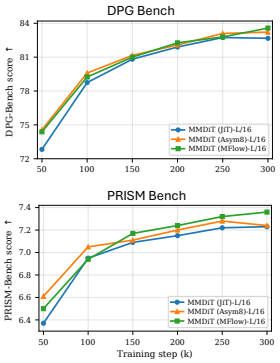


Figure 10: Text-to-Image scores over training steps.

Table 6: Comparison on text-to-image generation benchmarks. Parameter counts include only the image-generation model and exclude the text encoder. Higher is better for all metrics. Best results are in **bold**, and second-best results are underlined.

Model	Params	GenEval ↑	DPG ↑	PRISM ↑
PixelDiT-T2I (Yu et al. (2026))	1.3B	0.7359	84.58	69.6
PixNerd-XXL/16 (Wang et al. (2026a))	1.2B	0.7015	81.61	65.7
PixelFlow-XL/4 (Chen et al. (2025b))	882M	0.6463	77.90	67.6
PixelGen-XXL/16 (Ma et al. (2026c))	1.1B	0.7835	79.62	64.7
DeCo-XXL/16 (Ma et al. (2026b))	1.1B	0.8600	81.96	66.5
MiniT2I-L/16 (Wang et al. (2026c))	912M	0.8822	84.66	72.9
MMDiT (JiT)-L/16-SFT*	1.1B	0.8839	84.65	73.6
MMDiT (Asym8)-L/16-SFT*	1.1B	0.8832	84.61	73.9
MMDiT (MFLOW)-L/16-SFT* (Ours)	1.1B	0.8843	84.70	73.9
MMDiT (MFLOW)-XXL/16-SFT* (Ours)	13B-MoE (2.9B active)	**0.8898**	86.10	75.2
+ *ConvRefiner* (Ours)	13B-MoE (2.9B active)	0.8896	**86.16**	**75.4**

MFLOW Ablations. We ablated multiple axes of choices for MFLOW on ImageNet B-level model. We clamp the initializations to $[0.05, 0.95]$, as otherwise the sigmoid function will blow up $\boldsymbol{M}$. In terms of loss parameterization we see that using $\boldsymbol{u}_{\mathrm{G}}$ loss over $\boldsymbol{u}$ loss significantly deteriorates performance, bringing FID from 3.66 to 4.88. Training on $\boldsymbol{u}_{\mathrm{G}}$ creates a trivial incentive to reduce m_i toward $\boldsymbol{x}_0$-prediction without improving the recovered transport field (Appendix A.3). Subsequently, by varying the $\boldsymbol{M}$ initialization to the normalized-variance ('norm_var') as described in 3.2, we see slight improvement over 'asym8' initialization, FID reduces from 3.84 to 3.80. Given that 'norm_var' has a smoother dependence on the PCA direction's variance, we choose it as default. We also experiment with learning-rate of $\boldsymbol{M}$ parameters, and see that having a higher learning rate multiplier seems to help a bit bringing FID down from 3.80 to 3.74. We use this final setting (*init 'norm_var', $\boldsymbol{u}$-loss, learnable $\boldsymbol{M}$ w/ 10× lr-multiplier*) as our default unless specified otherwise.

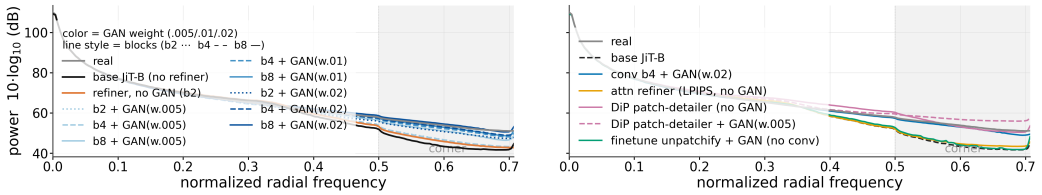


Figure 11: (Left) Increasing blocks of Conv-refiner and the weight of GAN loss helps recover high frequency. (Right) Other baselines refining does not help except DiP patch detailer(Chen et al. (2026b)). This probably reflects a convolutions bias.

***ConvRefiner* Ablations.** Our preliminary control experiment showed that JiT trained normally for 5% more epochs does not improve the frequency spectrum. In Fig. 11 (left) we investigate the effect

of number of blocks ($b \in 2, 4, 8$) and the weight of the adversarial loss ($w \in 0.005, 0.01, 0.02$). Firstly, we see that only using the refiner without the adversarial loss helps a bit. However, the frequency spectrum gets closer to real with more blocks and more GAN weight. Favoring a smaller lightweight network, we finally select $b = 4$ and $w = 0.02$ as our final setting. This improves pixel level details as seen in Fig. 23. Ablating the choice of refiner family (Fig. 11 (right)), we see that just finetuning the unpatchify layer or using a residual attention refiner has negligible effect. Another readily available convolutional refiner choice DiP patch-detailer(Chen et al. (2026b)) does improve the high-frequency and even overshoots with the GAN loss. However, it degrades the base JiT's FID 3.42 to 3.90 and 4.46 respectively, bolstering our choice of *ConvRefiner*.

6 Discussion

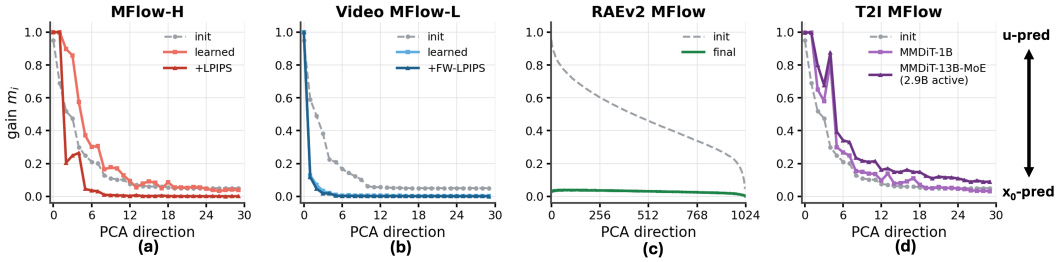


Figure 12: **Learned M values.** Variations in the MFlow gain curves across multiple axes: objectives – w/ vs w/o LPIPS, domains – images vs videos, representations – pixel vs latent $\boldsymbol{x}_0$-prediction models, model capacities – 1B vs 13B-MoE.

M-curve Analysis. Given the learned nature of the gain parameters m_i, we analyze how they vary across objectives, domains, representations, and model capacities. In Fig. 12, we visualize the learned m_i values ordered by decreasing PCA variance, which we call an M-curve. Here, $m_i = 0$ corresponds to $\boldsymbol{x}_0$-prediction and $m_i = 1$ to $\boldsymbol{u}$-prediction; the 'norm_var' initialization is shown in gray.

1. On ImageNet (a), adding LPIPS drives all but the leading few directions toward $\boldsymbol{x}_0$-prediction while also improving FID (Table 2), *suggesting less need for $\boldsymbol{u}$-like prediction beyond the dominant PCA directions*.
2. Moving to videos (b), the number of directions with non-negligible gains shrinks further, *showing that the preferred parameterization varies across domains*. We discuss in detail about the behaviour of LPIPS on M-curve in A.4.
3. With RAEv2 (Singh et al. (2026)) (c), whose latent model uses $\boldsymbol{x}_0$-prediction, the gains collapse almost entirely toward zero, recovering an approximately pure $\boldsymbol{x}_0$ parameterization. Yet forcing $\boldsymbol{u}$-prediction on a small leading subspace changes FID only marginally (Table 10), *indicating that a low-rank $\boldsymbol{u}$ subspace can still be tolerated*.
4. Finally, in text-to-image generation (d), increasing model capacity shifts more directions toward $\boldsymbol{u}$-prediction. Together with the RAEv2 result, this suggests that *model capacity influences how much $\boldsymbol{u}$-prediction a model can accommodate, whereas the statistics of the data representation determine where that $\boldsymbol{u}$-prediction is most useful*.

In particular, **for a fixed budget of N $\boldsymbol{u}$-predicted directions, the allocation should remain data-dependent, with the leading high-variance PCA directions providing a natural choice**. MFlow can exploit increased capacity or adapt to domain change, without requiring a manually chosen rank.

7 Conclusion

We revisit two design questions in vanilla pixel diffusion transformers: how their prediction parameterization should adapt to the data, and whether their minimalist architecture captures high-frequency. We introduce MFlow, a direction-adaptive parameterization that learns a continuous

interpolation between clean-data and velocity prediction across principal data directions, avoiding the need for manual subspace selection, saving parameterization search cycles. Our analysis of the learned M-curves further shows that the preferred parameterization changes systematically with both data statistics and model capacity. Beyond parameterization, we identify a persistent high-frequency deficit in JiT-style models and address it with *ConvRefiner*, a lightweight post-hoc convolutional module that restores high-frequency statistics while preserving generative quality. Together, these results suggest that both adaptive prediction targets and explicit treatment of fine-scale signal structure are important ingredients for scaling simple pixel-space diffusion transformers.

