

# DEEP COMPRESSION AUTOENCODER FOR EFFICIENT HIGH-RESOLUTION DIFFUSION MODELS

Anonymous authors

Paper under double-blind review

## ABSTRACT

We present Deep Compression Autoencoder (DC-AE), a new family of autoencoder models for accelerating high-resolution diffusion models. Existing autoencoder models have demonstrated impressive results at a moderate spatial compression ratio (e.g.,  $8\times$ ), but fail to maintain satisfactory reconstruction accuracy for high spatial compression ratios (e.g.,  $64\times$ ). We address this challenge by introducing two key techniques: (1) **Residual Autoencoding**, where we design our models to learn residuals based on the space-to-channel transformed features to alleviate the optimization difficulty of high spatial-compression autoencoders; (2) **Decoupled High-Resolution Adaptation**, an efficient decoupled three-phase training strategy for mitigating the generalization penalty of high spatial-compression autoencoders. With these designs, we improve the autoencoder’s spatial compression ratio up to 128 while maintaining the reconstruction quality. Applying our DC-AE to latent diffusion models, we achieve significant speedup without accuracy drop. For example, on ImageNet  $512 \times 512$ , our DC-AE provides  $19.1\times$  inference speedup and  $17.9\times$  training speedup on H100 GPU for UViT-H while achieving a better FID, compared with the widely used SD-VAE-f8 autoencoder. Our code and models will be released upon publication.

## 1 INTRODUCTION

Latent diffusion models (Rombach et al., 2022) have emerged as a leading framework and demonstrated great success in image synthesis (Labs, 2024; Esser et al., 2024). It employs an autoencoder model to project the images to the latent space to reduce the training and inference costs of diffusion models. For example, the predominantly adopted solution in current latent diffusion models (Rombach et al., 2022; Labs, 2024; Esser et al., 2024; Chen et al., 2024b;a) is to use an autoencoder with a spatial compression ratio of 8 (denoted as f8), which converts images of spatial size  $H \times W$  to latent features of spatial size  $\frac{H}{8} \times \frac{W}{8}$ . This spatial compression ratio is satisfactory for low-resolution image synthesis (e.g.,  $256 \times 256$ ). However, for high-resolution image synthesis (e.g.,  $1024 \times 1024$ ), further increasing the spatial compression ratio is critical, especially for diffusion transformer models (Peebles & Xie, 2023; Bao et al., 2023) that have quadratic computational complexity to the number of tokens.

The current common practice for further reducing the spatial size is downsampling on the diffusion model side. For example, in diffusion transformer models (Peebles & Xie, 2023; Bao et al., 2023), this is achieved by using a patch embedding layer with patch size  $p$  that compresses the latent features to  $\frac{H}{8p} \times \frac{W}{8p}$  tokens. In contrast, little effort has been made on the autoencoder side. The main bottleneck hindering the employment of high spatial-compression autoencoders is the reconstruction accuracy drop. For example, Figure 2 (a) shows the reconstruction results of SD-VAE (Rombach et al., 2022) on ImageNet  $256 \times 256$  with different spatial compression ratios. We can see that the rFID (reconstruction FID) degrades from 0.90 to 28.3 if switching from f8 to f64.

This work presents **Deep Compression Autoencoder (DC-AE)**, a new family of high spatial-compression autoencoders for efficient high-resolution image synthesis. By analyzing the underlying source of the accuracy degradation between high spatial-compression and low spatial-compression autoencoders, we find high spatial-compression autoencoders are more difficult to optimize (Section 3.1) and suffer from the generalization penalty across resolutions (Figure 3 b). To this end, we introduce two key techniques to address these two challenges. First, we propose **Resid-**




Figure 1: DC-AE accelerates diffusion models by increasing autoencoder’s spatial compression ratio.




Figure 2: (a) **Image Reconstruction Results on ImageNet 256x256**.  $f$  denotes the spatial compression ratio. When the spatial compression ratio increases, SD-VAE has a significant reconstruction accuracy drop (higher rFID) while DC-AE does not have this issue. (b) **ImageNet 512x512 Image Generation Results on UViT-S with Various Autoencoders**.  $p$  denotes the patch size. Shifting the token compression task to the autoencoder enables the diffusion model to focus more on the denoising task, leading to better FID. (c) **Comparison to SD-VAE-f8 on ImageNet 512x512 with UViT Variants**. DC-AE-f64p1 provides 19.1x higher inference throughput and 0.54 better ImageNet FID than SD-VAE-f8p2 on UViT-H.

