

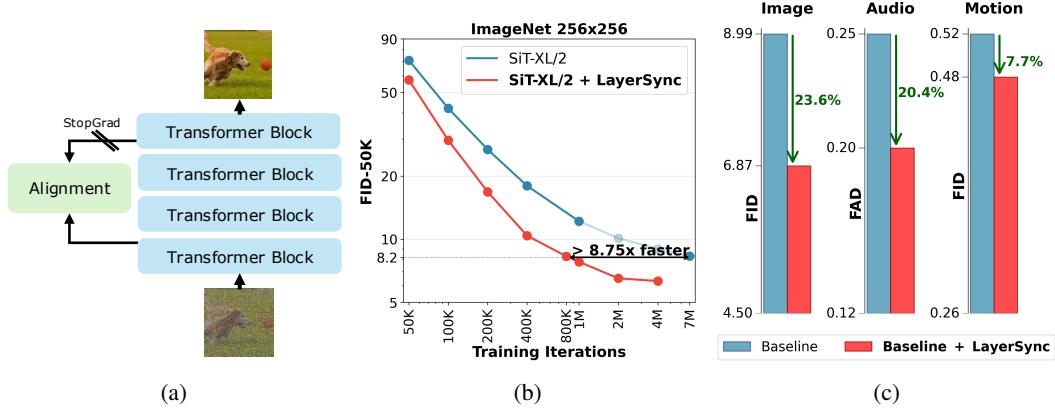
000 LAYERSYNC: SELF-ALIGNING INTERMEDIATE LAY- 001 ERS

002 **Anonymous authors**

003 Paper under double-blind review

009 ABSTRACT

011 We propose LayerSync, a domain-agnostic approach for improving the generation
 012 quality and the training efficiency of diffusion models. Prior studies have high-
 013 lighted the connection between the quality of generation and the representations
 014 learned by diffusion models, showing that external guidance on model interme-
 015 diate representations accelerates training. We reconceptualize this paradigm by
 016 regularizing diffusion models with their own intermediate representations. Build-
 017 ing on the observation that representation quality varies across diffusion model
 018 layers, we show that the most semantically rich representations can act as an in-
 019 trinsic guidance for weaker ones, reducing the need for external supervision. Our
 020 approach, LayerSync, is a self-sufficient, plug-and-play regularizer term with no
 021 overhead on diffusion model training and generalizes beyond the visual domain
 022 to other modalities. LayerSync requires no pretrained models nor additional data.
 023 We extensively evaluate the method on image generation and demonstrate its ap-
 024 plicability to other domains such as audio, video, and motion generation. We show
 025 that it consistently improves the generation quality and the training efficiency. For
 026 example, we speed up the training of flow-based transformer by over $8.75 \times$ on
 027 ImageNet dataset and improve the generation quality by 23.6%.



042 **Figure 1: LayerSync improves training efficiency and generation quality via internal repre-
 043 sentation alignment.** (a) LayerSync aligns deep and shallow layers. (b) LayerSync achieves over
 044 8.75 \times training acceleration on the ImageNet 256x256. (c) LayerSync consistently improves gen-
 045 eration quality across multiple modalities: by 23.6% on FID for images (ImageNet 256x256), 24%
 046 on FAD for audio (MTG-Jamendo), and 7.7% for FID on human motion (HumanML3D).

049 1 INTRODUCTION

051 Denoising generative models, such as diffusion (Ho et al., 2020; Song et al., 2020; Song & Ermon,
 052 2019) and flow matching models (Lipman et al., 2023), have demonstrated remarkable success in
 053 modeling complex data distributions, achieving state-of-the-art performance across a range of gen-
 erative tasks. However, this success comes at a significant computation cost. Thus, a new promising

line of research has emerged to improve the training efficiency of these models by improving the models' intermediate representations (Yu et al., 2024; Wang et al., 2025; Wang & He, 2025). It has been shown that the quality of a diffusion model's intermediate representations is intrinsically linked to its generative performance. As a result, explicitly guiding these representations accelerates training and improves generation quality (Yu et al., 2024).

Building on this insight, the most dominant approach (Yu et al., 2024; Wang et al., 2025) has been to leverage powerful external guidance from large pre-trained models, by aligning the internal features of a diffusion model with those of high-capacity vision models like DINOv2 (Oquab et al., 2023) or vision-language models (VLMs) like Qwen2-VL (Wang et al., 2024). These methods demonstrate that access to strong semantic features can accelerate training by an order of magnitude. While effective, this paradigm comes with several limitations. It introduces a dependency on massive external models that are themselves costly to train, require large amounts of data, and may not be available for domains beyond natural images. Additionally, this reliance on external data and parameters introduces extra overhead into the training pipeline. For instance, in the case of Wang et al. (2025), training is indeed faster in terms of iterations but involves calling a 9-billion-parameter VLM at each step. These limitations motivate the development of more self-contained and generalizable alternatives.

A recent step in this direction is the Dispersive Loss (Wang & He, 2025), a self-contained regularizer that encourages internal representations to spread out in the feature space, analogous to the repulsive force in contrastive learning (Oord et al., 2018). Although this approach demonstrates the potential for internal regularization, a substantial performance gap remains compared to methods that leverage external representations. In this paper, we propose a self-contained method with a more directed learning signal to reduce this gap.

Our work is motivated by two key observations: First, while diffusion models learn powerful representations, their quality is highly heterogeneous across the model's depth. As demonstrated by previous works (Mukhopadhyay et al., 2024; Ghadiyaram, 2025; Stracke et al., 2025) certain intermediate layers capture more semantically rich and useful information than others. Second, when models incorporate knowledge through external guidance using DINOv2 for instance, regularizing early layers seemed more effective than regularizing the deeper ones (Yu et al., 2024; Wang et al., 2025). Upon these two observations a clear opportunity presents itself: can the model's own strongest layers act as an intrinsic guidance to improve its weaker ones through self-alignment?

To this end, we propose LayerSync, a simple yet powerful regularization framework that aligns a model's own intermediate layers. LayerSync is a parameter-free, plug-and-play solution that operates without any external models or data, making it a truly self-contained method. LayerSync introduces negligible computational overhead, yet its effectiveness is substantial. Our experiments show that LayerSync consistently outperforms prior self-contained methods across all tested configurations. For image generation, LayerSync accelerates training on ImageNet 256×256 (Deng et al., 2009) by more than 8.75×. This leads to a new state-of-the-art in purely self-supervised image generation on ImageNet, demonstrating the strength of our self-alignment objective and substantially narrowing the gap between self-contained approaches and those relying on external guidance. Furthermore, due to its self-contained nature, LayerSync seamlessly generalizes to other modalities. Our experiment shows that for audio generation LayerSync leads to 21% improvement in FAD-10K on MTG-Jamendo (Bogdanov et al., 2019), to 7.7% improvement in FID for human motion generation on HumanML3D and 54.7% in FVD for video generation on CLEVRER (see Appendix A). To the best of our knowledge, it is the first time that a self-contained method proves to accelerate diffusion models training seamlessly across different domains.

Additionally, an analysis of the internal features confirms that LayerSync strengthens the model's representations, leading to a 32.4% improvement in classification and a 63.3% improvement in semantic segmentation.

Our main contributions are as follows:

- We introduce LayerSync, a minimalist, parameter-free, and self-contained regularization method that leverages a diffusion model's own layers as an intrinsic guidance via self-alignment.

- 108 • We demonstrate the domain-agnostic versatility of LayerSync by successfully applying it
109 to image, audio, human motion and video generation.
- 110 • We show that our self-supervised method not only accelerates training but also improves
111 the representations across the model’s layers.

113 2 RELATED WORK

116 **Representation learning with diffusion models.** Denoising Generative Models including both
117 diffusion (Ho et al., 2020; Song et al., 2020; Song & Ermon, 2019) and flow matching models
118 (Lipman et al., 2023), trained as multi-level denoising autoencoders (Vincent et al., 2008), naturally
119 give rise to discriminative representations. A line of work has specifically evaluated the quality
120 of these representations, showing that diffusion features can be effectively used across a variety of
121 tasks (Mukhopadhyay et al., 2024; Ghadiyaram, 2025; Stracke et al., 2025) and, in some cases,
122 achieve performance comparable to self-supervised representation learning methods (Stracke et al.,
123 2025). However, the quality of the representations varies across model layers, with the final layers,
124 just before the model begins decoding, consistently containing more semantically rich features
125 Ghadiyaram (2025); Xiang et al. (2023), regardless of whether the architecture is a U-Net (Ron-
126 neberger et al., 2015) or a Transformer (Vaswani et al., 2017). Our work is directly built upon those
127 insights. We demonstrate that the semantically rich representations in the intermediate layers can be
leveraged as a guidance signal to enhance the quality of earlier-layer representations.

128 **Representation regularization for improving diffusion models.** It has been shown that represen-
129 tation quality is closely linked to generative performance (Yu et al., 2024). One line of work im-
130 proves generation by regularizing model representations through alignment with strong pretrained
131 networks. For example, Yu et al. (2024) aligns diffusion features with self-supervised features
132 from DINOv2 (Oquab et al., 2023), while Wang et al. (2025) demonstrates that leveraging vi-
133 sion-language models (VLMs) (Wang et al., 2024) can yield further improvements. Although
134 such approaches accelerate training and enhance generation quality, they remain constrained by the
135 need for high-quality external representations, which are not readily available in non-visual domains.
136 Additionally, they introduce computational overhead, as pretrained models must be inferred at each
137 training step. Another group of work adopts self-supervised strategies that rely on EMA (Expo-
138 nential Moving Average) (Tervainen & Valpola, 2017) models to guide the representations. Zheng
139 et al. (2023) integrates a generative diffusion process with an auxiliary mask reconstruction task.
140 Zhu et al. (2024); Jiang et al. (2025) align representations between teacher and student encoders
141 in a joint embedding space. While being self-contained, such methods increase computational cost
142 by requiring an additional forward pass through the EMA model at each training step. Also, their
143 performance still lags behind methods that leverage external supervision. A recent work (Wang &
144 He, 2025) proposes dispersing representations in the feature space, analogous to the repulsive force
145 in contrastive learning. This approach introduces no additional training overhead. Similarly, we
146 present a self-contained, overhead-free solution; however, we leverage semantically richer internal
representations to guide and improve the learning of weaker ones.

147 3 METHOD

150 3.1 PRELIMINARIES

152 We adopt the generalized perspective of stochastic interpolants (Ma et al., 2024), which provides a
153 unifying framework for both flow-based and diffusion-based models. Here is a brief overview, we
154 refer to the Appendix Section N for more details.

155 Stochastic interpolants are generative models that learn to reverse a process that gradually converts
156 a data sample \mathbf{x}_0 into simple noise ϵ . This is achieved by defining a path between them:

$$157 \quad \mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \epsilon, \\ 158$$

159 where α_t and σ_t are functions of time controlling the mix of data and noise at time t . To generate
160 new data, the model must learn to travel backward along this path, from noise to data. The direction
161 and speed at any point \mathbf{x}_t and time t is given by a velocity field. The true velocity is the time
derivative of the path: $\dot{\alpha}_t \mathbf{x}_0 + \dot{\sigma}_t \epsilon$.

162 Since this true velocity is unknown during generation, a neural network $v_\theta(\mathbf{x}_t, t)$ is trained to predict
 163 it. The model learns by minimizing the velocity loss, which measures the squared difference between
 164 the predicted velocity and the ground-truth velocity:
 165

$$\mathcal{L}_{\text{velocity}}(\theta) := \mathbb{E}_{\mathbf{x}_0, \epsilon, t} \left[\|v_\theta(\mathbf{x}_t, t) - (\dot{\alpha}_t \mathbf{x}_0 + \dot{\sigma}_t \epsilon)\|^2 \right]. \quad (1)$$

168 Once trained, the model can generate new data by starting with a random noise sample and following
 169 the velocity field it has learned.
 170

172 3.2 DIFFUSION MODELS INTERMEDIATE REPRESENTATIONS

174 **Representation Hierarchy.** We investigate the representations learned by a pre-trained SiT model
 175 (Ma et al., 2024) on ImageNet (Deng et al., 2009). Through linear probing on downstream tasks
 176 (classification, segmentation) and Centered Kernel Alignment (CKA) with DINOv2 (Oquab et al.,
 177 2023) features, we observe a clear hierarchy in representation quality across layers. As shown
 178 in Figure 4c, deeper layers exhibit superior discriminative capabilities, consistent with established
 179 principles of hierarchical feature learning in deep networks (LeCun et al., 2015). This pattern of
 180 increasing semantic richness culminating in a peak before the final decoding blocks is a known
 181 characteristic of diffusion model representations (Mukhopadhyay et al., 2024).
 182

183 **Internal Block Structure of Diffusion Transformers.** Beyond the
 184 hierarchy of feature quality, we observe that the internal structure
 185 of Diffusion Transformers follows specific correlation patterns as
 186 shown in Figure 2 and discussed in (Raghu et al., 2021; An et al.,
 187 2024). These blocks are highly correlated and naturally segregate
 188 into three functional groups at convergence: (1) an initial group fo-
 189 cusing on local features, (2) a middle group of highly correlated
 190 blocks capturing global features, and (3) a final group acting as a
 191 decoder to project back to the latent space.
 192

193 We propose to internally regularize the network by aligning representations from its early layers with
 194 those from its own deeper, semantically richer layers. Beyond the immediate benefit of improving
 195 shallow layers, we hypothesize that this method may induce a recursive refinement process. The
 196 enhancement of early-layer features is expected to facilitate the learning of more robust deep-layer
 197 representations, which subsequently offer a more refined target for the internal alignment, potentially
 198 leading to a cascading improvement of the model’s feature space.
 199

200 3.3 LAYER SYNC

202 We propose a self-contained regularization approach, named **LayerSync**, designed to improve the
 203 generation quality and training dynamics of diffusion models. The core principle behind LayerSync
 204 is intra-model self-alignment, where the model is trained to guide itself. We use the context-rich
 205 deep layers as an “intrinsic guidance” to provide a direct signal to the earlier “weak” layers, thereby
 206 enhancing the model’s entire feature from within.
 207

208 LayerSync achieves this alignment by maximizing the similarity between the feature representations
 209 of designated strong and weak blocks. The similarity is computed for each patch in the represen-
 210 tation and then averaged over the whole patch sequence for each image. Let f_θ be the network
 211 transformer and let f_θ^k designate the network up to the k -th layer. Let \mathbf{x}_t be the input marginal dis-
 212 tribution at time t , with $t \sim \text{Uniform}(0, 1)$ and $x \sim \mathbf{x}_t$, we define the loss for LayerSync between
 213 layer k and k' with $k < k'$ as follows:
 214

$$\mathcal{L}_{\text{LayerSync}_{(k, k')}}(\theta) := -\mathbb{E}_{\mathbf{x}_t, t} \left[\frac{1}{N} \sum_{n=1}^N \text{sim} \left(f_\theta^k(x)^{[n]}, \text{stopgrad} \left(f_\theta^{k'}(x)^{[n]} \right) \right) \right], \quad (2)$$

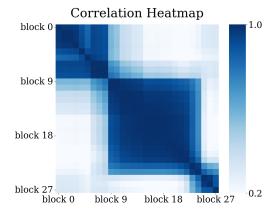


Figure 2: Correlation between the blocks of SiT-XL/2 at Convergence.