AI USE STATEMENT

In this work, we used AI tools only in supporting non-research roles under author supervision. Specifically, we used it to correct grammatical errors and improve the clarity of author-written text, and to assist with software tooling. In particular, we used it for plotting and visualization code used to produce some of the figures in this paper. We did not use AI tools for research ideations, for deriving results, for designing experiments, or for generating any scientific claims. We have reviewed all the AI-assisted text and figures for correctness.

REPRODUCIBILITY STATEMENT

We provide the MFLOW's closed-form velocity-recovery rule in Sec. 3.2 and Python-like pseudocode in Sec. A.2. The post-hoc *ConvRefiner* is specified as an architecture figure in Fig. 7, and its training objective is provided in Section 4 and Appendix A.7. Full training and sampling configurations like optimizer, learning-rate schedule, batch size, EMA, noise schedule, and the guidance interval, guidance scale, and number of sampling steps used at evaluation, are tabulated per model in Appendix A.1. Radially-averaged power spectrum (RAPS) evaluation protocol are described in Appendix A.6. All reported metrics (FID, FVD, RAPS, GenEval, DPG, Prism) are computed with standard public implementations we cite.

ACKNOWLEDGMENTS

We would like to thank Harsha Walvekar, Vatsal Agarwal, Prateksha Udhayanan, Namitha Padmanabhan, and Matthew Gwilliam for their help shaping this paper.

REFERENCES

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 *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 22669–22679. IEEE, 2023.

Hansheng Chen, Jan Ackermann, Minseo Kim, Gordon Wetzstein, and Leonidas Guibas. Asymmetric flow models. *arXiv preprint arXiv:2605.12964*, 2026a. URL https://arxiv.org/abs/2605.12964.

Junyu Chen, Han Cai, Junsong Chen, Enze Xie, Shang Yang, Haotian Tang, Muyang Li, and Song Han. Deep compression autoencoder for efficient high-resolution diffusion models. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 96539–96560, 2025a. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/f062da1973ac9ac61fc6d44dd7fa309f-Paper-Conference.pdf.

Shoufa Chen, Chongjian Ge, Shilong Zhang, Peize Sun, and Ping Luo. Pixelflow: Pixel-space generative models with flow. *arXiv preprint arXiv:2504.07963*, 2025b.

Zhennan Chen, Junwei Zhu, Xu Chen, Jiangning Zhang, Xiaobin Hu, Hanzhen Zhao, Chengjie Wang, Jian Yang, and Ying Tai. Dip: Taming diffusion models in pixel space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 36136–36146, June 2026b.

Katherine Crowson, Stefan Andreas Baumann, Alex Birch, Tanishq Mathew Abraham, Daniel Z Kaplan, and Enrico Shippole. Scalable high-resolution pixel-space image synthesis with hourglass diffusion transformers. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 9550–9575. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/crowson24a.html.

Lucas Degeorge, Paul Couairon, Arijit Ghosh, Alexei A. Efros, David Picard, and Vicky Kalogeiton. Balancing frequencies and pixels in flow matching, 2026. URL https://arxiv.org/abs/2609.02748.

Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.

Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.

Weichen Fan, Haiwen Diao, Penghao Wu, and Ziwei Liu. Show the signal, hide the noise: Spectral forcing for pixel-space diffusion, 2026. URL https://arxiv.org/abs/2606.15236.

Rongyao Fang, Aldrich Yu, Chengqi Duan, Linjiang Huang, Shuai Bai, Yuxuan Cai, Kun Wang, Si Liu, Xihui Liu, and Hongsheng Li. Flux-reason-6m & prism-bench: A million-scale text-to-image reasoning dataset and comprehensive benchmark. In *International Conference on Learning Representations*, volume 2026, pp. 63157–63186, 2026.

Zixuan Fu, Chong Wang, Lanqing Guo, Kailai Zhou, Jiahao Nie, and Bihan Wen. A frozen pixel-space diffusion model can guide itself with its own samples. *arXiv preprint arXiv:2607.29122*, 2026.

Songwei Ge, Thomas Hayes, Harry Yang, Xi Yin, Guan Pang, David Jacobs, Jia-Bin Huang, and Devi Parikh. Long video generation with time-agnostic vqgan and time-sensitive transformer. In *European conference on computer vision*, pp. 102–118. Springer, 2022.

Dhruba Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. *Advances in Neural Information Processing Systems*, 36: 52132–52152, 2023.

Zipeng Guo, Lichen Ma, Yu He, Xiaolong Fu, Jingling Fu, Junshi Huang, and Yan Li. Pixelu: A u-shaped transformer for efficient end-to-end pixel diffusion. In *European Conference on Computer Vision*, pp. 425–442. Springer, 2026.

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*, 30, 2017.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *arXiv preprint arXiv:2205.15868*, 2022.

Emiel Hoogeboom, Jonathan Heek, and Tim Salimans. simple diffusion: End-to-end diffusion for high resolution images. In *International Conference on Machine Learning*, pp. 13213–13232. PMLR, 2023.

Emiel Hoogeboom, Thomas Mensink, Jonathan Heek, Kay Lamerigts, Ruiqi Gao, and Tim Salimans. Simpler diffusion: 1.5 fid on imagenet512 with pixel-space diffusion. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 18062–18071. IEEE, 2025.

Xiwei Hu, Rui Wang, Yixiao Fang, Bin Fu, Pei Cheng, and Gang Yu. Ella: Equip diffusion models with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*, 2024.

Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *2017 IEEE conference on computer vision and pattern recognition (CVPR)*, pp. 5967–5976. Ieee, 2017.

Allan Jabri, David J. Fleet, and Ting Chen. Scalable adaptive computation for iterative generation. In *Proceedings of the 40th International Conference on Machine Learning*, ICML'23. JMLR.org, 2023.

Dengyang Jiang, Ruoyi Du, Zhennan Chen, Dongyang Liu, Zanyi Wang, Mingzhe Zheng, Xiangpeng Yang, Huanqia Cai, Aiming Hao, Yuming Jiang, Peng Gao, Harry Yang, and Steven Hoi. An empirical study of training pixel-space text-to-image diffusion models, 2026. URL https://arxiv.org/abs/2608.16887.

Qing Jin and Chaoyang Wang. Revisiting diffusion model predictions through dimensionality. *arXiv preprint arXiv:2601.21419*, 2026.

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35:26565–26577, 2022.

Diederik P Kingma and Ruiqi Gao. Understanding diffusion objectives as the ELBO with simple data augmentation. In *NeurIPS*, 2023. URL https://openreview.net/forum?id=NnMEadcdyD.

Xingjian Leng, Jaskirat Singh, Yunzhong Hou, Zhenchang Xing, Saining Xie, and Liang Zheng. Repa-e: Unlocking vae for end-to-end tuning with latent diffusion transformers. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 18262–18272. IEEE, 2025.

Tianhong Li and Kaiming He. Back to basics: Let denoising generative models denoise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 36115–36125, June 2026.

Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *2021 IEEE/CVF international conference on computer vision (ICCV)*, pp. 9992–10002. Ieee, 2021.

Zhiheng Liu, Weiming Ren, Xiaoke Huang, Shoufa Chen, Tianhong Li, Mengzhao Chen, Yatai Ji, Sen He, Jonas Schult, Belinda Zeng, Tao Xiang, Wenhu Chen, Ping Luo, Luke Zettlemoyer, and Yuren Cong. Tuna-2: Pixel embeddings beat vision encoders for multimodal understanding and generation, 2026. URL https://arxiv.org/abs/2604.24763.

Zhengxiong Luo, Dayou Chen, Yingya Zhang, Yan Huang, Liang Wang, Yujun Shen, Deli Zhao, Jingren Zhou, and Tieniu Tan. Videofusion: Decomposed diffusion models for high-quality video generation. *arXiv preprint arXiv:2303.08320*, 2023.

Lichen Ma, Zipeng Guo, Yu He, Xiaolong Fu, Luohang Liu, Jingling Fu, Junshi Huang, and Yan Li. Frequencybooster: Full-frequency modeling for high-fidelity pixel diffusion, 2026a. URL https://arxiv.org/abs/2605.17759.

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*, pp. 23–40. Springer, 2024.

Zehong Ma, Longhui Wei, Shuai Wang, Shiliang Zhang, and Qi Tian. Deco: Frequency-decoupled pixel diffusion for end-to-end image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 43600–43610, June 2026b.

Zehong Ma, Ruihan Xu, and Shiliang Zhang. Pixelgen: Improving pixel diffusion with perceptual supervision, 2026c. URL https://arxiv.org/abs/2602.02493.

Soumik Mukhopadhyay, Prateksha Udhayanan, and Abhinav Shrivastava. Scale space diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 35851–35860, June 2026.

Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pp. 8162–8171. PmLR, 2021.

Krunoslav Lehman Pavasovic, Théophane Vallaeys, Stéphane Mallat, Giulio Biroli, Luke Zettlemoyer, Brian Karrer, and Jakob Verbeek. Wait for the signal: Simple frequency-aware flow-matching, 2026. URL https://arxiv.org/abs/2607.28760.

William Peebles and Saining Xie. Scalable diffusion models with transformers. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4172–4182. IEEE, 2023.

Yixuan Ren, Hanyu Wang, Bo He, Hao Chen, and Abhinav Shrivastava. Nerv-diffusion: Diffuse implicit neural representation for video synthesis. In *International Conference on Learning Representations*, volume 2026, pp. 15548–15577, 2026.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *2022 IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pp. 10674–10685. ieee, 2022.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241. Springer, 2015.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.