**ual Autoencoding** (Figure 4) to alleviate the optimization difficulty of high spatial-compression autoencoders. It introduces extra non-parametric shortcuts to the autoencoder model to let the neural network modules learn residuals based on the space-to-channel operation. Second, we propose **Decoupled High-Resolution Adaptation** (Figure 6) to tackle the other challenge. It introduces a high-resolution latent adaptation phase and a low-resolution local refinement phase to avoid the generalization penalty while maintaining a low training cost.

With these techniques, we increase the spatial compression ratio of autoencoders to 32, 64, and 128 while maintaining good reconstruction accuracy (Table 2). The diffusion models can fully focus on the denoising task with our DC-AE taking over the whole token compression task, which delivers better image generation results than prior approaches (Table 3). For example, replacing SD-VAE-f8 with our DC-AE-f64, we achieve 17.9x higher H100 training throughput and 19.1x higher H100 inference throughput on UViT-H (Bao et al., 2023) while improving the ImageNet 512 x 512 FID from 3.55 to 3.01. Our pre-trained models and code will be released upon publication. We summarize our contributions as follows:

- We analyze the challenges of increasing the spatial compression ratio of autoencoders and provide insights into how to address these challenges.
- We propose Residual Autoencoding and Decoupled High-Resolution Adaptation that effectively improve the reconstruction accuracy of high spatial-compression autoencoders, making their reconstruction accuracy feasible for use in latent diffusion models.
- We build DC-AE, a new family of autoencoder models based on our techniques. DC-AE delivers significant training and inference speedup for latent diffusion models compared with prior autoencoder models.

## 2 RELATED WORK

**Autoencoder for Diffusion Models.** Training and evaluating diffusion models directly in high-resolution pixel space results in prohibitive computational costs. To address this issue, Rombach et al. (2022) proposes latent diffusion models that operate in a compressed latent space produced by pretrained autoencoders. The proposed autoencoder with 8x spatial compression ratio and 4




Figure 3: (a) High spatial-compression autoencoders are more difficult to optimize. Even with the same latent shape and stronger learning capacity, it still cannot match the f8 autoencoder’s rFID. (b) High spatial-compression autoencoders suffer from significant reconstruction accuracy drops when generalizing from low-resolution to high-resolution.

latent channels has been widely adopted in subsequent works (Peebles & Xie, 2023; Bao et al., 2023). Since then, follow-up works mainly focus on enhancing the reconstruction accuracy of the f8 autoencoder by increasing the number of latent channels (Esser et al., 2024; Dai et al., 2023; Labs, 2024). Additionally, to improve the reconstruction quality, especially for image editing tasks, Zhu et al. (2023) leverages a heavier decoder and incorporates task-specific priors. In contrast to prior works, our work focuses on an orthogonal direction, increasing the spatial compression ratio of the autoencoders (e.g., f64). To the best of our knowledge, our work is the first study in this critical but underexplored direction.

**Diffusion Model Acceleration.** Diffusion models have been widely used for image generation and showed impressive results (Labs, 2024; Esser et al., 2024). However, diffusion models are computationally intensive, motivating many works to accelerate diffusion models. One representative strategy is reducing the number of inference sampling steps by training-free few-step samplers (Song et al., 2021; Lu et al., 2022a;b; Zheng et al., 2023; Zhang & Chen, 2023; Zhang et al., 2023; Zhao et al., 2024b; Shih et al., 2024; Tang et al., 2024) or distilling-based methods (Meng et al., 2023; Salimans & Ho, 2022; Yin et al., 2024b;a; Song et al., 2023; Luo et al., 2023; Liu et al., 2022). Another representative strategy is model compression by leveraging sparsity (Li et al., 2022; Ma et al., 2024) or quantization (He et al., 2024; Fang et al., 2024; Li et al., 2023; Zhao et al., 2024a). Designing efficient diffusion model architectures (Li et al., 2024c; Liu et al., 2024; Cai et al., 2024) or inference systems (Li et al., 2024b; Wang et al., 2024) is also an effective approach for boosting efficiency. In addition, improving the data quality (Chen et al., 2024b;a) can boost the training efficiency of diffusion models.

All these techniques focus on the diffusion model while the autoencoder remains the same. Our work opens up a new direction for accelerating diffusion models, which can benefit both training and inference.