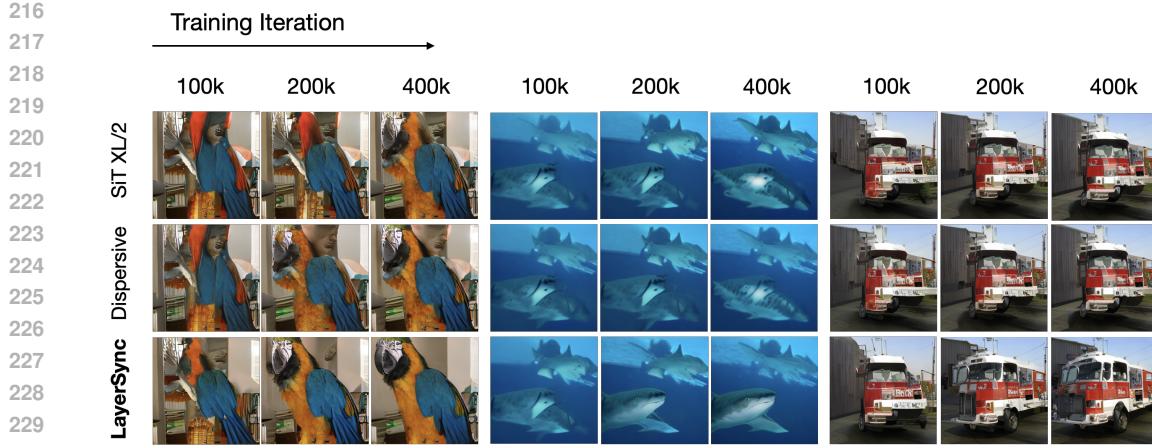


Figure 3: **LayerSync improves generation quality without relying on external representation.** We compare the images generated by SiT-XL/2 when regularized with dispersive and LayerSync. All the models are trained for 400K iterations, share the same noise, sampler, and number of sampling steps, and none of them use classifier-free guidance.

where n is the patch index and $\text{sim}(\cdot, \cdot)$ is a pre-defined similarity function. We experimented with different similarity functions and opted for cosine similarity in all our following experiments. This regularization term is integrated into the primary training objective as a weighted sum:

$$\mathcal{L} := \mathcal{L}_{\text{velocity}} + \lambda \mathcal{L}_{\text{LayerSync}}, \quad (3)$$

where the hyperparameter $\lambda > 0$ balances the standard denoising task with our internal representation alignment. Algorithm 1 in Appendix summarizes our proposed approach.

3.4 LAYER SELECTION

A key design consideration for LayerSync is the selection of the layers to align. We note that this is a shared characteristic with other representation guidance methods, in particular those that rely on external supervision. As detailed in Section 3.2, the layers of Diffusion Transformers naturally converge toward a structure in three groups of blocks focusing on local features, global features, and decoding. While our experiments indicate that LayerSync yields performance gains across a wide range of configurations, we observe that respecting this structure leads to optimal acceleration (see Section D).

Based upon these observations, we propose a selection strategy guided by three principles. First, drawing on established findings (Xiang et al., 2023) regarding the functional specialization of layers in generative transformers, we exclude the final 20% of blocks from being chosen as the reference layer as those are primarily specialized for low-level decoding tasks making them suboptimal as guidance targets. Second, based on the finding from (Raghu et al., 2021), we exclude the very first layers as having the earliest blocks focus on local features improves the model performance and generalization An et al. (2024). Third, to ensure a meaningful semantic gap between the representations, we enforce a minimum distance between the aligned and reference layers (e.g., 8 blocks for SiT-XL and 3 for SiT-B). We validate this heuristic and the robustness of the method to different layer selection through experiments as summarized in Tables 6, 11 and 12.

4 EXPERIMENTS

We conduct a comprehensive set of experiments to validate the effectiveness of LayerSync. Our evaluation is structured along three axes:

- We first study extensively the performance and training efficiency of LayerSync in large-scale class-conditional image generation (Section 4.1).

270

- 271 We then assess the domain-agnostic capabilities of our method by applying it to generative
- 272 tasks in audio (Section 4.2), human motion (Section 4.3), and video (Appendix Section A).
- 273
- 274 Finally, we perform an in-depth analysis to quantify the impact of LayerSync on the quality
- 275 and structure of the learned internal representations (Section 4.4).

276 **4.1 IMAGE GENERATION**

277 **Implementation details.** We strictly follow the setup in SiT (Ma et al., 2024). Specifically, we use
 278 ImageNet (Deng et al., 2009) and follow ADM (Dhariwal & Nichol, 2021) for data preprocessing.
 279 The processed image will have the resolution of 256×256 and is then encoded into a compressed
 280 vector $z \in \mathbb{R}^{32 \times 32 \times 4}$ using the Stable Diffusion VAE (Rombach et al., 2022). For model configu-
 281 rations, we use the B/2, L/2, and XL/2 architectures by Ma et al. (2024), which process inputs with
 282 a patch size of 2. More details about the architectures and the number of parameters are provided
 283 in the Section Q. We use cosine similarity between the patches as the similarity metric. Additional
 284 experimental details, including hyperparameter settings and computing resources, are provided in
 285 Section L.

286 **Evaluation metrics.** We report Frechet inception distance (FID; Heusel et al. (2017)), Inception
 287 Score (Salimans et al., 2016), Precision, and Recall (Kynkäanniemi et al., 2019) using 50,000
 288 samples. We provide details of each metric in Section O.

289 **Baselines.** We compare our results with Dispersive (Wang & He, 2025), the only self-contained,
 290 zero-cost method that accelerates training. For the sake of completeness, we also compare our
 291 method with several recent diffusion-based generation methods. For pixel-based approaches we
 292 compare with ADM (Dhariwal & Nichol, 2021), VDM++ (Kingma & Gao, 2023), Simple diffusion
 293 (Hoogeboom et al., 2023), CDM (Ho et al., 2022). For latent-based approaches we compare with
 294 LDM (Rombach et al., 2022), U-ViT-H/2 (Bao et al., 2023), DiffiT (Hatamizadeh et al., 2024),
 295 MDTv2-XL/2 (Gao et al., 2023), MaskDiT (Zheng et al., 2023), SD-DiT (Zhu et al., 2024), DiT
 296 (Peebles & Xie, 2023), and SiT (Ma et al., 2024). We also compare our approach with REPA (Zhang
 297 et al., 2025) and REED (Wang et al., 2025) which rely on external representations. **Additionally, we**
 298 **compare to two autoregressive methods, VAR (Tian et al., 2024) and D-JEPA (Chen et al., 2024).**
 299 **For both methods, we selected the models with similar number of parameters to SiT-XL. Finally, we**
 300 **compare our results with the EMA-based method SRA (Jiang et al., 2025).** A detailed description of
 301 each baseline is provided in Section P.

302 **Results.** As shown in Table 1, our method consistently improves diffusion transformer training and
 303 is more effective than Wang & He (2025). Our method results in $8.75 \times$ acceleration compared to
 304 SiT-XL baseline, reaching an FID of 8.29 after only 160 epochs, and in $4.7 \times$ acceleration compared
 305 to the baseline trained with Dispersive Loss. In table 2 we compare LayerSync with recent state-
 306 of-the-art diffusion model approaches. In particular, on SiT-XL/2, we reach FID 1.89 after 800
 307 epochs setting a new state-of-the-art in pure self-supervised generation, decreasing the gap with
 308 methods like Yu et al. (2024) that rely on external representations. We also qualitatively compare
 309 the progression of generation results in Figure 3, where we use the same initial noise across different
 310 models. **Additional comparison metrics with SRA are provided in Table 15.**

311 **4.2 AUDIO GENERATION**

312 **Implementations details.** We use the MTG-Jamendo dataset (Bogdanov et al., 2019), a large-scale
 313 collection containing over 55,000 full-length songs. For training, we process the audio by creating
 314 random 10-second samples, which are sampled at a standard rate of 44.1 kHz. We condition using
 315 the metadata provided with the dataset by conditioning the generation on the genre and instrument
 316 labels associated with each samples. Our audio generation model is an adaptation of the Scalable
 317 Interpolant Transformer (SiT-XL) (Ma et al., 2024), consistent with the 28-layer architecture used in
 318 our vision experiments. The model is configured to operate on patchified latent representations with
 319 a patch size of one. These latents are obtained from the pre-trained Variational Autoencoder (VAE)
 320 of the Stable Audio Open model (Evans et al., 2025), which provides a compact representation of
 321 the raw audio waveforms. The model was trained on 64 GH200 GPUs with a global batch size of
 322 1024. In our experiment we align layer 8 with 21 using cosine similarity between the patch-wise
 323 representations of these two layers.

Table 1: FID comparisons of class-conditional generation on ImageNet 256×256. No classifier-free guidance (CFG; Ho & Salimans (2021)) is used. The sampler used is the ODE-based Heun method, except for the last section, which uses the SDE-based Euler method following Ma et al. (2024)

Model	#Params	Epochs.	FID \downarrow
SiT-B/2	130M	80	36.19
+ Dispersive	130M	80	32.45 (10.3%)
+ LayerSync	130M	80	30.00 (17.1%)
SiT-L/2	458M	80	21.41
+ Dispersive	458M	80	16.68 (22.1%)
+ LayerSync	458M	80	14.83 (30.7%)
SiT-XL/2	675M	80	17.97
+ Dispersive	675M	80	15.95 (11.3%)
+ LayerSync	675M	80	11.24 (37.5%)
SiT-XL/2	675M	200	12.18
+ Dispersive	675M	200	10.64 (12.6%)
+ LayerSync	675M	200	8.28 (32.0%)
SiT-XL/2	675M	400	10.11
+ Dispersive	675M	400	8.81 (12.9%)
+ LayerSync	675M	400	6.94 (31.4%)
SiT-XL/2	675M	800	8.99
+ Dispersive	675M	800	8.08 (10.1%)
+ LayerSync	675M	800	6.87 (23.6%)
SiT-XL/2 (w/ SDE)	675M	1400	8.3
+ Dispersive	675M	800	7.71 (7.1%)
+ LayerSync	675M	≥ 1200	7.43 (10.5%)
+ LayerSync	675M	160	8.29 (0.1%)
+ LayerSync	675M	200	7.78 (6.3%)
+ LayerSync	675M	400	6.51 (21.6%)
+ LayerSync	675M	800	6.32 (23.9%)

Table 3: Quantitative results for audio generation on the MTG-Jamendo dataset. We report Fréchet Audio Distance (FAD) using CLAP embeddings. Our method significantly outperforms the baseline with no change in parameter count.

Method	#Params	Epoch	FAD (CLAP) \downarrow
SiT-XL (baseline)	756M	465	0.333
+ LayerSync (Ours)	756M	465	0.263 (21.0%)
SiT-XL (baseline)	756M	650	0.251
+ LayerSync (Ours)	756M	650	0.199 (20.7%)

Evaluation metrics. To quantitatively assess the quality and realism of the generated audio, we report the Fréchet Audio Distance (FAD)(Kilgour et al., 2019) with 10,000 samples using the widely-used CLAP embeddings (Zhao et al., 2023).

Results. LayerSync improves the final FAD-10K by 20.7% at 650 epochs as seen in Table 3. The model trained with LayerSync reaches the final performance of the baseline model around epoch 500 so 150 epochs earlier. The convergence speed is therefore improved by 23%.

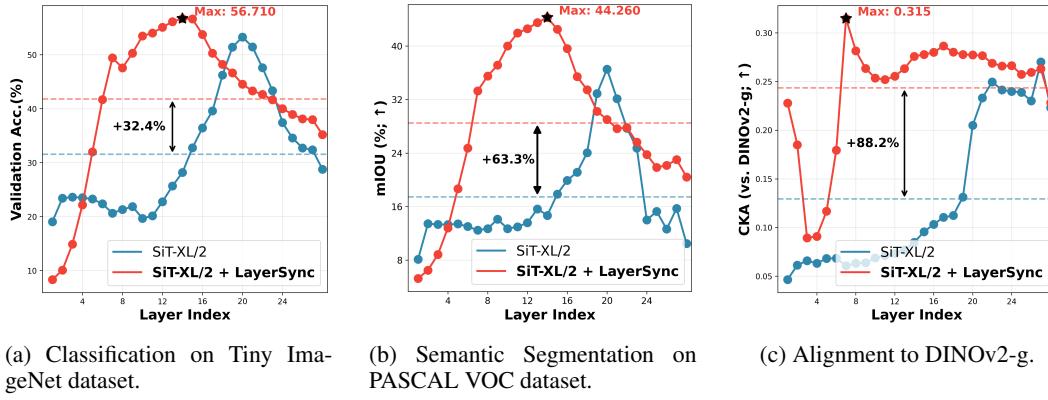
4.3 TEXT-CONDITION HUMAN MOTION GENERATION

To demonstrate that LayerSync can be applied in domains with limited datasets and compact architectures, we consider the task of human motion generation. Given a sentence that describes a motion as a sequence of actions, the task is to generate the corresponding human motion. Each motion sequence consists of a series of human poses, where each pose is represented by 22 joints defined as 3D points in space.

Implementation details. We follow the exact setup as MDM (Tevet et al., 2023) using a transformer with 8 layers. We use HumanML3D dataset (Guo et al., 2022a). We train the model with and without LayerSync for 600K iterations. We align layer 3 with 6. More details are provided in Section M

378
379 Table 4: Quantitative results for text-conditional human motion generation task on HumanML3D
380 dataset. LayerSync improves both FID and R-Precision.

Method	Iter.	FID \downarrow	R-Precision \uparrow
MDM	600K	0.5206	0.7202
+ LayerSync (Ours)	600K	0.4801 (7.7%)	0.7454 (3.4%)



385
386 (a) Classification on Tiny Image- (b) Semantic Segmentation on
387 net dataset. (c) Alignment to PASCAL VOC dataset.

388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
Figure 4: Assessing the quality of intermediate features shows that LayerSync improves average validation accuracy across layers (shown with dashed lines in the figures) for both classification and segmentation, and enhances alignment with DINOv2. In this experiment, layer 8 is aligned with layer 16.

Evaluation metrics. We report FID and R-Precision (top 3) (Kynkänniemi et al., 2019) as detailed in Section O.

Results. The results summarized in Table 4 show that LayerSync improves FID by 7.7% and R-Precision by 3.4%, confirming its effectiveness even with small architectures and limited datasets.

4.4 REPRESENTATION LEARNING

To evaluate the effect of LayerSync, we analyze the model’s intermediate representations. We compare SiT-XL/2 model trained with LayerSync for **160 epochs** against a baseline SiT-XL/2 trained for **1400 epochs** as they both have similar FID. This ensures that both models exhibit comparable generative performance, allowing us to isolate the impact of our regularization on the learned representations, independently of the final generation quality.

We consider linear probing for classification on Tiny ImageNet dataset (Deng et al., 2009), linear probing for segmentation on the PASCAL VOC dataset (Everingham et al., 2010), and Centered Kernel Alignment (CKA) (Kornblith et al., 2019) with DINOv2 embeddings (Oquab et al., 2023) to measure the distance between the model representations. More implementation details are provided in Section L.

Our empirical results, summarized in Figure 4, lead to two interesting observations. First, LayerSync induces a more homogeneous distribution of high-quality features across the network’s layers, leading to 32.4 % improvement in the average validation accuracy for classification, 63.3 % in average mIoU, and 88.2 % improvement in average alignment with DINOv2. Secondly, we observe not only a shift in the block with the best performance in downstream tasks but also an improvement in the best performing block. While an increase in mean performance is an intuitive consequence of regularizing weaker layers toward a high-performing one, the emergence of a new peak that significantly surpasses the baseline’s maximum is a non-trivial finding.