Jaeyo Shin, Jiwook Kim, and Hyunjung Shim. Representation alignment for just image transformers is not easier than you think, 2026. URL https://arxiv.org/abs/2603.14366.

Jaskirat Singh, Boyang Zheng, Zongze Wu, Richard Zhang, Eli Shechtman, and Saining Xie. Improved baselines with representation autoencoders. *arXiv preprint arXiv:2605.18324*, 2026.

Ivan Skorokhodov, Willi Menapace, Aliaksandr Siarohin, and Sergey Tulyakov. Hierarchical patch diffusion models for high-resolution video generation. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7569–7579. IEEE, 2024.

Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.

Wan Team. Wan: Open and advanced large-scale video generative models. 2025.

Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018.

Hanyu Wang, Saksham Suri, Yixuan Ren, Hao Chen, and Abhinav Shrivastava. Larp: Tokenizing videos with a learned autoregressive generative prior. In *International Conference on Learning Representations*, volume 2025, pp. 60671–60697, 2025.

Junke Wang, Yi Jiang, Zehuan Yuan, Binyue Peng, Zuxuan Wu, and Yu-Gang Jiang. Omnitokenizer: A joint image-video tokenizer for visual generation. *Advances in Neural Information Processing Systems*, 37:28281–28295, 2024.

Shuai Wang, Ziteng Gao, Chenhui Zhu, Weilin Huang, and Limin Wang. Pixnerd: Pixel neural field diffusion. In *International Conference on Learning Representations*, volume 2026, pp. 43559–43580, 2026a.

Shuai Wang, Zhi Tian, Weilin Huang, and Limin Wang. Ddt: Decoupled diffusion transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 40633–40642, 2026b.

Xianbang Wang, Hanhong Zhao, Yiyang Lu, Kangyang Zhou, Linrui Ma, and Kaiming He. Minit2i: A minimalist baseline for text-to-image generation, 2026c. URL https://peppaking8.github.io/#/post/minit2i.

Lijun Yu, Yong Cheng, Kihyuk Sohn, José Lezama, Han Zhang, Huiwen Chang, Alexander G Hauptmann, Ming-Hsuan Yang, Yuan Hao, Irfan Essa, et al. Magvit: Masked generative video transformer. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10459–10469. IEEE, 2023.

Lijun Yu, José Lezama, Nitesh Bharadwaj Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G Hauptmann, et al. Language model beats diffusion-tokenizer is key to visual generation. In *International Conference on Learning Representations*, volume 2024, pp. 765–783, 2024.

Sihyun Yu, Jihoon Tack, Sangwoo Mo, Hyunsu Kim, Junho Kim, Jung-Woo Ha, and Jinwoo Shin. Generating videos with dynamics-aware implicit generative adversarial networks. *arXiv preprint arXiv:2202.10571*, 2022.

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 *International Conference on Learning Representations*, 2025.

Yongsheng Yu, Wei Xiong, Weili Nie, Yichen Sheng, Shiqiu Liu, and Jiebo Luo. Pixeldit: Pixel diffusion transformers for image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14273–14282, June 2026.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE/CVF conference on computer vision and pattern recognition*, pp. 586–595. IEEE, 2018.

A Appendix

A.1 Hyperparameters

We provide our training hyperparameters in Table 7. For the ImageNet and UCF101 experiments we do not specifically change any of the base architecture, unless specified otherwise. Furthermore, we also provide evaluation time hyperparameters for these experiments in Table 8. Because the model parameterization changes over training for MFlow and hence it is not clear if the defaults work. We report inference steps, classifier-free guidance (CFG), sampling velocity-recovery clamp $\sigma_{\min}$, and CFG Interval.

We use the default JiT-B hyperparameters (Heun 50 steps, CFG 2.9, CFG Interval $[0.10, 1.0]$, $\sigma_{\text{min}} = 0.04$) for initial ImageNet B-level experiments to isolate the effect of training objective. For H-scale results we first conducted a coarse per-model CFG scan at $\sigma_{\min} = 0.04$, interval $[0.10, 1]$, followed by a finer CFG$\times\sigma_{\min}$ grid on the strongest (LPIPS) variant.

For UCF-101, we firstly observed that a smaller spatial patch size of 8×8 worked better than 16×16. For the Wan2.1 latent model we used the 'Wan-AI/Wan2.1-T2V-1.3B-Diffusers' VAE and a DiT based on JiT's modern architecture. For evaluation hyperparameters, we use a two-stage search because full-scale FVD evaluation is expensive. We first run a FVD@2K grid over CFG and $\sigma_{\min}$ with Heun/100 to find the basin. And then over a narrowed CFG$\times\sigma_{\min}$ grid around the coarse optimum we calculate the full-scale FVD@10K. On the baseline, we also found 100 steps of Heun works the best around the basin, so we use that for all the experiments. The MFlow + post-hoc Framewise-LPIPS loss model was trained for an additional 10% epochs. (We note that, the UCF101 results in Skorokhodov et al. (2024) is evaluated at a different spatial resolution of 256 with temporal subsampling and without a publicly available codebase or checkpoint it is not comparable. Hence, we omit them.)

We additionally port the three objectives to a large text-to-image (T2I) stack: a multimodal DiT (MMDiT) over a weight-free PixelVAE latent (16×16 pixel patchify $\rightarrow$ 768-d tokens) with a frozen Qwen3-VL-8B text encoder (hidden layers $\{10, 20, -1\}$). We train a dense 1.1B model and a 13B mixture-of-experts (2.9B active) on $\sim$900M images (multi-resolution $256/512$) with captions, then supervised-fine-tune for 4K steps on a high-quality mixture 33.5M aesthetic-filtered images. All variants share the same natural-image patch-PCA basis (uncentered, 16×16) as the ImageNet models, and MFlow again uses rational gains (c=0.05) with a $10\times$ gain learning rate. Unlike

the ImageNet-H and UCF settings, we do *not* run a per-model CFG$\times\sigma_{\min}$ search for T2I. We fix a single shared sampler for all variants (DDIM, 50 steps, CFG 3.0; Table 8), so the T2I numbers reflect out-of-the-box behavior.

Table 7: Configuration for our JiT / AsymFlow / MFLOW models. Architecture and optimizer follow JiT-B/L/H (Li & He (2026) Tab. 9) for image/video; the text-to-image (T2I) stack is a PixelVAE-latent MMDiT. We list the settings we update.

	Image (ImageNet-256)	**Video** (UCF-101, 16×128^2)	**Text-to-Image**
Architecture			
Backbone	JiT-B/H (JiT Tab. 9)	VideoJiT-B/L (= JiT depth/width)	MMDiT
Params	131M / 953M	131M / 458M	1.1B / 13B-MoE 2.9B act
Patch	16×16	$4\times8\times8$	16×16 (PixelVAE)
Sequence length	256 (+32 in-context)	1024	res-dep. (1024 @ 512px)
Bottleneck	128 (B), 256 (H)	256	– (none)
Text encoder	– (class label)	– (class label)	Qwen3-VL-8B (frozen)
Latent	pixel	pixel	pixel
Dropout	0 (B), 0.2 (H)	0	0
Training			
Optimizer	AdamW, $\beta{=}(0.9, 0.95)$	same	AdamW , $\beta{=}(0.95, 0.95)$
Epochs / warmup	600 / 5	3000 / 5	300K+4K steps / 1K
Effective batch	1024	512 (B), 256 (L)	2048[†]
Learning rate	2×10^{-4}	1×10^{-4}/5×10^{-5}	1×10^{-4}
LR schedule	constant	constant	warmup-stable-decay
EMA decay	{0.9996,0.9998,0.9999}	same	0.9995
Time sampler	logit-normal $\mu{=}-0.8, \sigma{=}0.8$	same	resolution-shifted[†]
Precision	bf16, no grad-clip	same	bf16, grad-clip 1.0
Class drop (CFG)	0.1	0.1	0.05 (text drop)
AsymFlow / MFLOW			
AsymFlow rank	$r{=}8$	$r{=}8$	$r{=}8$
MFLOW gains init	$m{=}\hat{\lambda}/(\hat{\lambda}{+}c)$, $c{=}0.05$	same	same
MFLOW θ	learnable, init 'norm_var'	same	learnable, init 'type' (clamp)
MFLOW loss space	velocity (u)	velocity (u)	velocity (u)
MFLOW gain LR mult.	$\times10$	$\times1$	$\times10$
PCA basis	fixed patch-PCA, ImageNet (uncentered, 16×16)	fixed tubelet-PCA, UCF-101 (uncentered, $4\times8\times8$)	fixed patch-PCA, natural img (uncentered, 16×16)
LPIPS			
Backbone / gate	VGG, $t\geq0.3$	VGG, $t\geq0.3$	VGG, $t\geq0.3$
Weight	0.1	0.2	0.1
Spatial extent	full frame	random $H/2\times W/2$	full frame
Frames	–	all 16 (framewise)	–

[†] Normalized time $\tau{=}t/T$ (with $t/T{=}\sigma(z),\ z{\sim}\mathcal{N}(0,1)$) is warped as $\tau \mapsto \dfrac{s\,\tau}{1+(s-1)\,\tau}$, where $s{=}\min\big(\sqrt{N/N_{\text{ref}}},\ s_{\max}\big)$, N is the latent token count ($HW\cdot\min(T,4)$; images $T{=}1$), $N_{\text{ref}}{=}256$ (16×16), and $s_{\max}{=}6$ (SD3/Hunyuan-style resolution shift).