### 3 METHOD

In this section, we first analyze why existing high spatial-compression autoencoders (e.g., SD-VAE-f64) fail to match the accuracy of low spatial-compression autoencoders (e.g., SD-VAE-f8). Then we introduce our Deep Compression Autoencoder (DC-AE) with *Residual Autoencoding* and *Decoupled High-Resolution Adaptation* to close the accuracy gap. Finally, we discuss the applications of our DC-AE to latent diffusion models.

#### 3.1 MOTIVATION

We conduct ablation study experiments to get insights into the underlying source of the accuracy gap between high spatial-compression and low spatial-compression autoencoders. Specifically, we consider three settings with gradually increased spatial compression ratio, from f8 to f64.

Each time the spatial compression ratio increases, we stack additional encoder and decoder stages upon the current autoencoder model. In this way, high spatial-compression autoencoders contain low spatial-compression autoencoders as sub-networks and thus have higher learning capacity.

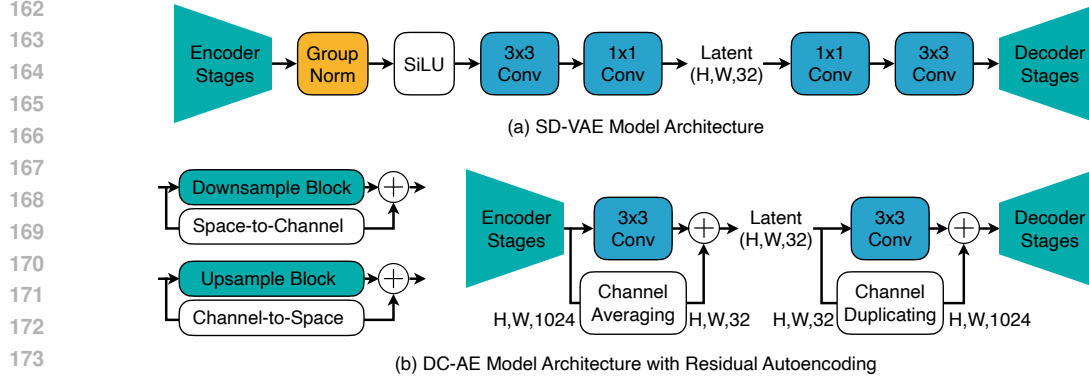


Figure 4: **Illustration of Residual Autoencoding.** It adds non-parametric shortcuts to let the neural network modules learn residuals based on the space-to-channel operation.

Additionally, we increase the latent channel number to maintain the same total latent size across different settings. We can then convert the latent to a higher spatial compression ratio one by applying a space-to-channel operation (Shi et al., 2016):  $H \times W \times C \rightarrow \frac{H}{p} \times \frac{W}{p} \times p^2 C$ .

We summarize the results in Figure 3 (a, gray dash line). Even with the same total latent size and stronger learning capacity, we still observe degraded reconstruction accuracy when the spatial compression ratio increases. It demonstrates that *the added encoder and decoder stages (consisting of multiple SD-VAE building blocks) work worse than a simple space-to-channel operation.*

Based on this finding, we conjecture *the accuracy gap comes from the model learning process: while we have good local optimums in the parameter space, the optimization difficulty hinders high spatial-compression autoencoders from reaching such local optimums.*

### 3.2 DEEP COMPRESSION AUTOENCODER

**Residual Autoencoding.** Motivated by the analysis, we introduce Residual Autoencoding to address the accuracy gap. The general idea is depicted in Figure 4. The core difference from the conventional design is that we explicitly let neural network modules learn the downsample residuals based on the space-to-channel operation to alleviate the optimization difficulty. Different from ResNet (He et al., 2016), the residual here is not identity mapping, but space-to-channel mapping.

In practice, this is implemented by adding extra non-parametric shortcuts on the encoder’s downsample blocks and decoder’s upsample blocks (Figure 4 b, left). Specifically, for the downsample block, the non-parametric shortcut is a space-to-channel operation followed by a non-parametric channel averaging operation to match the channel number. For example, assuming the downsample block’s input feature map shape is  $H \times W \times C$  and its output feature map shape is  $\frac{H}{2} \times \frac{W}{2} \times 2C$ , then the added shortcut is

$$\begin{aligned}
 H \times W \times C &\xrightarrow{\text{space-to-channel}} \frac{H}{2} \times \frac{W}{2} \times 4C \\
 &\xrightarrow{\text{split into two groups}} \underbrace{\left[ \frac{H}{2} \times \frac{W}{2} \times 2C, \frac{H}{2} \times \frac{W}{2} \times 2C \right]}_{\text{channel averaging}} \xrightarrow{\text{average}} \frac{H}{2} \times \frac{W}{2} \times 2C.
 \end{aligned}$$