We conclude that the representational benefits of LayerSync are not merely a byproduct of accelerated convergence. Even when the baseline model is afforded more than 8x larger training budget to match generative performance, its internal representations remain significantly inferior. We therefore hypothesize that LayerSync acts as a powerful structural regularizer that fundamentally

432 Table 5: **Quantitative evaluation of SiT-XL/2 trained on ImageNet 256×256 (Deng et al., 2009)**
 433 **comparing REPA and REPA + LayerSync at 200k steps.** The results show that LayerSync can be
 434 combined with approaches that rely on external representations, such as REPA, to further accelerate
 435 training.

Method	FID \downarrow	sFID \downarrow	Inception Score \uparrow
REPA	7.88	4.81	126.39
REPA + LayerSync	7.01	4.78	129.85

442 alters the model’s optimization trajectory. By imposing an internal semantic constraint, it guides
 443 the network to discover a more efficient and globally coherent feature hierarchy, one that remains
 444 inaccessible to the unconstrained model.

445 Furthermore, it is noteworthy that the quality of representations learned with LayerSync approaches
 446 that of models trained with powerful external guidance. Our peak classification accuracy, for ex-
 447 ample, is comparable to the results reported by (Yu et al., 2024). We interpret this as evidence for
 448 our initial hypothesis: LayerSync establishes a virtuous cycle that progressively refines the entire
 449 feature hierarchy. Notably, the layer of peak performance often shifts to align with the chosen refer-
 450 ence layer. For instance, when layer 8 was synced with layer 16, the new performance peak emerged
 451 at layer 16. One possible interpretation is that the alignment process redefines the model’s internal
 452 structure, effectively positioning the reference layer as the new frontier between feature encoding
 453 and decoding. We leave a deeper investigation and further evaluation of LayerSync’s impact on
 454 representations as future work.

455 4.5 LAYERSYNC COMBINED WITH EXTERNAL REPRESENTATIONS

456 In this experiment, we investigate whether LayerSync can be combined with external representation
 457 guidance to further accelerate training. Because LayerSync introduces negligible computational
 458 overhead, it serves as a versatile enhancement: it can be used as a standalone method when external
 459 models are unavailable or too costly, or combined with external representations to maximize per-
 460 formance. We find that the two approaches are synergistic; combining LayerSync with REPA (Yu
 461 et al., 2024) yields better performance than REPA alone. This suggests that the internal structural
 462 alignment of LayerSync (see Section 3.2) and the external semantic injection of REPA operate as
 463 complementary axes of improvement. The results in Table 5 are reported after 200k steps using 4
 464 nodes and a global batch size of 1024. Additional details provided in Section H

465 5 ABLATION STUDY

466 **Layer Selection.** To empirically validate the robustness of our layer selection strategy, we con-
 467 ducted an experiment with randomized layer pairings. For both SiT-XL and SiT-B architectures,
 468 we performed 10 independent training each with a different, randomly selected pair of layers fol-
 469 lowing our simple heuristic proposed in Section 3.4. The results in Table 6 demonstrate remarkable
 470 consistency. The low standard deviation in the FID (0.8 for SiT-XL) on both architecture confirms
 471 that the specific choice of layers is not a very sensitive hyperparameter. This robustness validates
 472 our claim that LayerSync is a practical, plug-and-play method that provides significant performance
 473 gains without necessitating an expensive search for optimal layer combinations.

474 **Effect of λ .** We examine the effect of the regularization coefficient λ on SiT B/2 in Table 7 and
 475 observe that our method is robust to a wide range of values for λ and consistently improves FID.

481 6 DISCUSSION

482 LayerSync is a regularization framework that promotes feature consistency across a model’s depth.
 483 It aligns intermediate layers by encouraging those with weaker representations to become more
 484 similar but not identical to those with richer features. This self-alignment propagates strong semantic
 485 information, which we found accelerates training and improves generative performance.

486
 487 Table 6: Performance of LayerSync with randomized layer pairings. Results show the mean FID
 488 and standard deviation (in parentheses) over 10 independent runs, confirming the robustness of Lay-
 489 erSync to layer selection.

Method	Model	Iterations	FID \downarrow (STD)
Baseline	SiT-B	400k	36.19
Dispersive	SiT-B	400k	32.45
LayerSync - Ours	SiT-B	400k	31.38 (0.7)
Baseline	SiT-XL	400k	17.98
Dispersive	SiT-XL	400k	15.59
LayerSync - Ours	SiT-XL	400k	12.24 (0.8)

490
 491
 492
 493
 494 Table 7: Ablation study for λ . We train SiT B/2 for 400K iterations while aligning block 2 with 8.
 495 We observe that our approach is robust for a wide range of λ . The baseline SiT B/2 has FID 36.19.
 496
 497

λ	0	0.1	0.2	0.3	0.5	0.7	Average (Std)
FID \downarrow	36.19	31.63	31.02	31.6	31.17	31.36	31.356 (0.27)
IS \uparrow	-	44.9	46.12	44.56	45.65	45.2	45.286 (0.61)

501
 502 This similarity between layers raises a natural question: does LayerSync make layers redundant,
 503 potentially allowing for model pruning? Our experiments (Appendix B) show that while models
 504 trained with LayerSync are more robust to layer removal than their baseline counterparts, per-
 505 formance still degrades significantly. This indicates that despite the improved alignment, each layer
 506 retains a unique function essential to the model’s capacity. Consequently, naively pruning a trained
 507 model did not prove superior to simply training a smaller architecture from scratch. [Similar findings](#)
 508 [have been reported for LLMs: intermediate transformer blocks often exhibit high correlation](#). While
 509 [removing these blocks has little effect on easier tasks such as question answering \(QA\), it degrades](#)
 510 [performance on more challenging tasks \(Gromov et al., 2024\)](#) Therefore, [high inter-block corre-](#)
 511 [lation does not necessarily imply that the blocks are redundant](#). However, the increase in resilience
 512 to layer removal is a finding that suggests that LayerSync may alter the functional contribution of
 513 layers in a way that could require further investigation.

514 We also wish to emphasize that the long-term effects of regularization might also demand further
 515 study. Although we did not observe the performance degradation seen in other methods with external
 516 guidance as reported in (Wang et al., 2025), future work could explore scheduling the LayerSync
 517 loss to preemptively address any potential long-term downsides.

518 Finally, the alignment loss function itself presents a key area for future research. We selected cosine
 519 similarity due to its strong empirical performance on images and its effective transfer to audio.
 520 However, developing novel alignment losses specifically engineered for different data domains, such
 521 as the hierarchical nature of text or the temporal patterns in time-series data, is an interesting and
 522 potentially impactful research direction.

523 7 CONCLUSION

524 In this paper, we introduced LayerSync, a simple yet novel self-supervised regularization method
 525 for improving diffusion transformers. We demonstrated that a model’s later-layer representations
 526 can effectively guide its earlier layers, enhancing feature quality and accelerating training at no
 527 additional cost. As a general framework, LayerSync requires no external guidance and is readily
 528 applicable to different data domains.

529 This work opens several avenues for future research in training efficiency, representation learning,
 530 and self-supervised learning. We believe the core principle of LayerSync is broadly applicable and
 531 encourage exploring its potential in other generative architectures beyond diffusion models.

540 8 REPRODUCIBILITY STATEMENT
541542 We are committed to open-sourcing the full codebase and all experiment configurations used in this
543 paper upon acceptance, to support transparency and facilitate future research in this area.
544545 9 ETHICS STATEMENT
546547 This work investigates a self-supervised method for accelerating the training of diffusion models
548 across multiple modalities. All experiments were conducted using publicly available datasets. No
549 personal or sensitive data was used, and we do not anticipate any direct ethical risks associated with
550 this research.
551552 10 ACKNOWLEDGMENT
553554 This research was supported by the Swiss National Science Foundation (SNSF) under Grant No.
555 10003100, and by Innosuisse – the Swiss Innovation Agency (Ref. 127.265 IP-ICT). Computational
556 resources were provided as a part of the Swiss AI Initiative by a grant from the Swiss National
557 Supercomputing Centre (CSCS) under project ID a144 on Alps.
558

559 560 REFERENCES

561 Jie An, De Wang, Pengsheng Guo, Jiebo Luo, and Alexander Schwing. On inductive biases that
562 enable generalization of diffusion transformers. *arXiv preprint arXiv:2410.21273*, 2024.
563564 Fan Bao, Shen Nie, Kaiwen Xue, Yue Cao, Chongxuan Li, Hang Su, and Jun Zhu. All are worth
565 words: A vit backbone for diffusion models. In *Proceedings of the IEEE/CVF conference on*
566 *computer vision and pattern recognition*, pp. 22669–22679, 2023.567 Dmitry Bogdanov, Minz Won, Philip Tovstogan, Alastair Porter, and Xavier Serra. The mtg-jamendo
568 dataset for automatic music tagging. In *Machine Learning for Music Discovery Workshop, Interna-*
569 *tional Conference on Machine Learning (ICML 2019)*, Long Beach, CA, United States, 2019.
570 URL <http://hdl.handle.net/10230/42015>.571 Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics
572 dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp.
573 6299–6308, 2017.
574575 Dengsheng Chen, Jie Hu, Xiaoming Wei, and Enhua Wu. Denoising with a joint-embedding predic-
576 tive architecture. *arXiv preprint arXiv:2410.03755*, 2024.577 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hi-
578 erarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,
579 pp. 248–255. Ieee, 2009.
580581 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances*
582 *in neural information processing systems*, 34:8780–8794, 2021.583 Zach Evans, Julian D Parker, CJ Carr, Zack Zukowski, Josiah Taylor, and Jordi Pons. Stable audio
584 open. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal*
585 *Processing (ICASSP)*, pp. 1–5. IEEE, 2025.
586587 Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman.
588 The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):
589 303–338, 2010.590 Shanghua Gao, Pan Zhou, Ming-Ming Cheng, and Shuicheng Yan. Mdtv2: Masked diffusion trans-
591 former is a strong image synthesizer. *arXiv preprint arXiv:2303.14389*, 2023.
592593 Phani Ghadiyaram. Revelio: interpreting and leveraging semantic information in diffusion models,
2025.

594 Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne West-
 595 phal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al.
 596 The “something something” video database for learning and evaluating visual common sense. In
 597 *Proceedings of the IEEE international conference on computer vision*, pp. 5842–5850, 2017.

598 Andrey Gromov, Kushal Tirumala, Hassan Shapourian, Paolo Glorioso, and Daniel A Roberts. The
 599 unreasonable ineffectiveness of the deeper layers. *arXiv preprint arXiv:2403.17887*, 2024.

600 Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and
 601 Li Cheng. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the*
 602 *28th ACM international conference on multimedia*, pp. 2021–2029, 2020.

603 Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating
 604 diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on*
 605 *Computer Vision and Pattern Recognition (CVPR)*, pp. 5152–5161, June 2022a.

606 Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating
 607 diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF conference on*
 608 *computer vision and pattern recognition*, pp. 5152–5161, 2022b.

609 Ali Hatamizadeh, Jiaming Song, Guilin Liu, Jan Kautz, and Arash Vahdat. Diffit: Diffusion vision
 610 transformers for image generation. In *European Conference on Computer Vision*, pp. 37–55.
 611 Springer, 2024.

612 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter.
 613 Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in*
 614 *neural information processing systems*, 30, 2017.

615 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on*
 616 *Deep Generative Models and Downstream Applications*, 2021.

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

619 Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans.
 620 Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning*
 621 *Research*, 23(47):1–33, 2022.

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

625 Dengyang Jiang, Mengmeng Wang, Liuzhuozheng Li, Lei Zhang, Haoyu Wang, Wei Wei, Guang
 626 Dai, Yanning Zhang, and Jingdong Wang. No other representation component is needed:
 627 Diffusion transformers can provide representation guidance by themselves. *arXiv preprint*
 628 *arXiv:2505.02831*, 2025.

629 Kevin Kilgour, Mauricio Zuluaga, Dominik Roblek, and Matthew Sharifi. Fréchet audio distance: A
 630 reference-free metric for evaluating music enhancement algorithms. In *Interspeech*, 2019. URL
 631 <https://api.semanticscholar.org/CorpusID:202725406>.

632 Diederik Kingma and Ruiqi Gao. Understanding diffusion objectives as the elbo with simple data
 633 augmentation. *Advances in Neural Information Processing Systems*, 36:65484–65516, 2023.

634 Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural
 635 network representations revisited. In *International conference on machine learning*, pp. 3519–
 636 3529. PMIR, 2019.

637 Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathe-
 638 matical statistics*, 22(1):79–86, 1951.

639 Tuomas Kynkänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved
 640 precision and recall metric for assessing generative models. *Advances in neural information*
 641 *processing systems*, 32, 2019.

648 Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444,
 649 2015.

650

651 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow
 652 matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.

653

654 Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl:
 655 A skinned multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pp. 851–866. Association for Computing Machinery, 2015.

656

657

658 Nanye Ma, Mark Goldstein, Michael S Albergo, Nicholas M Boffi, Eric Vanden-Eijnden, and Saining
 659 Xie. Sit: Exploring flow and diffusion-based generative models with scalable interpolant
 660 transformers. In *European Conference on Computer Vision*, pp. 23–40. Springer, 2024.

661

662 Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black.
 663 Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 5442–5451, 2019.

664

665 Soumik Mukhopadhyay, Matthew Gwilliam, Yosuke Yamaguchi, Vatsal Agarwal, Namitha Padmanabhan,
 666 Archana Swaminathan, Tianyi Zhou, Jun Ohya, and Abhinav Shrivastava. Do text-free diffusion
 667 models learn discriminative visual representations? In *European Conference on Computer Vision*, pp. 253–272. Springer, 2024.

668

669 Charlie Nash, Jacob Menick, Sander Dieleman, and Peter W Battaglia. Generating images with
 670 sparse representations. *arXiv preprint arXiv:2103.03841*, 2021.

671

672 Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.

673

674 Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov,
 675 Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning
 676 robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.

677

678 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.

679

680 Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy.
 681 Do vision transformers see like convolutional neural networks? In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, 2021.

682

683

684 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 685 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.

686

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

689

690

691 Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen.
 692 Improved techniques for training gans. *Advances in neural information processing systems*, 29,
 693 2016.

694

695 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.

696

697 Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution.
 698 *Advances in neural information processing systems*, 32, 2019.

699

700 Nick Stracke, Stefan Andreas Baumann, Kolja Bauer, Frank Fundel, and Björn Ommer. Cleandift:
 701 Diffusion features without noise. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 117–127, 2025.

702 Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethink-
 703 ing the inception architecture for computer vision. In *Proceedings of the IEEE conference on*
 704 *computer vision and pattern recognition*, pp. 2818–2826, 2016.

705

706 Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged con-
 707 sistency targets improve semi-supervised deep learning results. *Advances in neural information*
 708 *processing systems*, 30, 2017.

709 Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano.
 710 Human motion diffusion model. In *The Eleventh International Conference on Learning Repre-
 711 sentations*, 2023.