Table 8: Evaluation / sampling hyperparameters for our trained models (Tables 2, 3, 4, 5). Table 1 uses a single shared base JiT-B configuration while the rest are tuned per model.

Table	Model	Solver	Steps	CFG	$\sigma_{\min}$	CFG Interval
Table 3 (ImageNet, H, tuned)	JiT-H +LPIPS	Heun	50	2.50	0.04	$[0.10, 1.0]$
	AsymFlow-H +LPIPS	Heun	50	2.20	0.04	$[0.10, 1.0]$
	MFLOW-H +LPIPS	Heun	50	2.35	0.04	$[0.10, 1.0]$
Table 4 (UCF, B, tuned)	Video JiT-B	Heun	100	3.0	10^{-5}	$[0, 1.0]$
	Video AsymFlow-B	Heun	100	3.5	0.04	$[0, 1.0]$
	Video MFLOW-B	Heun	100	3.5	0.04	$[0, 1.0]$
Table 5 (UCF, L, tuned)	Video JiT-L	Heun	100	3.0	10^{-5}	$[0, 1.0]$
	Video AsymFlow-L	Heun	100	4.0	0.04	$[0, 1.0]$
	Video MFLOW-L	Heun	100	4.0	0.04	$[0, 1.0]$
	+post-hoc FW-LPIPS	Heun	100	3.0	0.04	$[0, 1.0]$
Tables 3* (T2I, shared)	MMDiT JiT-L +SFT	DDIM	50	3.0	10^{-5}	$[0, 1.0]$
	MMDiT AsymFlow-L +SFT	DDIM	50	3.0	0.04	$[0, 1.0]$
	MFLOW-L +SFT	DDIM	50	3.0	0.04	$[0, 1.0]$
	MFLOW-XXL (MoE) +SFT	DDIM	50	3.0	0.04	$[0, 1.0]$

A.2 MFLOW IMPLEMENTATION DETAILS

All operators are diagonal in the fixed patch-PCA basis $\boldsymbol{U}$, so every step is a change of basis (`img→patch coeffs @ U`), a per-direction scalar op, and a change back (`@ U`$^\top$ $\rightarrow$ `img`). Below, `x` is a batch of patch-token tensors, `t∈[0,1]` is flow matching timestep, `U (D,D)` and `lam (D,)` are the precomputed PCA basis/variances. These mirror the JiT codebase.

```
# One-time setup: fixed gains from lambda, or learnable
# params. lam_n is the normalized variance (lam_max = 1).
def setup_gains(lam):
  lam_n = lam.clamp(min=0) / lam.clamp(min=0).max()
  m_fixed = lam_n / (lam_n + c)
  return m_fixed

# Learnable variant (default): phi (D,) logits + scalar
# m_i = sigmoid(phi_i).
def gains():
  if learnable:
    m = sigmoid(phi)      # phi is init with torch.logit(m_fixed)
  else:
    m = m_fixed
  return m
```

```
# Graded target u_G = x0 - M eps (Using Eq.2), in the PCA basis.
def mflow_target(x0, eps, t):
  m  = gains()
  ce = to_patch_coeffs(eps) @ U          # eps in PCA coords
  Me = (m * ce) @ U.T                    # M eps = U diag(m) U^T eps
  return x0 - to_img(Me)                 # u_G
```

```
# Exact recovery of x0-hat from the output u_G (Using Eq.5,
    Using Eq.6).
def mflow_to_x0(u_M, x_t, t):
  m   = gains(t)
  omt = clamp_min(1 - t, t_eps)          # (1 - t)
  cu  = to_patch_coeffs(u_M) @ U         # u_M coeffs
  cz  = to_patch_coeffs(x_t) @ U         # x_t coeffs
  denom = clamp_min(omt + m * t, t_eps)  # Delta_i = (1-t)+m_i t
  x0_c  = (omt * cu + m * cz) / denom    # x0-hat coeffs (eq. 11)
  return to_img(x0_c @ U.T)
```

```
# Velocity for sampling (Using Eq.6).
def mflow_to_velocity(u_M, x_t, t):
  x0_hat = mflow_to_x0(u_M, x_t, t)
  return (x0_hat - x_t) / clamp_min(1 - t, t_eps)
```

```
# Training step: velocity-space flow-matching loss.
def training_loss(net, x0, t):
  eps = randn_like(x0)
  x_t = t * x0 + (1 - t) * eps           # forward interpolant
  u_M = net(x_t, t, cond)                # raw output ~= u_G
  u_pred = mflow_to_velocity(u_M, x_t, t)  # -> full velocity
  loss   = weight(t) * mse(u_pred, x0 - eps)  # Using Eq.7
  return loss
```

```
# Sampling: Euler/Heun on velocity, integrating t: 0 -> 1.
```

```
def sample(net, steps):
  x = randn(...)                          # t = 0, pure noise
  for t0, t1 in schedule(steps):
    u = mflow_to_velocity(net(x, t0, cond), x, t0)
    x = heun_step(x, u, t0, t1, net)
  return x
```

A.3 WHY MFLOW USES u-PREDICTION

MFLOW predicts the graded target

$$\boldsymbol{u}_{\mathrm{G}} = \boldsymbol{x}_0 - \boldsymbol{M}\boldsymbol{\epsilon}, \tag{8}$$

where the direction-wise gains $m_i = \mathrm{sigmoid}(\phi_i)$ are optimized jointly with the network. A natural alternative would be to directly regress this target:

$$\mathcal{L}_{\mathrm{G}}(\boldsymbol{\theta}, \boldsymbol{\phi}) = \mathbb{E}\Big[w(t)\|\hat{\boldsymbol{u}}_{\mathrm{G}} - \boldsymbol{u}_{\mathrm{G}}\|_2^2\Big]. \tag{9}$$

Here, $\boldsymbol{u}_{\mathrm{G}}$ is not detached from the computation graph, so this objective does propagate gradients to the gains. We show, however, that it introduces an undesirable dependence on the scale of the chosen parameterization.

Recall that

$$\boldsymbol{x}_t = t\boldsymbol{x}_0 + (1-t)\boldsymbol{\epsilon}, \qquad \boldsymbol{u} = \boldsymbol{x}_0 - \boldsymbol{\epsilon}. \tag{10}$$

In PCA direction i, the graded target can be rewritten as

$$\tilde{u}_{\mathrm{G},i} = (1-m_i)\tilde{x}_{t,i} + \Delta_i(t)\tilde{u}_i, \qquad \Delta_i(t) \coloneqq (1-t) + tm_i. \tag{11}$$

Consequently, the velocity recovered from a network prediction $\hat{\tilde{u}}_{\mathrm{G},i}$ is

$$\hat{\tilde{u}}_i = \frac{\hat{\tilde{u}}_{\mathrm{G},i} - (1-m_i)\tilde{x}_{t,i}}{\Delta_i(t)}. \tag{12}$$

Combining Eqs. equation 11 and equation 12 gives the exact error relation

$$\hat{\tilde{u}}_{\mathrm{G},i} - \tilde{u}_{\mathrm{G},i} = \Delta_i(t)\Big(\hat{\tilde{u}}_i - \tilde{u}_i\Big). \tag{13}$$

Thus, direct graded-target supervision is equivalent to

$$\mathcal{L}_{\mathrm{G}} = \mathbb{E}\left[w(t)\sum_i \Delta_i(t)^2\Big(\hat{\tilde{u}}_i - \tilde{u}_i\Big)^2\right]. \tag{14}$$

The same recovered-velocity error therefore receives a different penalty depending on m_i and t.

The gains receive gradients. Because the target is not detached, the per-sample derivative of the direct loss is

$$\left.\frac{\partial \ell_{\mathrm{G},i}}{\partial m_i}\right|_{\hat{\tilde{u}}_{\mathrm{G},i}} = 2w(t)\Big(\hat{\tilde{u}}_{\mathrm{G},i} - \tilde{u}_{\mathrm{G},i}\Big)\tilde{\epsilon}_i. \tag{15}$$

Equivalently,

$$\frac{\partial \ell_{\mathrm{G},i}}{\partial \phi_i} = 2w(t)\Big(\hat{\tilde{u}}_{\mathrm{G},i} - \tilde{u}_{\mathrm{G},i}\Big)\tilde{\epsilon}_i\, m_i(1-m_i). \tag{16}$$

This gradient is generally nonzero, but its sign depends on the current prediction residual. The issue with direct supervision is therefore not that the gains cannot learn, but that the objective admits a parameterization-scale shortcut.