Accordingly, for the upsample block, the non-parametric shortcut is a channel-to-space operation followed by a non-parametric channel duplicating operation

$$\begin{aligned}
 \frac{H}{2} \times \frac{W}{2} \times 2C &\xrightarrow{\text{channel-to-space}} H \times W \times \frac{C}{2} \\
 &\xrightarrow{\text{duplicate}} \underbrace{\left[ H \times W \times \frac{C}{2}, H \times W \times \frac{C}{2} \right]}_{\text{channel duplicating}} \xrightarrow{\text{concat}} H \times W \times C.
 \end{aligned}$$

In addition to the downsample and upsample blocks, we also change the middle stage design following the same principle (Figure 4 b, right).





238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254

Figure 5: Autoencoder already learns to reconstruct content and semantics without GAN loss, while GAN loss improves local details and removes local artifacts. We replace the GAN loss full training with lightweight local refinement training which achieves the same goal and has lower training cost.

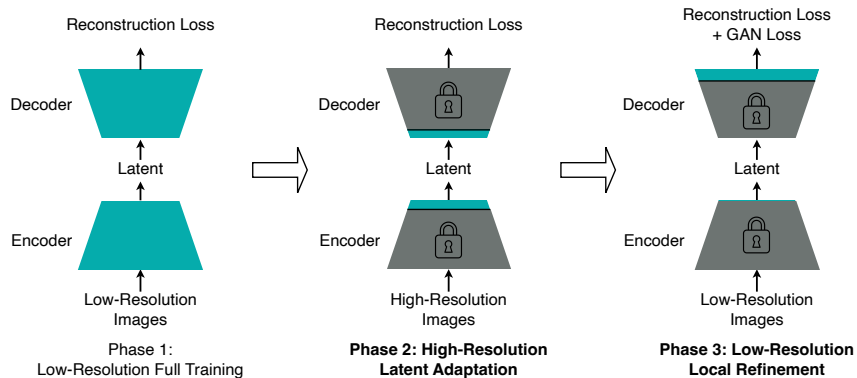


Figure 6: Illustration of Decoupled High-Resolution Adaptation.

Figure 3 (a) shows the comparison with and without our Residual Autoencoding on ImageNet  $256 \times 256$ . We can see that Residual Autoencoding effectively improves the reconstruction accuracy of high spatial-compression autoencoders.

**Decoupled High-Resolution Adaptation.** Residual Autoencoding alone can address the accuracy gap when handling low-resolution images. However, when extending it to high-resolution images, we find it not sufficient. Due to the large cost of high-resolution training, the common practice for high-resolution diffusion models is directly using autoencoders trained on low-resolution images (e.g.,  $256 \times 256$ ) (Chen et al., 2024b;a). This strategy works well for low spatial-compression autoencoders. However, high spatial-compression autoencoders suffer from a significant accuracy drop. For example, in Figure 3 (b), we can see that f64 autoencoder’s rFID degrades from 0.50 to 7.40 when generalizing from  $256 \times 256$  to  $1024 \times 1024$ . In contrast, the f8 autoencoder’s rFID improves from 0.51 to 0.19 under the same setting. Additionally, we also find this issue more severe when using a higher spatial compression ratio. In this work, we refer to this phenomenon as the *generalization penalty of high spatial-compression autoencoders*. A straightforward solution to

<b>ImageNet 512×512 (Class-Conditional)</b>					
Diffusion Model	Autoencoder	Patch Size	#Tokens	FID (w/o CFG) ↓	FID (w/ CFG) ↓
UViT-S [1]	SD-VAE-f8	8	64	125.08	95.93
	SD-VAE-f16	4	64	115.32	88.06
	SD-VAE-f32	2	64	107.33	76.57
	DC-AE-f64	1	64	<b>67.30</b>	<b>35.96</b>

Table 1: Ablation Study on Patch Size and Autoencoder’s Spatial Compression Ratio.

address this issue is conducting training on high-resolution images. However, it suffers from a large training cost and unstable high-resolution GAN loss training.

We introduce Decoupled High-Resolution Adaptation to tackle this challenge. Figure 6 demonstrates the detailed training pipeline. Compared with the conventional single-phase training strategy (Rombach et al., 2022), our Decoupled High-Resolution Adaptation has two key differences.