712

713 Keyu Tian, Yi Jiang, Zehuan Yuan, BINGYUE PENG, and Liwei Wang. Visual autoregressive
 714 modeling: Scalable image generation via next-scale prediction. In *The Thirty-eighth Annual*
 715 *Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=gojL67Cfs8>.

716

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

720

721 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
 722 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural informa-
 723 tion processing systems*, 30, 2017.

724

725 Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and
 726 composing robust features with denoising autoencoders. In *Proceedings of the 25th international
 727 conference on Machine learning*, pp. 1096–1103, 2008.

728

729 Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu,
 730 Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative
 731 models. *arXiv preprint arXiv:2503.20314*, 2025.

732

733 Chenyu Wang, Cai Zhou, Sharut Gupta, Zongyu Lin, Stefanie Jegelka, Stephen Bates, and Tommi
 734 Jaakkola. Learning diffusion models with flexible representation guidance. *arXiv preprint*
 735 *arXiv:2507.08980*, 2025.

736

737 Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu,
 738 Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the
 739 world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024.

740

741 Runqian Wang and Kaiming He. Diffuse and disperse: Image generation with representation regu-
 742 larization. *arXiv preprint arXiv:2506.09027*, 2025.

743

744 Weilai Xiang, Hongyu Yang, Di Huang, and Yunhong Wang. Denoising diffusion autoencoders are
 745 unified self-supervised learners. In *Proceedings of the IEEE/CVF International Conference on*
 746 *Computer Vision*, pp. 15802–15812, 2023.

747

748 Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang,
 749 Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models
 750 with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024.

751

752 Kexin Yi, Chuang Gan, Yunzhu Li, Pushmeet Kohli, Jiajun Wu, Antonio Torralba, and Joshua B
 753 Tenenbaum. Clevrer: Collision events for video representation and reasoning. *arXiv preprint*
 754 *arXiv:1910.01442*, 2019.

755

756 Sihyun Yu, Sangkyung Kwak, Huiwon Jang, Jongheon Jeong, Jonathan Huang, Jinwoo Shin, and
 757 Saining Xie. Representation alignment for generation: Training diffusion transformers is easier
 758 than you think. *arXiv preprint arXiv:2410.06940*, 2024.

759

760 Xiangdong Zhang, Jiaqi Liao, Shaofeng Zhang, Fanqing Meng, Xiangpeng Wan, Junchi Yan, and
 761 Yu Cheng. Videorepa: Learning physics for video generation through relational alignment with
 762 foundation models. *arXiv preprint arXiv:2505.23656*, 2025.

756 Tianqi Zhao, Ming Kong, Tian Liang, Qiang Zhu, Kun Kuang, and Fei Wu. Clap: Contrastive
757 language-audio pre-training model for multi-modal sentiment analysis. In Ioannis Kompat-
758 siaris, Jiebo Luo, Nicu Sebe, Angela Yao, Vasileios Mazaris, Symeon Papadopoulos, Adrian
759 Popescu, and Zi Helen Huang (eds.), *ICMR*, pp. 622–626. ACM, 2023. URL <http://dblp.uni-trier.de/db/conf/mir/icmr2023.html#ZhaoKLZK023>.

760
761 Hongkai Zheng, Weili Nie, Arash Vahdat, and Anima Anandkumar. Fast training of diffusion models
762 with masked transformers. *arXiv preprint arXiv:2306.09305*, 2023.

763
764 Rui Zhu, Yingwei Pan, Yehao Li, Ting Yao, Zhenglong Sun, Tao Mei, and Chang Wen Chen. Sd-dit:
765 Unleashing the power of self-supervised discrimination in diffusion transformer. In *Proceedings
766 of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8435–8445, 2024.

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810 A VIDEO GENERATION
811812 We study the effectiveness of LayerSync for both video generation and fine-tuning existing video
813 diffusion models on new datasets.
814815 **Training from Scratch on CLEVRER.** We train SiT-XL model with 3D patchification of size
816 (1,2,2) on the CLEVRER dataset (Yi et al., 2019). Due to the high computational cost of video
817 training from scratch, we limit the run to 24k steps on 16 GPUs, using this as a proof of concept to
818 demonstrate the effectiveness of LayerSync.
819820 **Fine-tuning on SSv2.** We use SSv2 dataset (Goyal et al., 2017) to finetune CogVidX-2B video
821 generation model (Yang et al., 2024) and Wan2.1 1.3B foundation model (Wan et al., 2025). Each
822 video has 33 frames. The frames are normalized and processed with Stable Diffusion VAE (Rom-
823 bach et al., 2022). Instead of cosine similarity we use the TRD similarity metric proposed in Zhang
824 et al. (2025), as it has been shown that it is more effective than cosine similarity for finetuning. We
825 align layer 4 with layer 24. Each model is fine-tuned for a single epoch (1,100 steps) using 16 GPUs
826 and a global batch size of 160 videos.
827828 **TRD loss definition:** The loss is based on (Zhang et al., 2025) and has two terms. The spatial term
829 is a per-frame similarity metric, and the temporal term focuses on the consistency between the frame
830 and is a cross-frame similarity metric. Assuming \mathbf{y}_v being the intermediate representation, we first
831 reshape it to $\mathbb{R}^{f \times (hw) \times D}$, with f being the number of frames, hw being the size of the token and D
832 being the representation dimension, then the spatial term is calculated as:
833

834
$$y_{\text{spatial}}^{d,i,j} = \frac{\mathbf{y}_v^{d,i} \cdot \mathbf{y}_v^{d,j}}{\|\mathbf{y}_v^{d,i}\| \|\mathbf{y}_v^{d,j}\|}, \quad (4)$$

835

836 where $i, j \in [1, hw]$ index spatial positions and d being the frame index.
837838 And the temporal term is calculated as:
839

840
$$y_{\text{temp}}^{d,i,j,e} = \frac{\mathbf{y}_v^{d,i} \cdot \mathbf{y}_v^{e,j}}{\|\mathbf{y}_v^{d,i}\| \|\mathbf{y}_v^{e,j}\|}, \quad \forall e \in [1, f] \setminus \{d\}, j \in [1, hw]. \quad (5)$$

841

842 The final TRD loss term is :
843

844
$$\mathcal{L}_{\text{TRD}} = \underbrace{\frac{1}{f(hw)^2} \sum_{d=1}^f \sum_{i,j=1}^{hw} |h_{\text{spatial}}^{d,i,j} - y_{\text{spatial}}^{d,i,j}|}_{\text{Spatial component}} + \underbrace{\frac{1}{f(hw)^2(f-1)} \sum_{d=1}^f \sum_{e \neq d} \sum_{i,j=1}^{hw} |h_{\text{temp}}^{d,i,j,e} - y_{\text{temp}}^{d,i,j,e}|}_{\text{Temporal component}}, \quad (6)$$

845

846 where h can be either an external representation or, in our case the representation of a different layer.
847848 **Baselines.** We compare LayerSync with Dispersive (Wang & He, 2025). We apply the dispersive
849 loss at 25% depth which is what yielded the best results in original work. We refer to fine-tuning
850 without any extra guidance as vanilla.
851852 **Evaluation metrics.** We rely on Fréchet Video Distance (FVD; Unterthiner et al. (2018)) for eval-
853 uation. We generate 5000 videos of 33 frames for evaluation on finetuned models and 16 frames for
854 SiT-XL.
855856 **Results.** As shown in Table 8, LayerSync consistently outperforms all baselines across both fine-
857 tuning and from-scratch training setups, achieving the lowest FVD scores in every scenario. When
858 fine-tuning large pre-trained models on SSv2, LayerSync improves FVD by 19.1% over the vanilla
859 baseline for CogVideoX-2B and by 22.8% for Wan2.1, demonstrating its effectiveness in enhancing
860 temporal coherence and sample quality during adaptation. In the from-scratch CLEVRER experi-
861 ment, LayerSync achieves a 54.7% reduction in FVD compared to the vanilla baseline. This sub-
862 stantial gain highlights LayerSync’s ability to serve as a strong inductive bias, improving learning
863 efficiency and generation quality.
864865 These consistent improvements across model scales and training regimes underscore the generality
866 and robustness of LayerSync as a self-contained regularization strategy.
867

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

Table 8: **FVD scores (\downarrow) for video generation.** We observe that LayerSync consistently improves FVD for both finetuning and training from scratch.

	CogVideoX-2B (SSv2)	Wan2.1 (SSv2)	SiT-XL (CLEVERER)
Vanilla	371.88	363.98	265.50
Dispersive	342.10	372.43	165.12
LayerSync (Ours)	300.91	280.78	120.13

B DROPPING BLOCKS

To investigate the effect of dropping blocks, we train SiT XL/2 and SiT XL/2 with LayerSync (aligning layers 7 and 16) for 120k iterations. We then drop four blocks in between the aligned layers. Quantitative results are presented in Table 9, and qualitative examples are shown in Figure 5, indicating that the model trained with LayerSync is more robust to block drop.

We also experimented dropping blocks outside the aligned layers. As summarized in Table 9 and Figure 6, this leads to a more significant degradation in sample quality, suggesting that the drop of blocks outside the synced range has a more detrimental effect. Although LayerSync improves robustness to dropped blocks, doing so still results in an increase in FID.

Table 9: Comparison of FID, sFID, Inception Score (IS), Precision, and Recall when dropping specific blocks from the model. The model trained with LayerSync is more robust to block removal.

	Skipped blocks	FID \downarrow	sFID \downarrow	IS \uparrow	Precision \uparrow	Recall \uparrow
SiT XL/2	-	37.03	5.49	35.41	0.53	0.61
SiT XL/2 + Layer Sync	-	25.72	5.05	48.49	0.61	0.59
SiT XL/2	[9,11,13,15]	211.66	93.92	4.02	0.01	0.10
SiT XL/2 + Layer Sync	[9,11,13,15]	55.07	7.85	23.04	0.39	0.63
SiT XL/2 + Layer Sync	[9,11,13,15,21]	86.11	18.30	16.79	0.29	0.46
SiT XL/2 + Layer Sync	[1,9,11,13,15]	92.84	22.28	15.38	0.26	0.44



Figure 5: Qualitative comparison of generated samples from SiT XL/2 and SiT XL/2 with LayerSync when layers 7 and 16 are synced. After dropping blocks [9, 11, 13, 15], we observe that LayerSync helps preserve visual quality despite block removal.

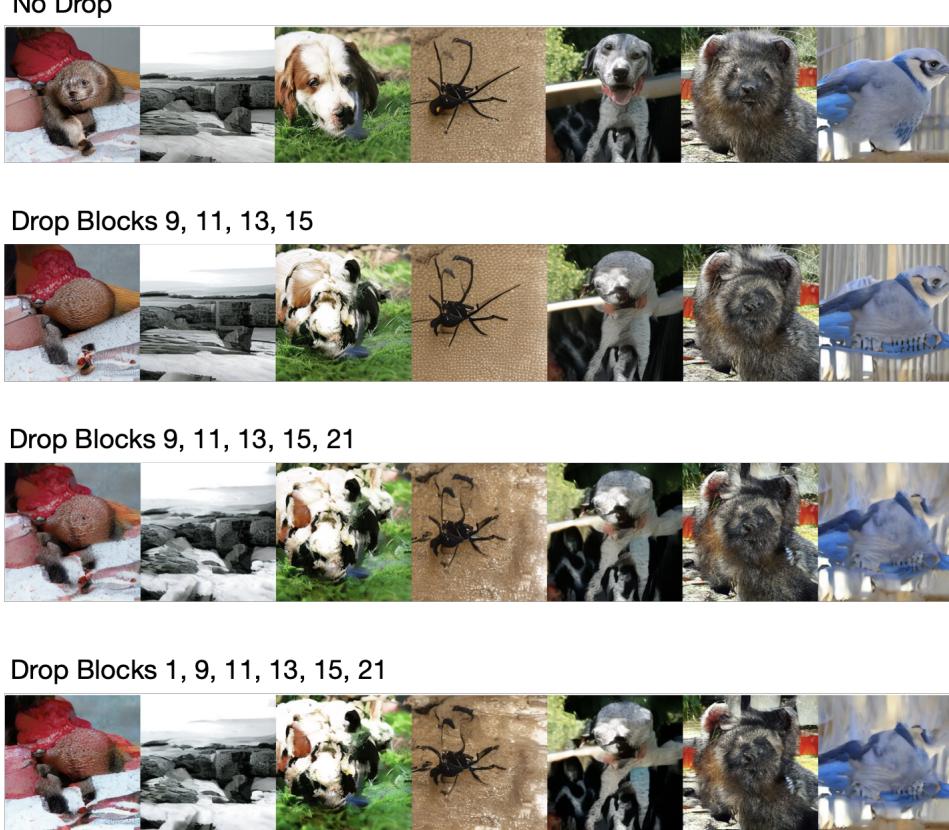


Figure 6: Qualitative results when dropping blocks from SiT XL/2 with LayerSync, where layers 7 and 16 are synced. Dropping blocks outside the synced layers leads to a more noticeable degradation in sample quality.

C ALIGNMENT CONSIDERING THE Timestep

We apply alignment on SiTXL/2 for the last 75 %, 50 % and 25 % of timesteps. The model is trained on ImageNet 256 × 256 for 80 epochs. The results summarized in Table 10 shows that the best performance is achieved when alignment is applied on all the timesteps, which confirms that improving the weak representations is beneficial regardless of the timestep.

Table 10: Applying LayerSync on specific timesteps shows that alignment is beneficial for all the timesteps.

Timestep	FID ↓
25%	18.28
50%	18.77
75%	17.68
100%	16.03

D ABLATION ON BLOCK SELECTION

Ablation study on block selection for SiT-XL/2 and SiT-L/2 summarized in Tables 11 and 12 shows that LayerSync consistently improves the generation quality, but the gain is suboptimal when the distance between the blocks is low or when aligning with the decoder blocks (very last layers). All the models are trained for 100K iterations with 16 GPUs and a batch size of 1024.

972
 973 Table 11: **Ablation study on block selection for SiT-XL/2.** The results show that LayerSync
 974 consistently improves the generation quality, but the gain is suboptimal when the distance between
 975 the blocks is low or when aligning with the decoder blocks (very last layers).

Model	IS \uparrow	FID \downarrow	sFID \downarrow
Vanilla SiT-XL/2	50.408	26.534	5.035
LayerSync 6 \leftarrow 18	75.221	15.386	4.672
LayerSync 7 \leftarrow 17	75.727	15.433	4.694
LayerSync 7 \leftarrow 15	74.987	15.740	4.611
LayerSync 9 \leftarrow 17	73.791	16.058	4.695
LayerSync 8 \leftarrow 18	73.549	16.078	4.589
LayerSync 8 \leftarrow 17	72.821	16.276	4.638
LayerSync 8 \leftarrow 20	71.477	16.568	4.662
LayerSync 9 \leftarrow 20	71.016	16.661	4.695
LayerSync 8 \leftarrow 19	71.094	16.680	4.658
LayerSync 6 \leftarrow 14	72.750	16.697	4.695
LayerSync 10 \leftarrow 18	70.209	16.999	4.638
LayerSync 9 \leftarrow 19	69.831	17.091	4.701
LayerSync 11 \leftarrow 21	68.952	17.096	4.673
LayerSync 7 \leftarrow 21	69.438	17.126	4.685
LayerSync 11 \leftarrow 19	68.464	17.615	4.628
LayerSync 9 \leftarrow 22	68.127	17.823	4.717
LayerSync 8 \leftarrow 21	67.633	18.032	4.705
LayerSync 6 \leftarrow 8	57.956	22.630	4.903
LayerSync 6 \leftarrow 9	56.552	23.307	4.892
LayerSync 6 \leftarrow 7	55.717	23.924	4.954
LayerSync 11 \leftarrow 23	61.851	19.966	4.769
LayerSync 12 \leftarrow 23	59.692	20.711	4.737
LayerSync 15 \leftarrow 23	55.235	23.030	4.783

1000
 1001 Table 12: **Ablation study on block selection for SiT-L/2.** The results show that LayerSync
 1002 consistently improves the generation quality, but the gain is suboptimal when the distance between
 1003 the blocks is low or when aligning with the decoder blocks (very last layers).