A parameterization-scale shortcut. Consider any fixed recovered-velocity predictor $\bar{\tilde{u}}_i(\boldsymbol{x}_t, t)$. For every value of m_i, define the corresponding graded prediction

$$\hat{\tilde{u}}_{\mathrm{G},i}^{(m_i)} = (1-m_i)\tilde{x}_{t,i} + \Delta_i(t)\bar{\tilde{u}}_i. \tag{17}$$

Substituting Eq. equation 17 into the recovery map always yields the same predictor:

$$\frac{\hat{\tilde{u}}_{\mathrm{G},i}^{(m_i)} - (1-m_i)\tilde{x}_{t,i}}{\Delta_i(t)} = \bar{\tilde{u}}_i. \tag{18}$$

Hence, changing m_i along this family does not change the recovered transport field. Nevertheless, the direct graded-target loss becomes

$$\ell_{\mathrm{G},i}(m_i) = w(t)\,\Delta_i(t)^2(\bar{\tilde{u}}_i - \tilde{u}_i)^2. \tag{19}$$

Its derivative along this equivalent-predictor family is

$$\frac{\mathrm{d}\ell_{\mathrm{G},i}}{\mathrm{d}m_i} = 2w(t)\,t\,\Delta_i(t)(\bar{\tilde{u}}_i - \tilde{u}_i)^2 \geq 0. \tag{20}$$

Therefore, whenever $t > 0$ and the recovered velocity has nonzero error, decreasing m_i reduces the graded-target loss while leaving the recovered velocity unchanged. The direct objective can thus improve by making the target more $\boldsymbol{x}_0$-like, rather than by learning a more accurate transport field.

This does not imply that every instantaneous minibatch gradient must decrease m_i. Instead, Eq. equation 20 shows that the joint optimization problem contains a loss-reducing direction that is unrelated to transport accuracy.

For the binary AsymFlow parameterization,

$$\Delta_i(t)^2 = \begin{cases} 1, & m_i = 1, \\ (1-t)^2, & m_i = 0. \end{cases} \tag{21}$$

Direct $\boldsymbol{u}_{\mathrm{G}}$ supervision therefore suppresses recovered-velocity errors in the clean-image-prediction subspace by $(1-t)^2$. Near the clean endpoint, $t \to 1$, these errors become almost unpenalized, even though they are amplified when the prediction is converted back to velocity.

Velocity-space supervision. MFLOW instead applies the standard flow-matching objective after exact velocity recovery:

$$\mathcal{L}_{\boldsymbol{u}}(\boldsymbol{\theta}, \boldsymbol{\phi}) = \mathbb{E}\left[w(t)\sum_i \left(\hat{\tilde{u}}_i - \tilde{u}_i\right)^2\right]. \tag{22}$$

Using Eq. equation 13, this is equivalently

$$\mathcal{L}_{\boldsymbol{u}} = \mathbb{E}\left[w(t)\sum_i \frac{\left(\hat{\tilde{u}}_{\mathrm{G},i} - \tilde{u}_{\mathrm{G},i}\right)^2}{\Delta_i(t)^2}\right]. \tag{23}$$

The inverse factor from velocity recovery cancels the artificial $\Delta_i(t)^2$ scaling of the graded target. In particular, along the equivalent-predictor family in Eq. equation 17,

$$\frac{\mathrm{d}}{\mathrm{d}m_i}\left[w(t)(\bar{\tilde{u}}_i - \tilde{u}_i)^2\right] = 0. \tag{24}$$

Velocity-space supervision therefore evaluates all parameterizations using the same transport objective. The gains can still adapt because they alter the conditioning and finite-capacity approximation problem faced by the network, but they cannot reduce the objective merely by shrinking the scale of the regression target.

A.4 SURPRISING BEHAVIOUR LPIPS ON THE LEARNED $\boldsymbol{M}$-CURVE

Figure 12(a) shows that LPIPS substantially changes the prediction parameterization learned by MFLOW. Without LPIPS, the gains decrease approximately monotonically with PCA variance: high-variance directions receive more $\boldsymbol{u}$-like prediction, while lower-variance directions become increasingly $\boldsymbol{x}_0$-like. This behavior is consistent with the normalized-variance initialization. With LPIPS, however, the leading gains become non-monotonic. Using zero-based indexing, PC0 and PC1 retain high gains, PC2–PC3 exhibit a pronounced dip, and PC4 rises again.

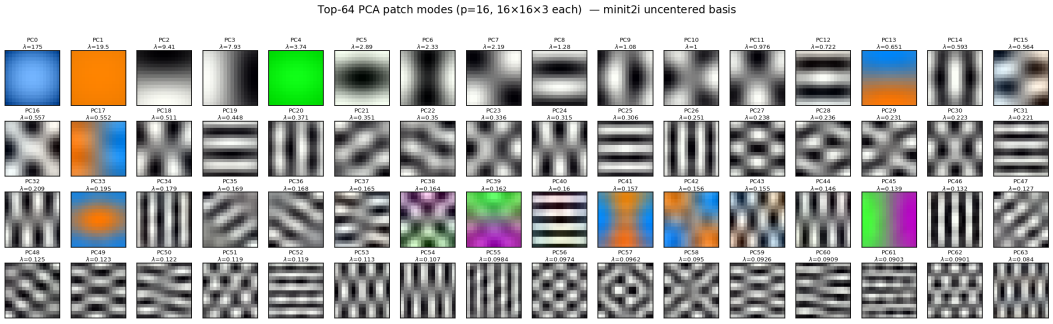


Figure 13: **Leading modes of the uncentered patch-PCA basis.** We reshape the first 64 eigenvectors of the $16 \times 16 \times 3$ patch-PCA basis into RGB patches; the value λ denotes the variance of each direction. The leading ordering interleaves patch-global or very low-frequency modes with spatially varying modes. In particular, PC0 is dominated by low-frequency patch-level tone, PC1 is an approximately uniform chromatic mode, PC2–PC3 encode spatial gradients, and PC4 returns to an approximately uniform chromatic mode. Eigenvector signs are arbitrary, so the interpretation is based on spatial structure rather than the displayed color polarity.

The PCA modes in Fig. 13 provide a structural interpretation of this pattern. PC0 represents a patch-global, low-spatial-frequency component, while PC1 is nearly constant across the patch and primarily changes its color. In contrast, PC2 and PC3 contain vertical and horizontal spatial gradients. PC4 then returns to a nearly spatially uniform chromatic component. The leading sequence is therefore approximately

$$\text{patch-global}, \quad \text{patch-global}, \quad \text{spatial}, \quad \text{spatial}, \quad \text{patch-global},$$

which closely matches the learned gain pattern

$$\text{high}, \quad \text{high}, \quad \text{low}, \quad \text{low}, \quad \text{high}.$$

In particular, PC4 has substantially lower variance than PC2 and PC3 ($\lambda_4 = 3.74$, compared with $\lambda_2 = 9.41$ and $\lambda_3 = 7.93$), yet LPIPS assigns it a larger gain. The spike at PC4 is therefore not explained by variance ordering alone; it coincides with the return of a patch-global chromatic mode after two spatial-gradient modes.

This effect is also more specific than a preference for color. Several later PCA directions contain strong chromatic variation but are spatially non-uniform, and they do not receive the same gain increase. The observed pattern instead suggests that LPIPS distinguishes between patch-global tone or color changes and spatially varying patterns.

One possible explanation follows from the conditional-mean behavior of squared-error regression. In a finite-capacity model, a more $\boldsymbol{x}_0$-like parameterization can shrink uncertain components toward their conditional means. For a spatially varying component, such shrinkage may primarily smooth a local gradient or texture. For a patch-global component, however, it can coherently alter the brightness, contrast, or color of the entire patch. Because LPIPS is applied to the recovered clean image, it may make these perceptually coherent shifts more costly. The learned parameterization responds by retaining a larger velocity-like component for the patch-global modes, thereby reducing their tendency toward mean reversion.

The controlled comparison establishes that adding LPIPS produces the non-monotonic $\boldsymbol{M}$-curve and that the unusually high gains align with patch-global PCA modes. The mean-reversion explanation is a plausible interpretation of this alignment rather than a causal mechanism directly established by the ablation.

A.5 A Capacity-Constrained View of $\boldsymbol{x}_0$- and $\boldsymbol{u}$-Prediction

The choice between clean-data and velocity prediction involves two competing considerations. Clean-data prediction avoids asking the model to reproduce full-dimensional Gaussian noise, which can be advantageous when a high-dimensional representation has much lower intrinsic dimension.

Velocity prediction, on the other hand, provides better numerical conditioning near the clean endpoint. Our RAEv2 experiment illustrates this trade-off particularly clearly.

RAEv2 predominantly favors clean-data prediction. In our clean-latent RAEv2 setup, the pure $\boldsymbol{x}_0$-prediction baseline obtains an FID of 1.05, while AsymFlow with $r = 8$ and MFlow obtain 1.08 and 1.09, respectively, after 80 epochs (Table 10). Consistently, the learned MFlow gains in Fig. 12(c) move almost entirely toward zero, recovering an approximately pure $\boldsymbol{x}_0$ parameterization. These results indicate that, for this representation, model, and training budget, the network prefers to avoid predicting noise in most latent directions.

At the same time, the AsymFlow result is only 0.03 FID behind the pure-$\boldsymbol{x}_0$ baseline, despite using full velocity prediction in its leading eight PCA directions. This result does not show that velocity prediction is superior in this setting; pure $\boldsymbol{x}_0$-prediction remains best in Table 10. It instead shows that a small velocity-prediction subspace can be accommodated at little measured cost, even in a representation that strongly favors $\boldsymbol{x}_0$-prediction overall.