First, we decouple the GAN loss training from the full model training and introduce a dedicated local refinement phase for the GAN loss training. In the local refinement phase (Figure 6, phase 3), we only tune the head layers of the decoder while freezing all the other layers. The intuition of this design is based on the finding that the reconstruction loss alone is sufficient for learning to reconstruct the content and semantics. Meanwhile, the GAN loss mainly improves local details and removes local artifacts (Figure 5). Achieving the same goal of local refinement, only tuning the decoder’s head layers has a lower training cost and delivers better accuracy than the full training.

Moreover, the decoupling prevents the GAN loss training from changing the latent space. This approach enables us to conduct the local refinement phase on low-resolution images without worrying about the generalization penalty. This further reduces the training cost of phase 3 and avoids the highly unstable high-resolution GAN loss training.

Second, we introduce an additional high-resolution latent adaptation phase (Figure 6, phase 2) that tunes the middle layers (i.e., encoder’s head layers and decoder’s input layers) to adapt the latent space for alleviating the generalization penalty. In our experiments, we find only tuning middle layers is sufficient for addressing this issue (Figure 3 b) while having a lower training cost than high-resolution full training (memory cost: 153.98 GB  $\rightarrow$  67.81 GB)<sup>1</sup>.

### 3.3 APPLICATION TO LATENT DIFFUSION MODELS

Applying our DC-AE to latent diffusion models is straightforward. The only hyperparameter to change is the patch size (Peebles & Xie, 2023). For diffusion transformer models (Peebles & Xie, 2023; Bao et al., 2023), increasing the patch size  $p$  is the common approach for reducing the number of tokens. It is equivalent to first applying the space-to-channel operation to reduce the spatial size of the given latent by  $p\times$  and then using the transformer model with a patch size of 1.

Since combining a low spatial-compression autoencoder (e.g., f8) with the space-to-channel operation can also achieve a high spatial compression ratio, a natural question is how it compares with directly reaching the target spatial compression ratio with DC-AE.

We conduct ablation study experiments and summarize the results in Table 1. We can see that directly reaching the target spatial compression ratio with the autoencoder gives the best results among all settings. In addition, we also find that shifting the spatial compression ratio from the diffusion model to the autoencoder consistently leads to better FID.

## 4 EXPERIMENTS

### 4.1 SETUPS

**Implementation Details.** We use a mixture of datasets to train autoencoders (baselines and DC-AE), containing ImageNet (Deng et al., 2009), SAM (Kirillov et al., 2023), MapillaryVistas (Neuhold et al., 2017), and FFHQ (Karras et al., 2019). For ImageNet experiments, we exclusively

<sup>1</sup>Assuming the input resolution is  $1024 \times 1024$  and the batch size is 12.

<b>ImageNet 256×256</b>	Latent Shape	Autoencoder	rFID ↓	PSNR ↑	SSIM ↑	LPIPS ↓
f32c32	8×8×32	SD-VAE [36]	2.64	22.13	0.59	0.117
		DC-AE	<b>0.69</b>	<b>23.85</b>	<b>0.66</b>	<b>0.082</b>
f64c128	4×4×128	SD-VAE [36]	26.65	18.07	0.41	0.283
		DC-AE	<b>0.81</b>	<b>23.60</b>	<b>0.65</b>	<b>0.087</b>
<b>ImageNet 512×512</b>	Latent Shape	Autoencoder	rFID ↓	PSNR ↑	SSIM ↑	LPIPS ↓
f64c128	8×8×128	SD-VAE [36]	16.84	19.49	0.48	0.282
		DC-AE	<b>0.22</b>	<b>26.15</b>	<b>0.71</b>	<b>0.080</b>
f128c512	4×4×512	SD-VAE [36]	100.74	15.90	0.40	0.531
		DC-AE	<b>0.23</b>	<b>25.73</b>	<b>0.70</b>	<b>0.084</b>
<b>FFHQ 1024×1024</b>	Latent Shape	Autoencoder	rFID ↓	PSNR ↑	SSIM ↑	LPIPS ↓
f64c128	16×16×128	SD-VAE [36]	6.62	24.55	0.68	0.237
		DC-AE	<b>0.23</b>	<b>31.04</b>	<b>0.83</b>	<b>0.061</b>
f128c512	8×8×512	SD-VAE [36]	179.71	18.11	0.63	0.585
		DC-AE	<b>0.41</b>	<b>31.18</b>	<b>0.83</b>	<b>0.062</b>
<b>MapillaryVistas 2048×2048</b>	Latent Shape	Autoencoder	rFID ↓	PSNR ↑	SSIM ↑	LPIPS ↓
f64c128	32×32×128	SD-VAE [36]	7.55	22.37	0.68	0.262
		DC-AE	<b>0.36</b>	<b>29.57</b>	<b>0.84</b>	<b>0.075</b>
f128c512	16×16×512	SD-VAE [36]	152.09	17.82	0.67	0.594
		DC-AE	<b>0.38</b>	<b>29.70</b>	<b>0.84</b>	<b>0.074</b>