Model	IS	FID	sFID
LayerSync 6 \leftarrow 15	64.638	19.165	4.827
LayerSync 8 \leftarrow 16	62.794	19.566	4.856
LayerSync 7 \leftarrow 15	62.854	19.663	4.841
LayerSync 6 \leftarrow 14	62.616	19.945	4.791
LayerSync 5 \leftarrow 13	62.158	20.048	4.778
LayerSync 7 \leftarrow 16	61.4	20.055	4.865
LayerSync 7 \leftarrow 18	60.533	20.324	4.820
LayerSync 6 \leftarrow 16	61.304	20.444	4.794
LayerSync 5 \leftarrow 12	60.807	20.492	4.775
LayerSync 5 \leftarrow 11	61.104	20.696	4.821
LayerSync 5 \leftarrow 15	60.362	20.734	4.824
LayerSync 5 \leftarrow 17	61.178	20.771	4.878
LayerSync 5 \leftarrow 8	50.473	26.511	5.084
LayerSync 5 \leftarrow 7	49.948	27.004	5.171
LayerSync 5 \leftarrow 6	47.19	28.741	5.21
LayerSync 11 \leftarrow 20	43.253	30.57	5.172
LayerSync 17 \leftarrow 20	43.563	31.056	5.308
LayerSync 16 \leftarrow 20	43.242	31.24	5.316
LayerSync 13 \leftarrow 20	42.535	31.407	5.182
LayerSync 14 \leftarrow 20	41.626	32.35	5.282
LayerSync 15 \leftarrow 20	41.381	32.622	5.403

1026 **E LAYERSYNC VS. INCREASING LEARNING RATE**
1027

1028 To show that LayerSync impact is not simply due to an increase in the gradient magnitude, we
1029 designed two sets of experiments below. All the models are trained with 16 GPUs, batch size
1030 1024 for 100k iterations (80 epochs). We then report the gradient norms and the FIDs of different
1031 configurations. We consistently observed that the gradient norm with LayerSync is actually smaller
1032 than that of the baseline, indicating that our method is not simply equivalent to using a larger learning
1033 rate.

1034 **(1) Global learning rate increase.**
1035

1036 Starting from the default learning rate of 1×10^{-4} , we trained models with higher learning rates
1037 2×10^{-4} and 5×10^{-4} . For learning rates above 5×10^{-4} , the model diverges. While increasing
1038 the learning rate can partially accelerate training, the resulting FID scores remain worse than
1039 those obtained with LayerSync with the default learning rate 1×10^{-4} as shown in Table 13. The
1040 visualization of gradient norm is provided in Figure 7a.

1041 Table 13: **Effect of Global Learning Rate.** Simply increasing the global learning rate improves
1042 FID slightly but does not match the performance gains of LayerSync.

Method	lr.	FID
SiT-XL/2	1×10^{-4}	26.53
SiT-XL/2	2×10^{-4}	24.95
LayerSync	1×10^{-4}	16.03

1043 **(2) Higher learning rate on early blocks only.**
1044

1050 We then increased the learning rate only for the first 8 blocks and compared this to a model trained
1051 with LayerSync aligning layers 8–16 at a global learning rate of 1×10^{-4} . Again, for learning
1052 rates above 1×10^{-3} training diverges. As summarized in Table 14, for 2×10^{-4} , 5×10^{-4} , and
1053 1×10^{-3} , the FID improvements do not match those obtained with LayerSync. The visualization of
1054 the gradient norms is provided in Figure 7b.

1055 Table 14: **Effect of Early-Layer Learning Rate.** Increasing the learning rate specifically on early
1056 layers (first 8 blocks) is beneficial but still underperforms compared to LayerSync.

Method	lr. Early Layers	General lr.	FID
SiT-XL/2	1×10^{-4}	1×10^{-4}	26.53
SiT-XL/2	2×10^{-4}	1×10^{-4}	19.24
SiT-XL/2	5×10^{-4}	1×10^{-4}	24.63
LayerSync	1×10^{-4}	1×10^{-4}	16.03

1064 **F LAYERSYNC VS SELF-REPRESENTATION ALIGNMENT USING EMA**
1065

1067 We compare our LayerSync approach with the concurrent work SRA (Jiang et al., 2025) in terms of
1068 training time and computational overhead. We computed the metrics for SiT-XL/2 using 4 GH200
1069 GPUs and a batch size of 32 per GPU. Results are reported in Table 15. We show that LayerSync
1070 requires 25.5% fewer Flops, is 40.5% faster in real-time, and reaches an FID 5% higher.

1071 Table 15: **Comparison between LayerSync and SRA.** LayerSync results in lower FID while being
1072 less computationally expensive.

Method	FID \downarrow	Wall-clock time/step \downarrow	GFlops \downarrow
SiT-XL/2+SRA	1.58	0.617	30762
SiT-XL/2+LayerSync	1.50	0.367	22910

1079 **G THE VIRTUOUS CYCLE AND THE EVOLUTION OF REPRESENTATIONS**

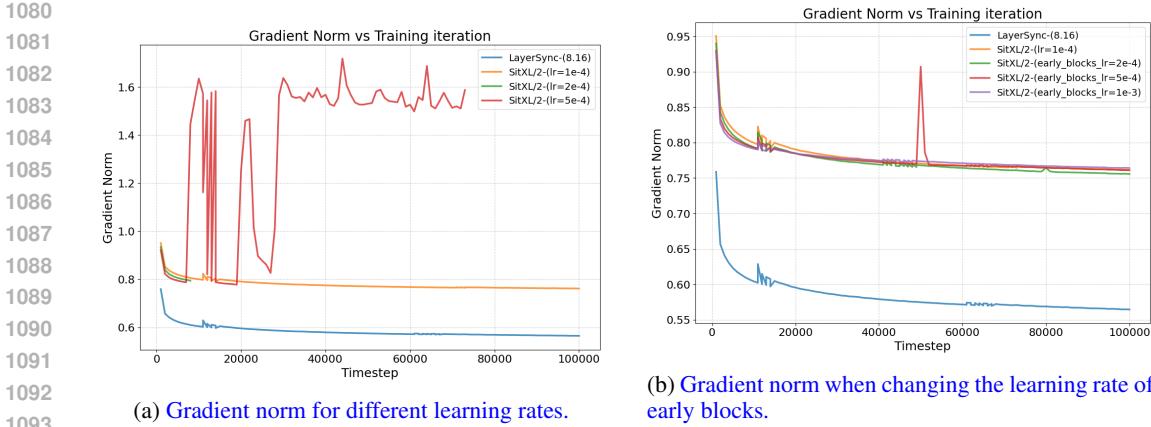


Figure 7: **Gradient norm visualization revealed that LayerSync impact is not similar to simply increasing the learning rate as the model trained with LayerSync has lower gradient norm than the baseline.**

To provide empirical support for the "virtuous cycle" and analyze the evolution of the feature hierarchy, we conducted two additional studies evaluating the segmentation performance (mIoU) of internal representations throughout the network on PASCAL VOC (Everingham et al., 2010).

First, we tracked the evolution of representations throughout the training process (from 100k to 600k steps). As shown in Figure 8a, we observe that while the relative structure of feature quality across layers remains stable, the performance monotonically improves across the entire hierarchy.

Our second study evaluates models trained with different alignment targets at the same training stage (100k steps) (see Figure 8b). This comparison provides the most compelling evidence for our hypothesis. When we synchronize an early block (e.g., block 6) with a deeper target (block 18 vs. block 10), we observe two critical effects:

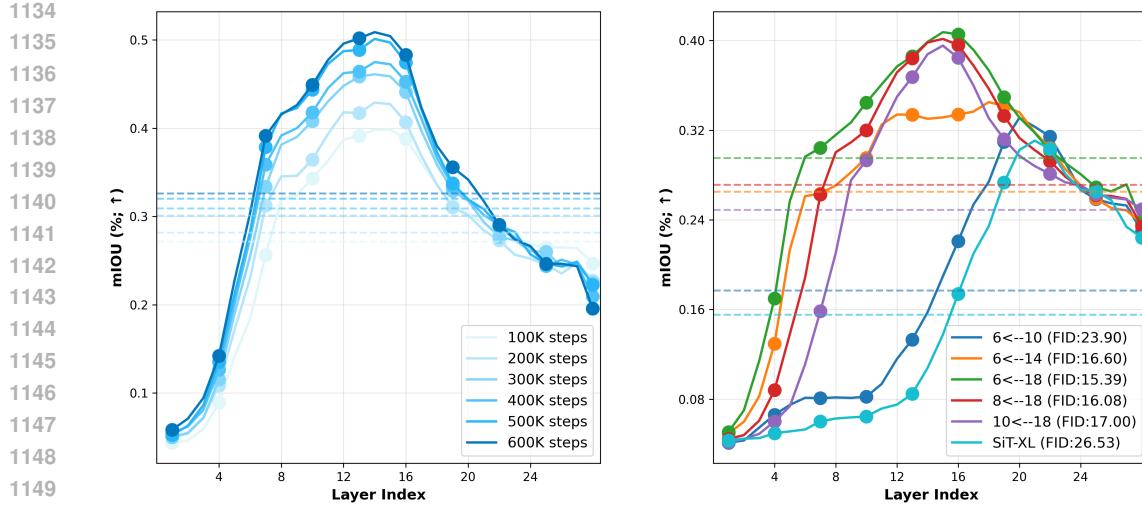
- **Global Improvement:** The model guided by the deeper layer achieves superior downstream performance and lower FID (15.39 vs. 23.90), indicating that guidance from deeper layers correlates with better overall representations.
- **Accelerated Maturation:** Notably, using a deeper target shifts the peak performance of the network to earlier layers. By effectively "pulling" semantic richness from the deep target to the earlier block, the early layers appear to acquire higher-level features sooner in the depth hierarchy.

These observations are consistent with the hypothesized virtuous cycle: guiding an early block with a stronger target improves its representation, which in turn provides higher-quality input to subsequent layers. This likely facilitates the learning of stronger deep representations, which then serve as even better guides, progressively refining the entire hierarchy.

H OPTIMAL PLACEMENT OF EXTERNAL GUIDANCE WHEN COMBINED WITH LAYERSYNC

To maximize the synergy between internal and external alignment, we investigated the optimal depth for applying REPA. The results are summarized in Table 16 and show that applying REPA *before* the synchronization range leads to no significant synergies. The most effective strategy integrates the external signal *between* the aligned layers. For instance, using LayerSync to align layers 8 and 16 while applying REPA at layer 10 yields the best performance. All the models are trained for 50k iterations with 16 GPUs and batch size 1024.

I EVOLUTION OF BLOCK STRUCTURE



(a) Evolution of representations throughout the training process assessed by semantic segmentation on PASCAL VOC dataset.

(b) Evaluation of representations for different layers alignments assessed by semantic segmentation on PASCAL VOC dataset.

Figure 8: **Evaluation of the representations assessed by semantic segmentation on PASCAL VOC dataset.** (a) Evolution throughout the training process. (b) Evaluation of representation for different layer alignment.

Table 16: **Qualitative comparison between different combinations of REPA and LayerSync.** The results show that Combining LayerSync with REPA can further accelerate the training and the best place to apply REPA is between the syncing layers.

Method	REPA Layer	LayerSync Layer	FID ↓
SiTXL/2	–	–	59.45
SiTXL/2 + LayerSync	–	8–16	46.26
SiTXL/2 + REPA	7	–	46.06
SiTXL/2 + REPA + LayerSync	7	8–16	43.55
SiTXL/2 + REPA + LayerSync	10	8–16	29.68

A comparison between the block structure of the SiT-XL/2 and SiT-XL/2 + LayerSync is provided in Figure 9, showing that LayerSync imposes the structural equilibrium early in training.

J LAYERSYNC - ALGORITHM

We present the algorithmic formulation of LayerSync in Algorithm 1.

Algorithm 1 LayerSync

Require: Weak Representation $Z_k \in \mathbb{R}^{B \times P \times D}$, Strong Representation $Z_{k'} \in \mathbb{R}^{B \times P \times D}$ where B is the batch size, P the number of patches and D the feature dimension.

```

1:  $Z_k^{\text{norm}} \leftarrow \text{normalize}(Z_k, \text{dim} = -1)$   $\triangleright$  L2-normalize embeddings
2:  $Z_{k'}^{\text{norm}} \leftarrow \text{normalize}(Z_{k'}, \text{dim} = -1)$   $\triangleright$  L2-normalize embeddings
3:  $\mathcal{L}_{\text{LayerSync}} \leftarrow -\text{similarity}(\sum_{j=1}^P Z_k^{\text{norm}}[:, j] \cdot Z_{k'}^{\text{norm}}[:, j])$   $\triangleright$  Negative similarity across patches
4: return  $\mathcal{L}_{\text{LayerSync}}$ 

```

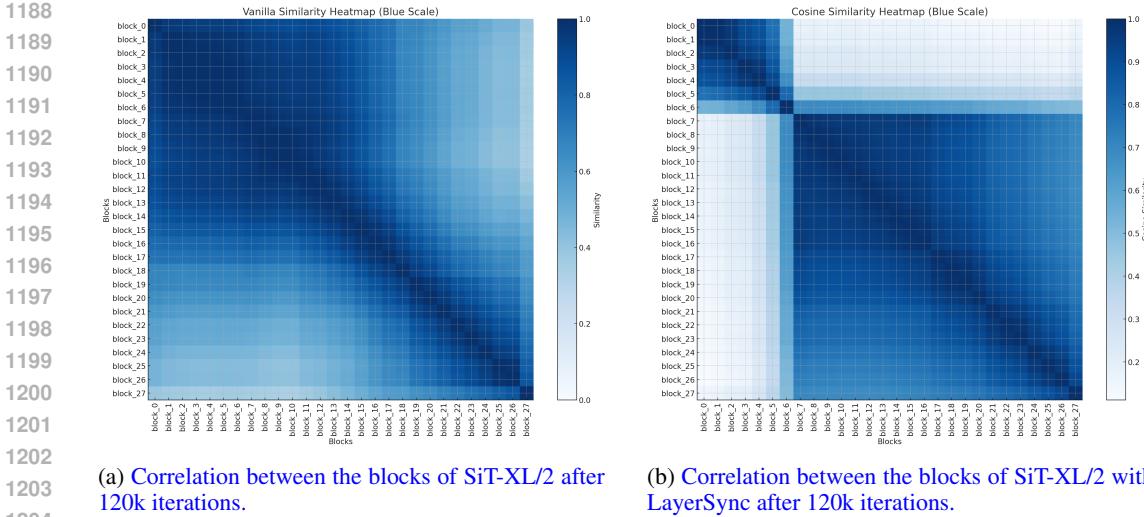


Figure 9: **Evolution of the Block Structure.** (a) Early in training (120k iterations), the vanilla model lacks this structure. (b) LayerSync imposes such structural equilibrium early in training (120k iterations).