Numerical conditioning versus target complexity. The benefit of retaining a velocity-like component can be understood from the recovery map. Under our convention,

$$\boldsymbol{x}_t = t\boldsymbol{x}_0 + (1-t)\boldsymbol{\epsilon}, \qquad t = 1 \text{ is the clean endpoint.}$$

For pure $\boldsymbol{x}_0$-prediction, the velocity is recovered as

$$\hat{\tilde{u}}_i = \frac{\hat{\tilde{x}}_{0,i} - \tilde{x}_{t,i}}{1-t}. \tag{25}$$

Consequently, an error $\delta_i^x = \hat{\tilde{x}}_{0,i} - \tilde{x}_{0,i}$ in the clean-data prediction produces a velocity error

$$\hat{\tilde{u}}_i - \tilde{u}_i = \frac{\delta_i^x}{1-t}. \tag{26}$$

The factor $1/(1-t)$ becomes ill-conditioned as $t \to 1$.

For MFlow, the corresponding relation is

$$\hat{\tilde{u}}_i - \tilde{u}_i = \frac{\hat{\tilde{u}}_{\mathrm{G},i} - \tilde{u}_{\mathrm{G},i}}{\Delta_i(t)}, \qquad \Delta_i(t) = (1-t) + tm_i. \tag{27}$$

When $m_i = 0$, this reduces to the same $1/(1-t)$ amplification as $\boldsymbol{x}_0$-prediction. For any $m_i > 0$, however,

$$\lim_{t \to 1} \Delta_i(t) = m_i,$$

so the clean-endpoint singularity is removed in that direction. Thus, conditional on comparable target-prediction errors, a nonzero velocity component gives better-conditioned velocity recovery near clean samples.

This conditioning advantage comes with an opposing cost. Increasing m_i changes the target from clean-data prediction toward

$$\tilde{u}_{\mathrm{G},i} = \tilde{x}_{0,i} - m_i\tilde{\epsilon}_i,$$

requiring the finite-capacity network to retain a larger Gaussian-noise component. Full velocity prediction is therefore well conditioned during recovery but can be difficult to approximate in high-dimensional spaces, whereas pure $\boldsymbol{x}_0$-prediction removes the noise-prediction burden but retains the clean-endpoint amplification. The optimal parameterization must balance these two effects.

A soft velocity budget. This trade-off suggests viewing MFlow as learning both *how much* velocity prediction the model can support and *where* it should be allocated. A useful summary of the former is the soft velocity rank

$$r_{\mathrm{eff}} \coloneqq \mathrm{tr}(\boldsymbol{M}) = \sum_i m_i. \tag{28}$$

Pure $\boldsymbol{x}_0$-prediction has $r_{\mathrm{eff}} = 0$, full $\boldsymbol{u}$-prediction has $r_{\mathrm{eff}} = D$, and rank-r AsymFlow has $r_{\mathrm{eff}} = r$. For MFlow, this quantity is continuous and learned.

Our results suggest a capacity-constrained principle: *use velocity-like prediction in as many directions as the model can accurately support, and retain clean-data-like prediction in the remaining*

Table 9: **Learnable gains on UCF-101 (B-level).** Results use the default sampling hyperparameters, without an evaluation sweep.

Method	FVD↓
Video MFLOW, fixed gains	169.4
Video MFLOW, learnable gains	**155.5**

Table 10: **RAEv2 results.** H-level models trained on ImageNet for 80 epochs.

Method	FID↓
RAEv2 Singh et al. (2026)	**1.05**
RAEv2 + Asym8	1.08
RAEv2 + MFLOW	1.09

directions. The RAEv2 curve indicates a small velocity budget for this high-dimensional latent representation, whereas the text-to-image curves show that increasing model capacity shifts more directions toward velocity prediction. We therefore interpret model capacity as influencing the available velocity budget, while the data representation and training objective determine how that budget should be distributed.

Why leading PCA directions are a natural default. Suppose that only N directions can be assigned substantial velocity gains. In the absence of perceptual or other auxiliary losses, the leading PCA directions provide a natural initial allocation. In the PCA basis,

$$\mathrm{Var}(\tilde{x}_{0,i}) = \lambda_i, \qquad \mathrm{Var}(\tilde{\epsilon}_i) = 1, \tag{29}$$

where λ_i decreases with i, while the isotropic noise variance is identical in every direction. The velocity target therefore becomes increasingly noise-dominated as λ_i decreases. Allocating a limited velocity budget to high-variance directions concentrates the noise-prediction burden in the dominant data subspace, while keeping the many low-variance, predominantly off-manifold directions closer to $\boldsymbol{x}_0$-prediction. It also improves near-clean conditioning in the directions carrying the largest fraction of the data variation.

This PCA ordering should nevertheless be viewed as an initialization prior rather than a universal rule. Auxiliary objectives can favor directions according to their visual function rather than variance, as shown by the non-monotonic LPIPS $\boldsymbol{M}$-curve in Fig. 12(a). We also observe that a non-leading PCA block can outperform the first eight directions in one ImageNet pilot. MFlow therefore initializes gains using normalized variance but lets the model modify both the effective velocity budget and its direction-wise allocation through end-to-end training.

A.6 RADIALLY-AVERAGED POWER SPECTRUM

To quantify how faithfully a model reproduces the spatial-frequency content of natural images, we compare the azimuthally-averaged 2D power spectrum of generated and real images. Each image is converted to luma ($Y = 0.299R + 0.587G + 0.114B$), mean-subtracted to remove the DC component, and multiplied by a separable 2D Hann window to suppress boundary leakage. We compute the squared magnitude of its 2D DFT, $P(\mathbf{f}) = |\mathcal{F}g(\mathbf{f})|^2$, center it (fftshift), and azimuthally average it, *i.e.*, average over all orientations at each fixed radius $r = \|\mathbf{f}\|$, to obtain a 1D profile $P(f)$ as a function of radial frequency f (cycles/pixel). Averaging $P(f)$ over $N = 4000$ images gives the RAPS of a distribution.

A.7 REFINERS

For our *ConvRefiner* we use a t-conditioned PatchGAN discriminator (Isola et al. (2017)) (NLayerDiscriminator, 3 layers, ndf=64) with a hinge adversarial loss, applied to the refined $\hat{\boldsymbol{x}}_0$. We have two types of batches – all-t batch and near-clean batch ($t \sim \mathcal{U}(0.9, 0.98)$). For the all-$t$ batch the flow-matching and LPIPS($t \geq 0.3$) losses are normally applied. For the near-clean batch, additional GAN loss is applied with a 0.02 weight, if not otherwise specified, and the flow-matching and LPIPS loss are downweighted by a factor of 0.25. For Attention Refiner we used a 2 blocks of swin-style (Liu et al. (2021)) windowed attention residual refiner in the clean image space, as a drop in replacement for the *ConvRefiner*. Note that DiP patch-detailer also takes as input $\boldsymbol{x}_t$ apart from the model output.

A.8 LIMITATIONS.

Although MFLOW is a general formulation that transfers across pixel-space image and video generation, high-dimensional representation latents such as RAEv2 Singh et al. (2026), and text-to-image



Figure 14: **ImageNet Loss Curves at H-level.** We usually observe that the $\boldsymbol{u}$-loss of MFLOW training is very slightly better.

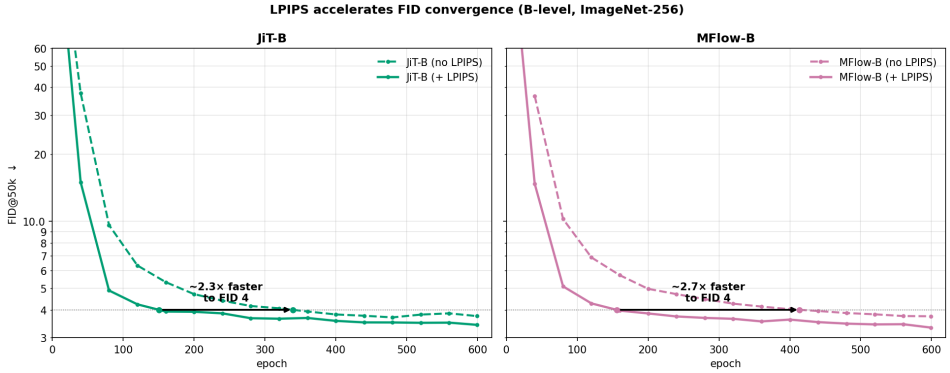


Figure 15: **Faster convergence on ImageNet with LPIPS at B-level.** This corresponds to the results noted in Table 2.

models of different capacities, its improvements over strong fixed parameterizations are sometimes modest. One possible reason is the joint optimization of the gains and network weights. While the underlying velocity target remains fixed, updating $\boldsymbol{M}$ continuously changes the network's direct output parameterization $\boldsymbol{u}_{\mathrm{G}} = \boldsymbol{x}_0 - \boldsymbol{M}\boldsymbol{\epsilon}$, creating an effectively moving regression target and allowing the gains and network to co-adapt. Indeed, under the same velocity-space objective, MFLOW attains a slightly lower training loss than JiT and AsymFlow on the ImageNet H-level experiment (Fig. 21); we observe similar trends for B-level models with and without LPIPS. However, the lower diffusion loss does not consistently translate into better FID. This mismatch suggests that improved regression conditioning need not directly improve distribution-level sample quality, and its precise cause remains unclear. Staged or decoupled optimization of the gains is an important direction for future work.