K FLOP COMPARISON

We compare the computational complexity of Dispersive Loss, which computes pairwise distances, with LayerSync in Table 17. LayerSync is more efficient in terms of computational complexity as the pairwise comparisons in Dispersive Loss result in a quadratic cost with respect to batch size.

Table 17: Comparison of computational complexity between the Dispersive Loss (pairwise distances) and LayerSync. B is the batch size and D is the feature dimension.

	FLOPs	Scaling w.r.t. Batch Size
Dispersive	$\mathcal{O}(B^2 D)$	Quadratic (B^2)
Layer Sync	$\mathcal{O}(BD)$	Linear (B)

L IMAGE GENERATION EXPERIMENTAL DETAILS

We use a node of 4 GH200 GPUs and a batch size of 256. The details of hyper parameters and sampler are provided in Tables 18 and 19.

Classification. We use the Tiny ImageNet dataset (Deng et al., 2009), upsample the images to 256×256 , and train linear classification heads for 50 epochs. Performance is evaluated on the validation set.

Segmentation. For segmentation, we use the PASCAL VOC dataset (Everingham et al., 2010) and train linear heads for 25 epochs.

CKA. For CKA evaluations (Kornblith et al., 2019), we use 4,000 samples from ImageNet 256×256 .

Table 18: Hyperparameter setup for main experiments.

	Table 1 (SiT-B)	Table 1 (SiT-L)	Table 1 (SiT-XL)	Table 2
Architecture				
Input dim.	$32 \times 32 \times 4$			
Num. layers	12	24	28	28
Hidden dim.	768	1024	1152	1152
Num. heads	12	16	16	16
LayerSync				
λ	0.3	0.2	0.2	0.2
Syncing layers	(4,7)	(8,18)	(8,16)	(8,16)
$\text{sim}(\cdot, \cdot)$	cos. sim.	cos. sim.	cos. sim.	cos. sim.
Optimization				
Batch size	256	256	256	256
Optimizer	AdamW	AdamW	AdamW	AdamW
lr	0.0001	0.0001	0.0001	0.0001
Interpolants				
α_t	t	t	t	t
σ_t	$1 - t$	$1 - t$	$1 - t$	$1 - t$
Training objective	v-prediction	v-prediction	v-prediction	v-prediction
Sampler	ODE Heun	ODE Heun	ODE Heun	SDE Euler–Maruyama
Sampling steps	250	250	250	250
Guidance	–	–	–	1.37

Table 19: Hyperparameter setup for figures and ablation experiments.

	Figure 3 (SiT-XL)	Table 6 (SiT-XL and SiT-B)	Table 7 (SiT-B)
Architecture			
Input dim.	$32 \times 32 \times 4$	$32 \times 32 \times 4$	$32 \times 32 \times 4$
Num. layers	28	–	12
Hidden dim.	1152	–	768
Num. heads	16	–	12
LayerSync			
λ	0.2	0.3	–
Alignment depth	(8,16)	–	(2,8)
$\text{sim}(\cdot, \cdot)$	cos. sim.	cos. sim.	cos. sim.
Optimization			
Training iteration	400K	400K	400K
Batch size	256	256	256
Optimizer	AdamW	AdamW	AdamW
lr	0.0001	0.0001	0.0001
Interpolants			
α_t	t	t	t
σ_t	$1 - t$	$1 - t$	$1 - t$
Training objective	v-prediction	v-prediction	v-prediction
Sampler	ODE Heun	ODE Heun	ODE Heun
Sampling steps	250	250	250
Guidance	–	–	–

1296 M HUMAN MOTION GENERATION EXPERIMENTAL DETAILS
12971298 **Task.** Given a sentence that describes a motion as a sequence of actions, the task is to generate
1299 a corresponding human motion. Each motion sequence consists of a series of human poses, where
1300 each pose is represented by 22 joints defined as 3D points in space.
13011302 **Dataset.** We rely on HumanML3D dataset (Guo et al., 2022a) that contains 44,970 motion anno-
1303 tations across 14,646 motion sequences from the AMASS (Mahmood et al., 2019) and HumanAct12
1304 (Guo et al., 2020) datasets, along with corresponding text descriptions, and is widely used for the
1305 task of text-conditional human motion generation. Motions in the HumanML3D dataset follow the
1306 skeleton structure of SMPL (Loper et al., 2015) with 22 joints. Each pose \mathbf{p} in the motion sequence
1307 is represented by a vector of size 237,
1308

1309
$$(r^a, \dot{r}^x, \dot{r}^z, r^y, j^p, j^v, j^r, c^f),$$

1310 where $r^a \in \mathbb{R}$ is the root (pelvis joint) angular velocity along the Y-axis; $(\dot{r}^x, \dot{r}^z) \in \mathbb{R}$ are the root
1311 linear velocities in the XZ-plane; $r^y \in \mathbb{R}$ is the root height; $j^p \in \mathbb{R}^{3j}$, $j^v \in \mathbb{R}^{3j}$, and $j^r \in \mathbb{R}^{6j}$ are
1312 the local joint positions, velocities, and rotations in the root space, with j indicating the number of
1313 joints; $c^f \in \mathbb{R}^4$ represents foot-ground contact features.
13141315 **Implementation details.** We use the exact setup as MDM (Tevet et al., 2023), we train up to
1316 600K iterations using a H100 GPU. We sync block 3 with block 6.
13171318 N STOCHASTIC INTERPOLANTS
13191320 We adopt the generalized perspective of stochastic interpolants (Ma et al., 2024) which provides a
1321 unifying framework for both flow-based and diffusion-based models.
13221323 At the core of these models is a process that gradually transforms a real data sample $\mathbf{x}_0 \sim p(\mathbf{x})$ into
1324 a simple noise sample $\epsilon \sim \mathcal{N}(0, I)$. This process is defined by:
1325

1326
$$\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \epsilon, \quad (7)$$

1327

1328 where α_t and σ_t are functions of time, respectively decreasing and increasing, that control the mix of
1329 data and noise, satisfying the boundary conditions $\alpha_0 = \sigma_T = 1$, and $\alpha_T = \sigma_0 = 0$. The generative
1330 process aims to reverse this path. This can be model through a deterministic trajectory commonly
1331 described as the probability flow ordinary differential equation (PF-ODE).
1332

1333
$$\dot{\mathbf{x}}_t = v(\mathbf{x}_t, t), \quad (8)$$

1334

1335 where $v(\mathbf{x}_t, t)$ is the velocity field, specifying the direction and magnitude of movement at any point
1336 \mathbf{x}_t at any time t to go from noise back to data. The velocity fields is defined as the time derivative of
1337 the interpolant:
1338

1339
$$v(\mathbf{x}, t) = \dot{\mathbf{x}}_t \Big|_{\mathbf{x}_t = \mathbf{x}} = \dot{\alpha}_t \mathbb{E}[\mathbf{x}_0 \mid \mathbf{x}_t = \mathbf{x}] + \dot{\sigma}_t \mathbb{E}[\epsilon \mid \mathbf{x}_t = \mathbf{x}]. \quad (9)$$

1340

However, since those conditional expectations are intractable, a model $v_\theta(\mathbf{x}_t, t)$ is trained to approximate it by minimizing the flow matching loss defined as:
1341

1343
$$\mathcal{L}_{\text{velocity}}(\theta) := \mathbb{E}_{\mathbf{x}_0, \epsilon, t} \left[\|v_\theta(\mathbf{x}_t, t) - \dot{\alpha}_t \mathbf{x}_0 - \dot{\sigma}_t \epsilon\|^2 \right]. \quad (10)$$

1344

1346 The data is then generated by integrating equation 8 from $t=1$ to $t=0$ using any standard ODE solver
1347 starting from a random noise sample $\mathbf{x}_1 \sim \mathcal{N}(0, \mathbf{I})$. There exists also an alternative way to model the
1348 reverse process using Stochastic Differential Equation (SDE). The SDE shares the same marginal
1349 probability densities $p_t(\mathbf{x})$ as the PF-ODE but follows a stochastic, rather than deterministic, trajectory.
The general form of this reverse SDE is:
1350

1350

$$d\mathbf{x}_t = \left(v(\mathbf{x}_t, t) - \frac{1}{2} w_t s(\mathbf{x}_t, t) \right) dt + \sqrt{w_t} d\mathbf{w}_t \quad (11)$$

1353

1354 where w_t is a diffusion coefficient and $d\mathbf{w}_t$ is a standard Wiener process, and $s(\mathbf{x}_t, t)$ is the score
 1355 function, defined as the gradient of the log-density of the data. The velocity and the score are not
 1356 independent, they are two sides of the same coin as the score can be derived from the velocity field
 1357 and vice versa.

1358

O EVALUATION METRICS DETAILS.

1359

O.1 IMAGE

1360

- **FID.** Heusel et al. (2017) measures the distance between the real and generated data distributions in the feature space of a pretrained Inception-v3 network (Szegedy et al., 2016). It computes the Fréchet distance (Heusel et al., 2017) between two multivariate Gaussians fitted to the feature embeddings, capturing both the quality and diversity of generated samples. Lower values indicate better performance.
- **sFID.** Nash et al. (2021) compares local image patches instead of global image statistics. By focusing on patch-level embeddings, sFID provides a more fine-grained evaluation of spatial consistency and local realism in the generated samples.
- **Inception Score.** Salimans et al. (2016) computes the Kullback–Leibler (KL) divergence (Kullback & Leibler, 1951) between conditional and marginal label distributions predicted by an Inception network.
- **Precision and Recall.** Kynkänniemi et al. (2019) measures the fraction of generated samples that lie within the support of the real data distribution in feature space. High recall reflects the diversity of generated samples, indicating that the model captures the variability of the real data distribution.

1377

O.2 AUDIO

1378

1379

1380

1381

1382

1383

1384

O.3 VIDEO

1385

1386

1387

1388

1389

1390

O.4 MOTION

1391

1392

1393

1394

1395

P EXTENDED RELATED WORK

1396

1397

In what follows, we summarize the main baseline methods used in our evaluation:

1398

1399

1400

1401

1402

1403

- **ADM** (Dhariwal & Nichol, 2021): Builds upon U-Net-based diffusion models by introducing classifier-guided sampling, allowing fine-grained control over the trade-off between generation quality and diversity.
- **VDM++** (Kingma & Gao, 2023): Proposes an adaptive noise schedule that adjusts dynamically during training, improving convergence and sample quality.

- **Simple diffusion** (Hoogeboom et al., 2023): Simplifies both the noise schedule and architectural components, enabling high-resolution image generation with improved computational efficiency.
- **CDM** (Ho et al., 2022): Introduces cascaded diffusion models that progressively refine images from low to high resolution using super-resolution stages, achieving better detail synthesis.
- **LDM** (Rombach et al., 2022): Trains diffusion models in a compressed latent space learned by a VAE, drastically reducing training cost while maintaining image fidelity.
- **U-ViT** (Bao et al., 2023): Combines ViT-based backbones with U-Net-style skip connections in the latent space, bridging the benefits of transformers and convolutional inductive biases.
- **DiffiT** (Hatamizadeh et al., 2024): Enhances transformer-based diffusion models using time-aware multi-head self-attention, boosting sample efficiency and reducing training time.
- **MDTv2** (Gao et al., 2023): Employs an asymmetric encoder-decoder transformer architecture with U-Net-inspired shortcuts in the encoder and dense skip connections in the decoder, improving video generation quality and coherence.
- **MaskDiT** (Zheng et al., 2023): Introduces masked modeling into diffusion transformers by training with an auxiliary mask reconstruction objective, leading to better efficiency and generalization.
- **SD-DiT** (Zhu et al., 2024): Builds on MaskDiT by incorporating a self-supervised discrimination objective using momentum encoding, enhancing the semantic richness of internal representations.
- **DiT** (Peebles & Xie, 2023): Proposes a pure transformer architecture for diffusion, using AdaLN-zero modules to stabilize training and scale to large model sizes efficiently.
- **SiT** (Ma et al., 2024): Investigates the link between training efficiency and flow-based perspectives by transitioning from discrete-time diffusion to continuous flow matching, showing improved sample quality and convergence rates.
- **D-JEPA** (Chen et al., 2024): Integrates Joint-Embedding Predictive Architectures (JEPA) into generative modeling by reframing masked image modeling as a generalized next-token prediction task, utilizing diffusion or flow matching loss to model per-token probability distributions in a continuous space.
- **VAR** (Tian et al., 2024): Redefines autoregressive image generation as a coarse-to-fine “next-scale prediction” process, diverging from standard raster-scan next-token prediction, allowing for faster inference and scaling laws similar to Large Language Models.
- **SRA** (Jiang et al., 2025): an EMA-based method that enables diffusion transformers to enhance their own representation learning and generation quality by aligning latent outputs from earlier, noisier layers with those from later, cleaner layers, eliminating the need for external guidance models.

Q DETAILS OF SiT MODEL

The architecture of the SiT block is provided in Figure 10 and more details on the model parameters are summarized in Table 20.

Table 20: The number of transformer layers, hidden dimensionality, and number of attention heads for SiT models used in our experiments.

Config	#Layers	Hidden dim	#Heads
B/2	12	768	12
L/2	24	1024	16
XL/2	28	1152	16

```

1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488

```

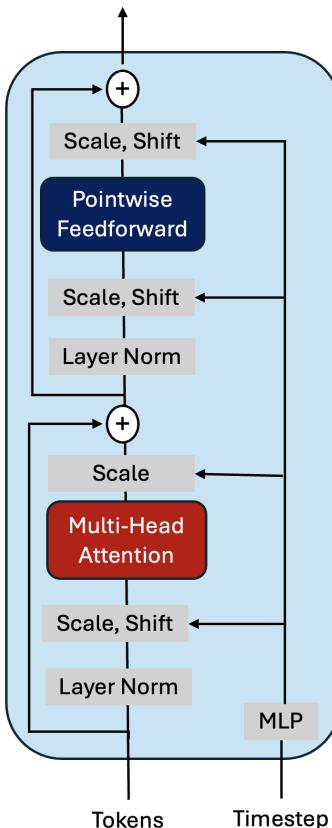


Figure 10: Visualization of a single SiT block.

R ATTENTION MAPS PCA OVER LAYERS

We visualize the learned representations by applying PCA to the features of SiT-XL/2 models trained on ImageNet 256×256 . We add different levels of noise to the input image and visualize the resulting features. We compare two variants: the baseline SiT-XL/2 and SiT-XL/2 with **LayerSync**, where block 8 is synced with block 16. Both models are trained for 400K iterations on a single node with 4 GH100 GPUs. Our results show that LayerSync results in more discriminative features, particularly in the earlier blocks.

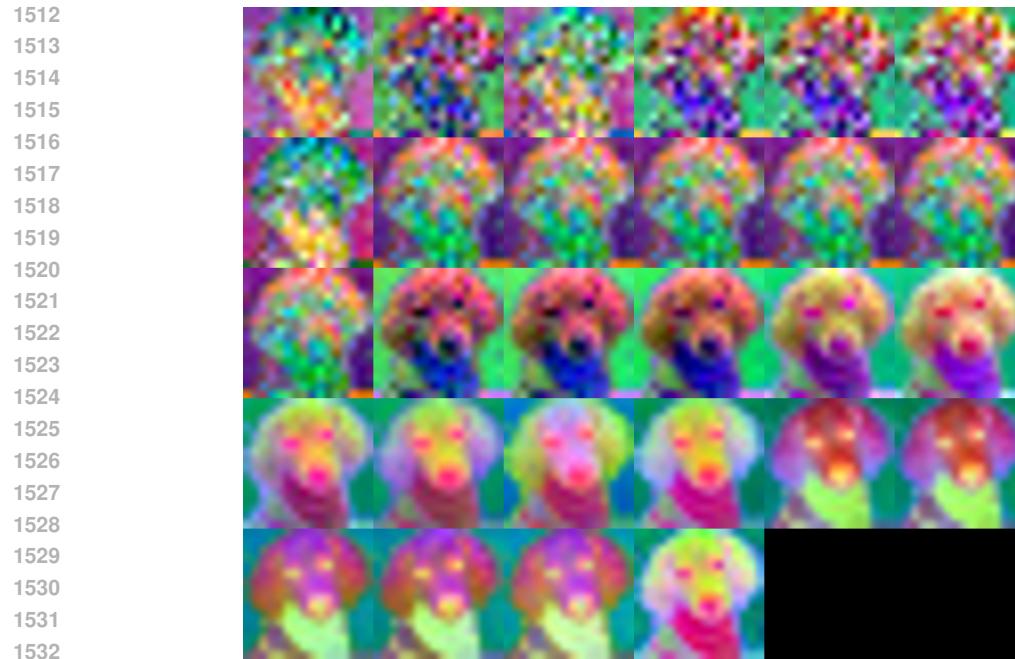


Figure 11: **Input image to the model**

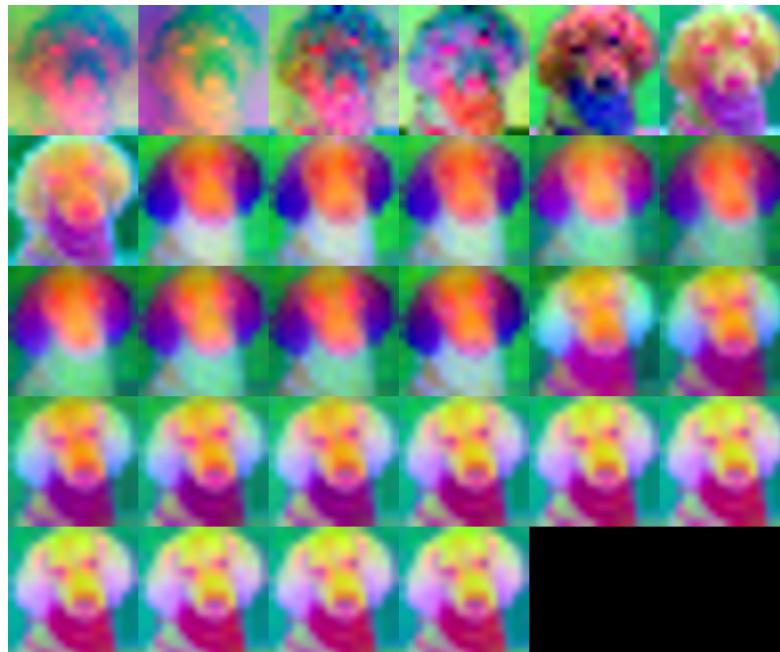
```

1509
1510
1511

```

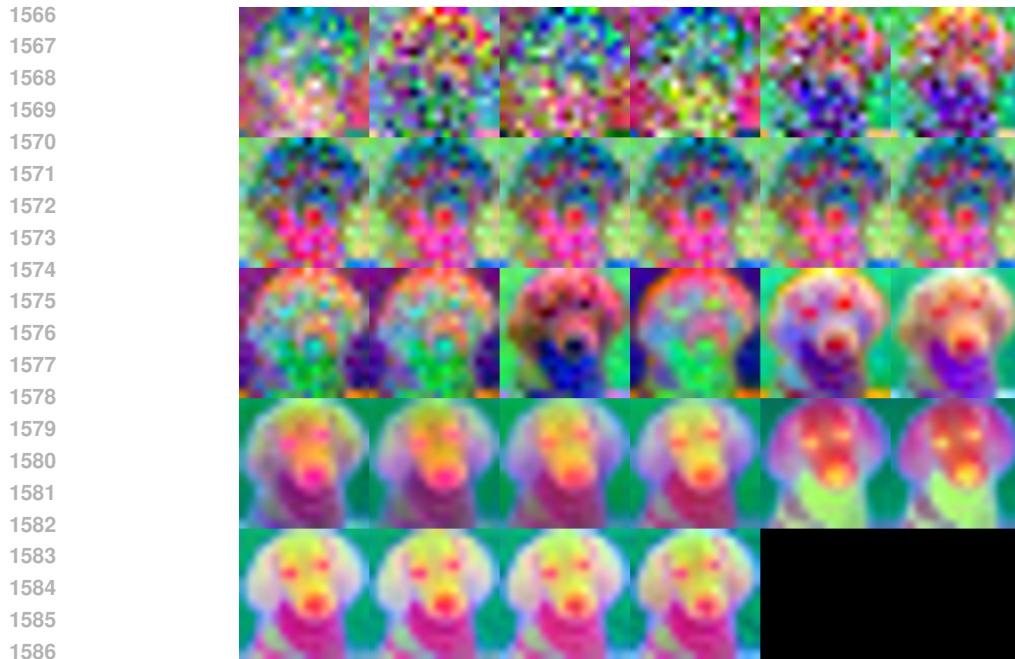


1533 **Figure 12: Visualization of SiT-XL/2 model features with 10% noise added to the input image.**
1534 The top-left plot shows the features from the first block, and subsequent blocks are visualized row
1535 by row, ending with the final block in the bottom-right corner.
1536



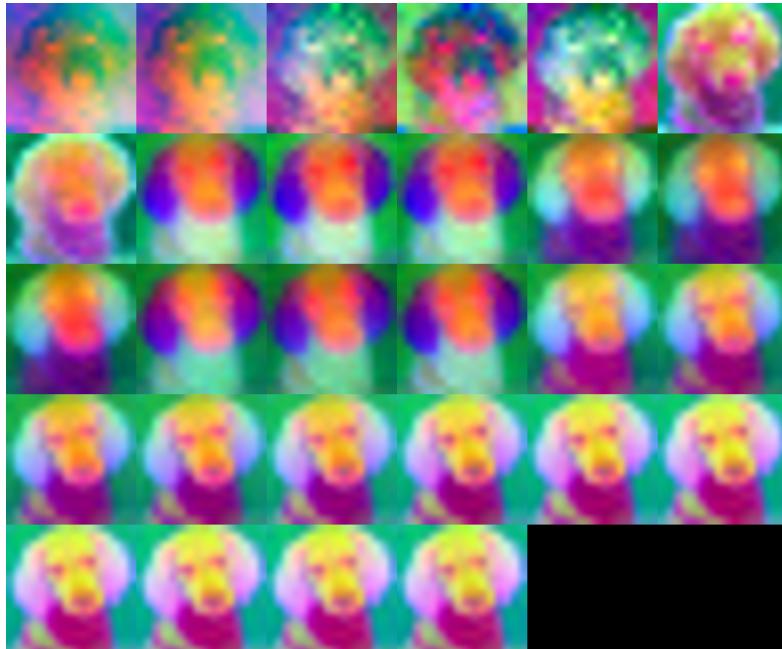
1561 **Figure 13: Visualization of SiT-XL/2 model + LayerSync features with 10% noise added to the**
1562 **input image.** The top-left plot shows the features from the first block, and subsequent blocks are visualized row
1563 by row, ending with the final block in the bottom-right corner.
1564

1565



1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613

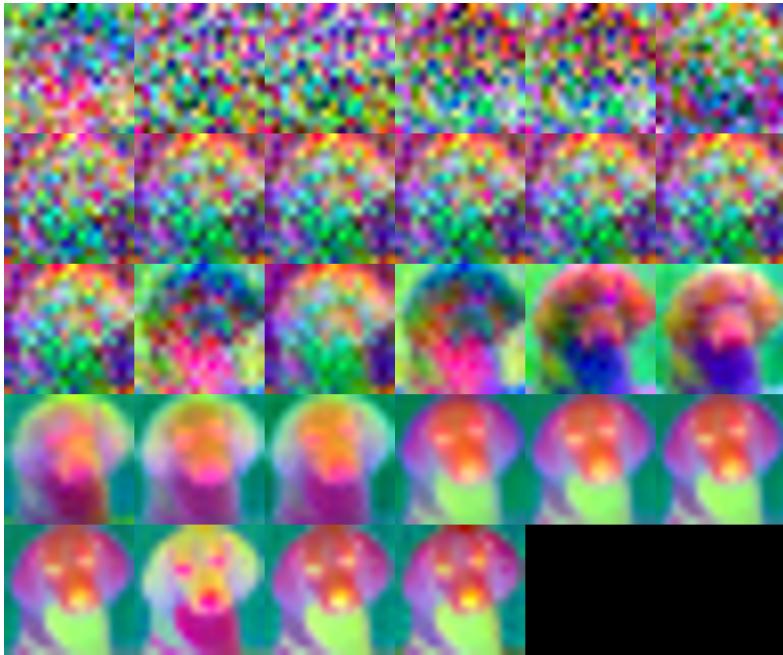
Figure 14: Visualization of SiT-XL/2 model features with 30% noise added to the input image. The top-left plot shows the features from the first block, and subsequent blocks are visualized row by row, ending with the final block in the bottom-right corner.



1614
1615
1616
1617
1618
1619

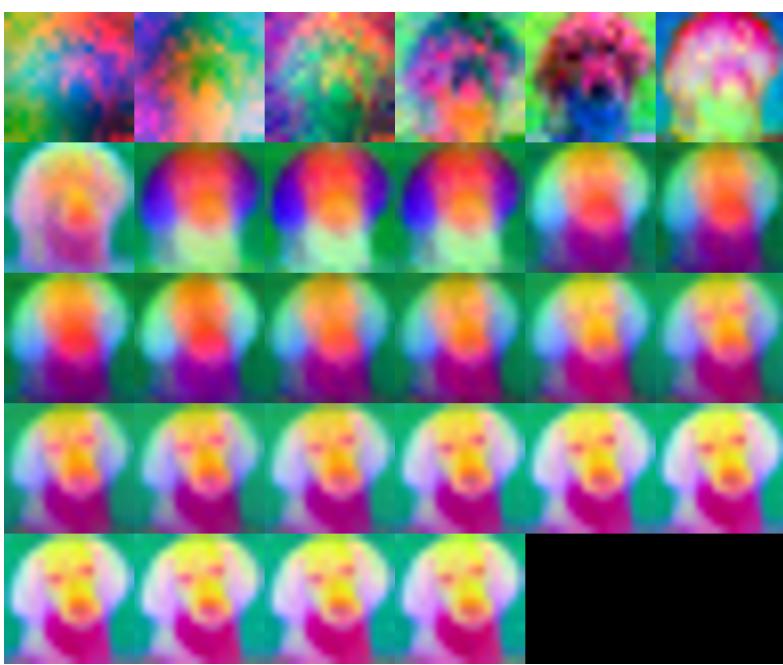
Figure 15: Visualization of SiT-XL/2 model + LayerSync features with 30% noise added to the input image. The top-left plot shows the features from the first block, and subsequent blocks are visualized row by row, ending with the final block in the bottom-right corner.

1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640



1641 **Figure 16: Visualization of SiT-XL/2 model features with 50% noise added to the input image.**
1642 The top-left plot shows the features from the first block, and subsequent blocks are visualized row
1643 by row, ending with the final block in the bottom-right corner.
1644

1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665

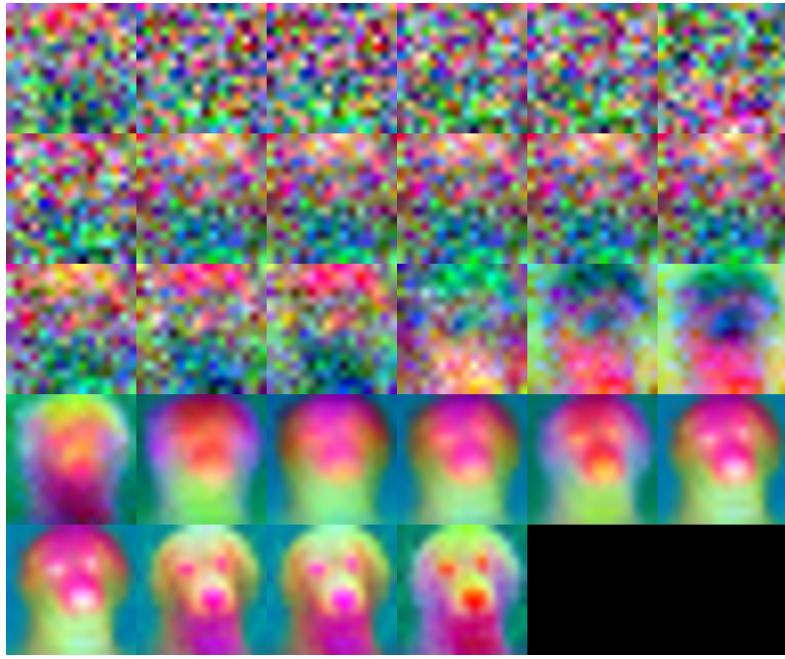


1666 **Figure 17: Visualization of SiT-XL/2 model + LayerSync features with 50% noise added to the**
1667 **input image.** The top-left plot shows the features from the first block, and subsequent blocks are visualized row
1668 by row, ending with the final block in the bottom-right corner.
1669
1670

S QUALITATIVE EXAMPLES

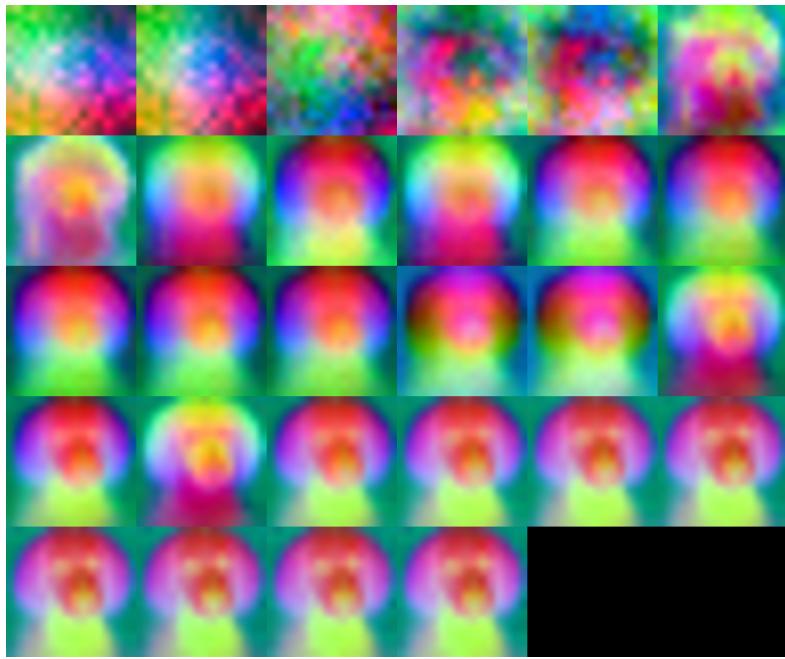
1671 We provide qualitative examples in Figure 22. The model is trained for 800 on ImageNet dataset
1672 (Deng et al., 2009) and the samples are generated using classifier-free guidance with a scale of 4 and
1673