ConvRefiner reveals a complementary evaluation limitation. It substantially closes the high-frequency spectral gap and brings the generated spectrum closer to that of real images, yet this improvement produces little or no change in FID. This suggests that FID is relatively insensitive to the local frequency statistics modified by the refiner. Consequently, spectral fidelity and FID should be viewed as complementary rather than interchangeable measures; texture-sensitive metrics and human evaluation are needed to determine how strongly the recovered high-frequency content translates to perceived image quality.

Figure 16: **Qualitative Text-to-Image Results on DPG-Bench.** Images generated by our method and baselines on DPG-Bench prompts.



Figure 17: **Qualitative Text-to-Image Results on PRISM-Bench.** Images generated by our method and baselines on PRISM-Bench prompts.

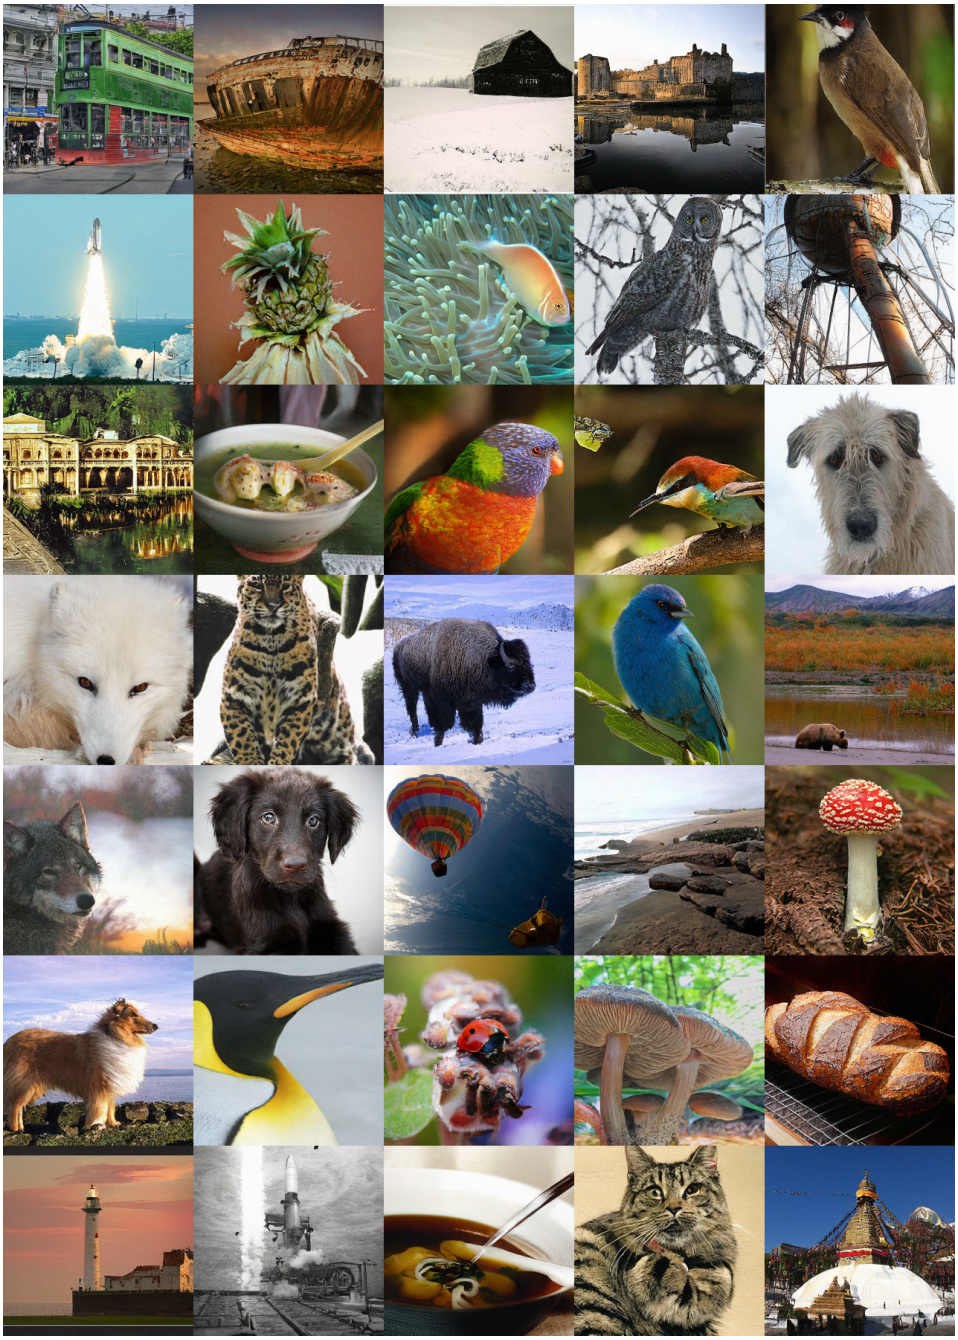


Figure 18: **Qualitative ImageNet Results.** Images generated by MFLOW +*ConvRefiner* H-level model.



Figure 19: **More Class-wise Uncurated Qualitative ImageNet Results.** Images generated by MFLOW +*ConvRefiner* H-level model.

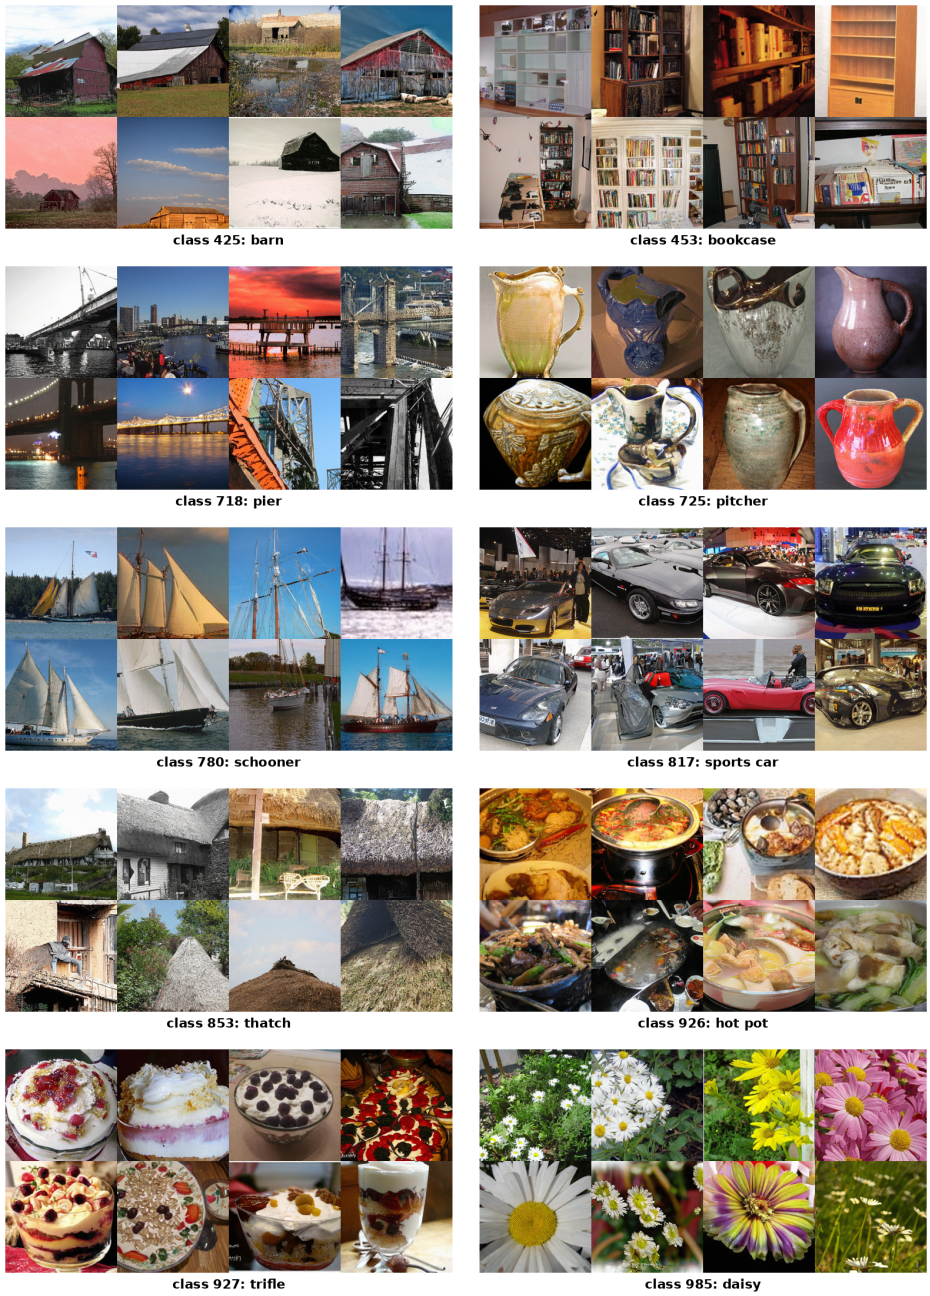


Figure 20: **More Class-wise Uncurated Qualitative ImageNet Results.** Images generated by MFLOW +*ConvRefiner* H-level model.