1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694



1695 **Figure 18: Visualization of SiT-XL/2 model features with 70% noise added to the input image.**
1696 The top-left plot shows the features from the first block, and subsequent blocks are visualized row
1697 by row, ending with the final block in the bottom-right corner.
1698

1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720

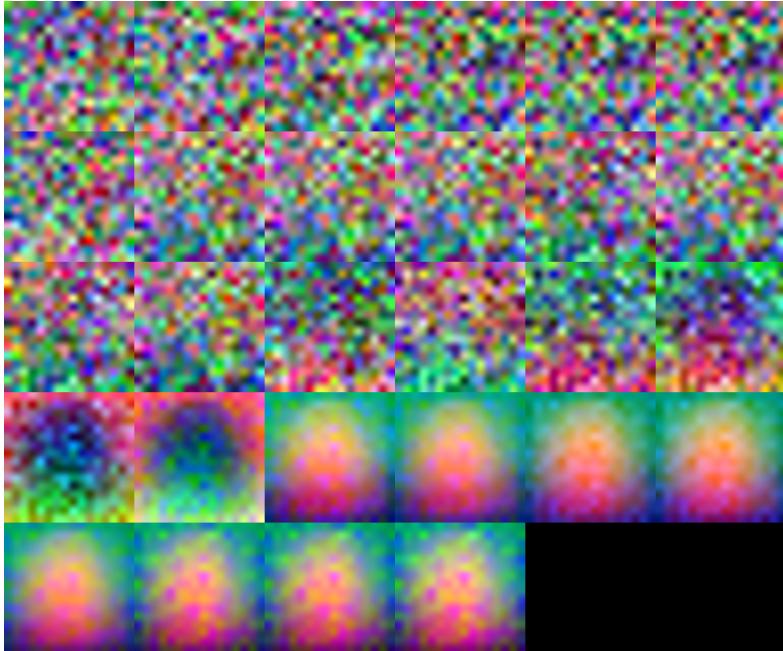


1721 **Figure 19: Visualization of SiT-XL/2 model + LayerSync features with 70% noise added to the**
1722 **input image.** The top-left plot shows the features from the first block, and subsequent blocks are visualized row
1723 by row, ending with the final block in the bottom-right corner.
1724
1725

1726
1727

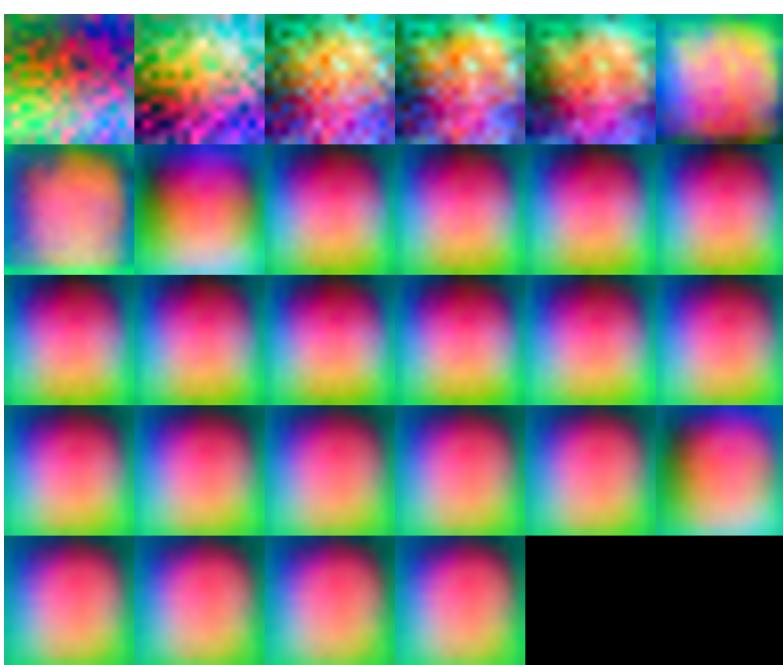
the ODE Heun sampler. Additional qualitative comparisons between the baseline SiT-XL/2, SiT-XL/2 regularized with Dispersive, and SiT-XL/2 regularized with LayerSync trained on ImageNet dataset (Deng et al., 2009) are shown in Figure 23. All models are trained for 400K iterations and

1728
 1729
 1730
 1731
 1732
 1733
 1734
 1735
 1736
 1737
 1738
 1739
 1740
 1741
 1742
 1743
 1744
 1745
 1746
 1747
 1748



1749 **Figure 20: Visualization of SiT-XL/2 model features with 90% noise added to the input image.**
 1750 The top-left plot shows the features from the first block, and subsequent blocks are visualized row
 1751 by row, ending with the final block in the bottom-right corner.
 1752

1753
 1754
 1755
 1756
 1757
 1758
 1759
 1760
 1761
 1762
 1763
 1764
 1765
 1766
 1767
 1768
 1769
 1770
 1771
 1772
 1773



1774 **Figure 21: Visualization of SiT-XL/2 model + LayerSync features with 90% noise added to the**
 1775 **input image.** The top-left plot shows the features from the first block, and subsequent blocks are visualized row
 1776 by row, ending with the final block in the bottom-right corner.
 1777

1778
 1779
 1780
 1781

share the same noise, sampler, and number of sampling steps. The samples are generated using ODE Heun sampler and no classifier-free guidance is used. LayerSync improves generation quality without relying on external representation.

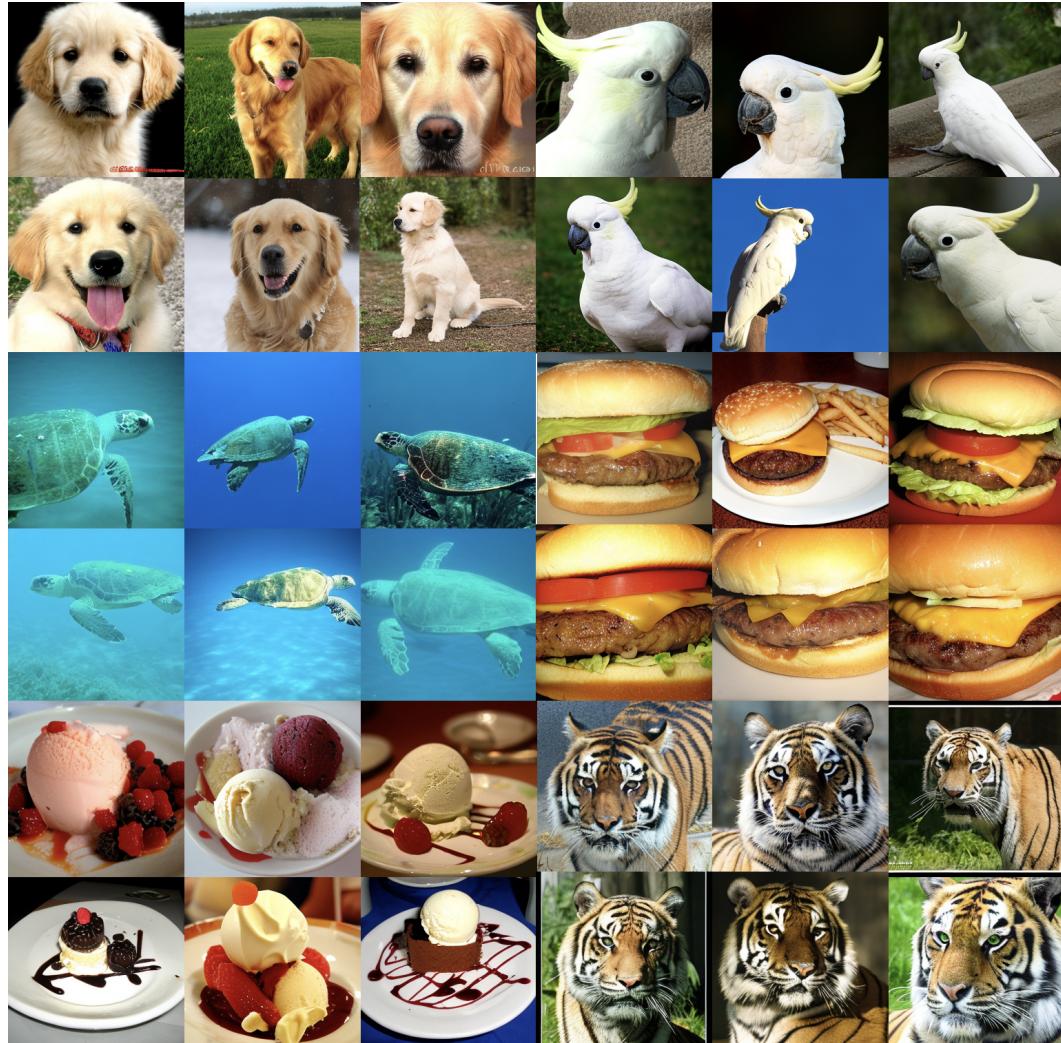


Figure 22: **Selected samples from the SiT XL/2 with LayerSync on ImageNet 256×256.** We use classifier free guidance with a cfg of 4.0.

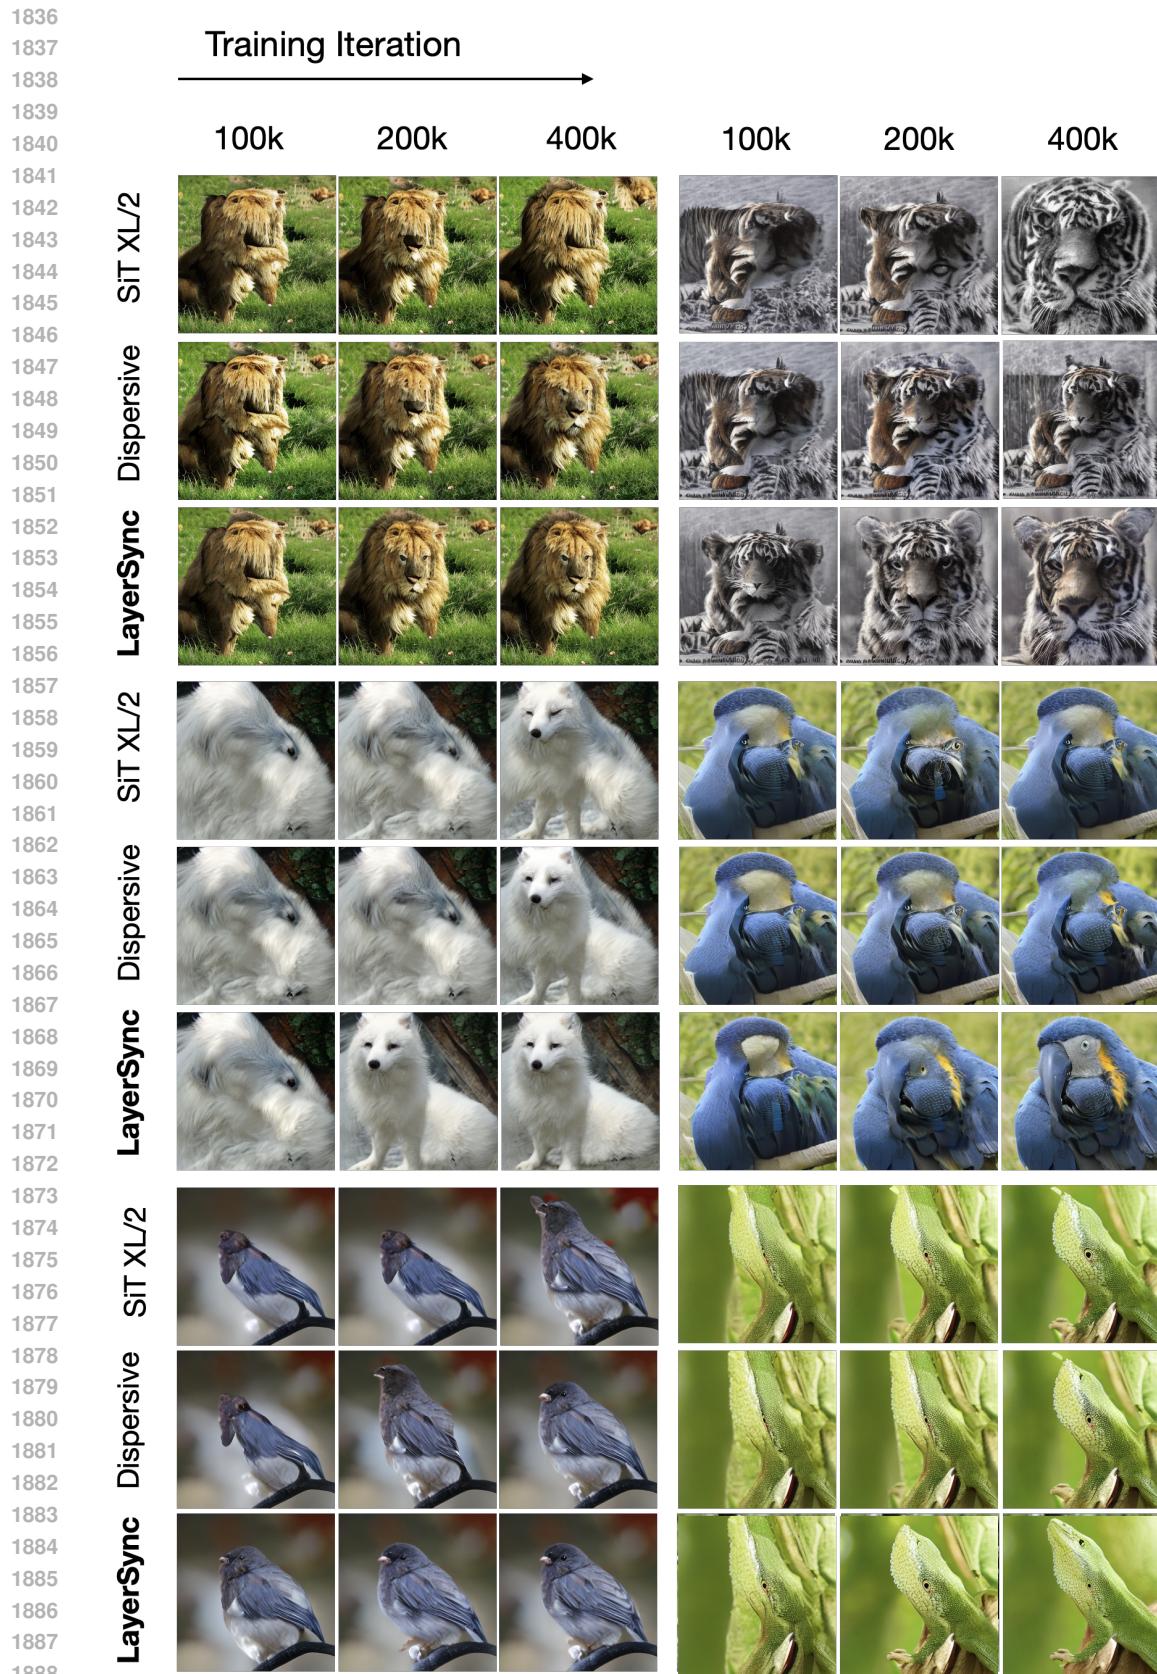


Figure 23: Qualitative comparison of SiT-XL/2 when regularized with Dispersive and LayerSync.