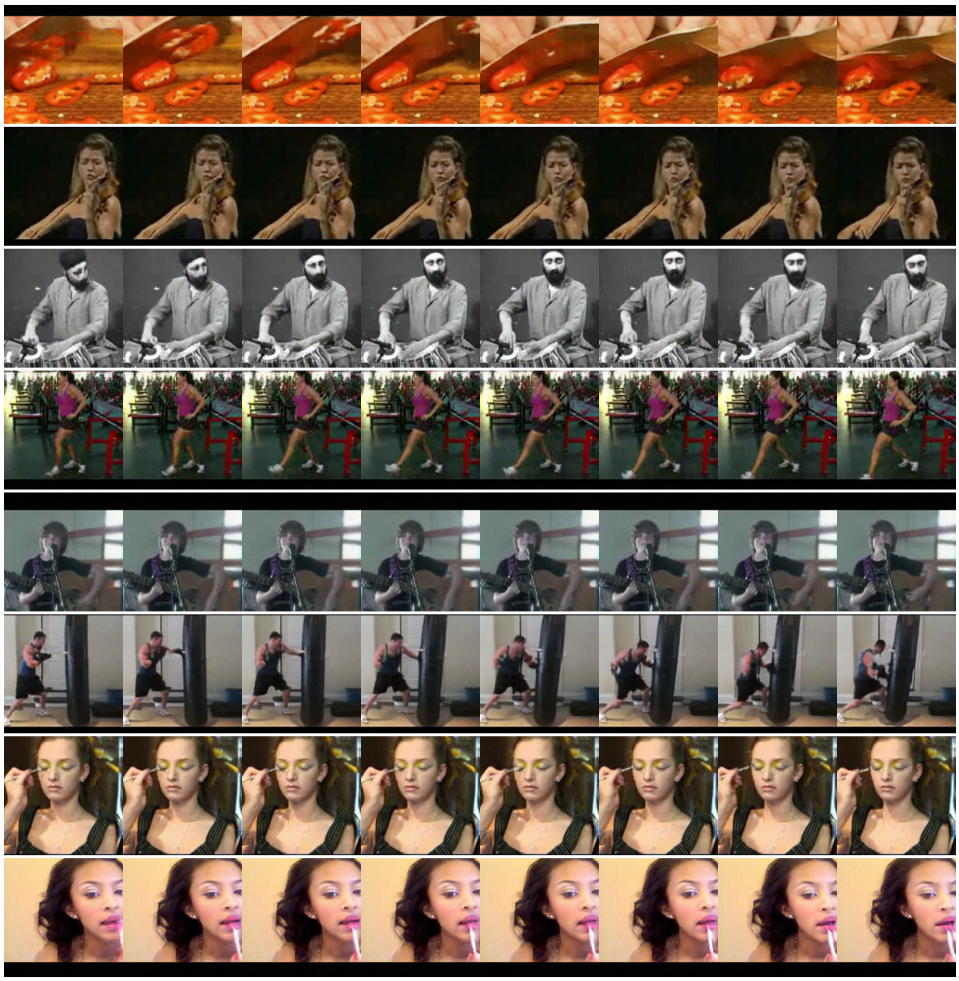


Figure 21: **Qualitative UCF-101 Results.** Frames generated by MFLOW +post-hoc FW-LPIPS L-level model.



Figure 22: **Visualizing effect of *ConvRefiner*.** Here we visualize the effect of adding our post-training to MFLOW-H+LPIPS on ImageNet generations. We also note the average power gain in the high-frequency(0.5–0.7) region. We see overall improvement in texture details.

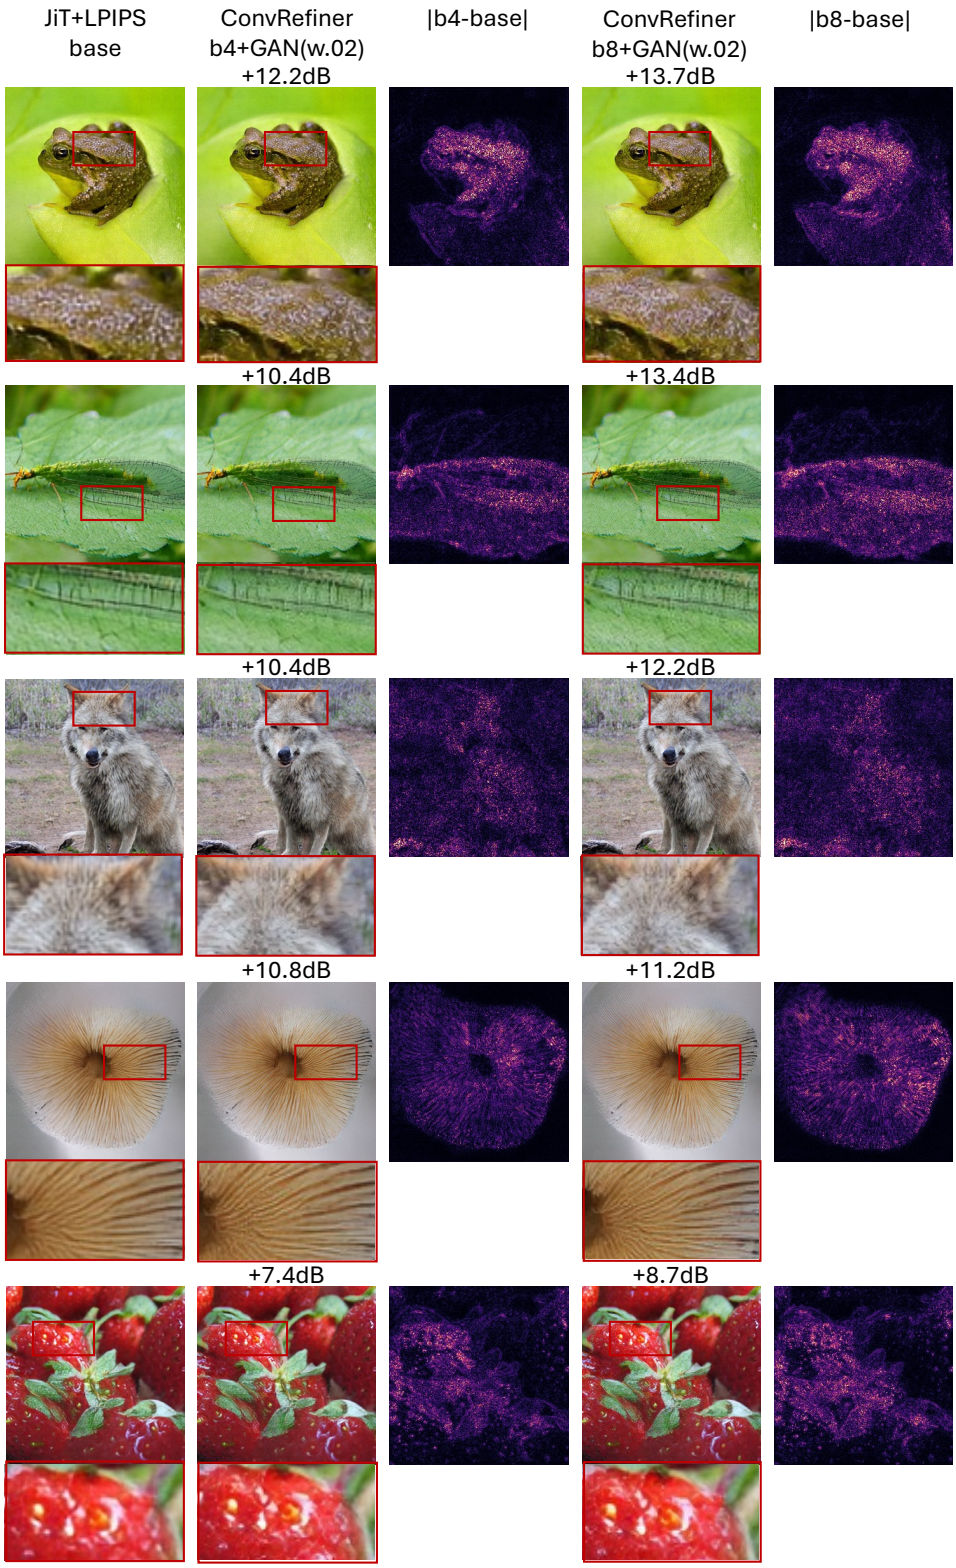


Figure 23: **Qualitative Ablation for *ConvRefiner*.** Here we visualize the effect of adding our post-training to JiT-B on ImageNet generations. We also note the average power gain in the high-frequency(0.5–0.7) region. We also show the visual improvements – we see that b4 adds additional texture details, while b8 sometimes overshoots in texture.