Table 2: Image Reconstruction Results.

use the ImageNet training split to train autoencoders and diffusion models. The model architecture is similar to SD-VAE (Rombach et al., 2022) except for our new designs discussed in Section 3.2. In addition, we use the original autoencoders instead of the variational autoencoders for our models, as they perform the same in our experiments and the original autoencoders are simpler. We also replace transformer blocks with EfficientViT blocks (Cai et al., 2023) to make autoencoders more friendly for handling high-resolution images while maintaining similar accuracy.

For image generation experiments, we apply autoencoders to diffusion transformer models including DiT (Peebles & Xie, 2023) and UViT (Bao et al., 2023). We follow the same training settings as the original papers. We consider three settings with different resolutions, including ImageNet (Deng et al., 2009) for  $512 \times 512$  generation, FFHQ (Karras et al., 2019) and MJHQ (Li et al., 2024a) for  $1024 \times 1024$  generation, and MapillaryVistas (Neuhold et al., 2017) for  $2048 \times 2048$  generation.

**Efficiency Profiling.** We profile the training and inference throughput on the H100 GPU with PyTorch and TensorRT respectively. The latency is measured on the 3090 GPU with batch size 2. The training memory is profiled using PyTorch, assuming a batch size of 256. We use fp16 for all cases. For simplicity, we assume the number of sampling steps is 1.

## 4.2 IMAGE COMPRESSION AND RECONSTRUCTION

Table 2 summarizes the results of DC-AE and SD-VAE (Rombach et al., 2022) under various settings (f represents the spatial compression ratio and c denotes the number of latent channels). DC-AE provides significant reconstruction accuracy improvements than SD-VAE for all cases. For example, on ImageNet  $512 \times 512$ , DC-AE improves the rFID from 16.84 to 0.22 for the f64c128 autoencoder and 100.74 to 0.23 for the f128c512 autoencoder.

In addition to the quantitative results, Figure 7 shows image reconstruction samples produced by SD-VAE and DC-AE. Reconstructed images by DC-AE demonstrate a better visual quality than SD-VAE’s reconstructed images. In particular, for the f64 and f128 autoencoders, DC-AE still maintains a good visual quality for small texts and human faces.

Diffusion Model	Autoencoder	Patch Size	Throughput (image/s) ↑		Latency (ms) ↓	Memory (GB) ↓	FID ↓	
			Training	Inference			w/o CFG	w/ CFG
UViT-S [1]	SD3-VAE-f8 [8]	2	352	2984/T	3.8T	13.8	164.34	143.82
	Flux-VAE-f8 [18]	2	352	2984/T	3.8T	13.8	106.07	84.73
	SDXL-VAE-f8 [35]	2	352	2991/T	3.8T	13.8	51.03	26.38
	Asym-VAE-f8 [53]	2	352	2991/T	3.8T	13.8	51.96	24.57
	SD-VAE-f8 [36]	2	352	2991/T	3.8T	13.8	51.96	24.57
	SD-VAE-f16 [36]	2	1550	12881/T	1.3T	4.0	76.86	44.22
	SD-VAE-f32 [36]	1	1551	12883/T	1.3T	4.0	70.23	38.63
	DC-AE-f32	1	1553	12850/T	1.3T	4.0	<b>46.12</b>	<b>18.08</b>
	DC-AE-f64	1	<b>6295</b>	<b>53774/T</b>	<b>0.7T</b>	<b>1.5</b>	67.30	35.96
	DC-AE-f64 <sup>†</sup>	1	<b>6295</b>	<b>53774/T</b>	<b>0.7T</b>	<b>1.5</b>	61.84	30.63
DiT-XL [34]	Flux-VAE-f8 [18]	2	54	416/T	31.7T	56.3	27.35	8.72
	Asym-VAE-f8 [53]	2	54	424/T	31.7T	56.2	11.55	2.95
	SD-VAE-f8 [36]	2	54	424/T	31.7T	56.2	12.03	3.04
	DC-AE-f32	1	<b>241</b>	<b>2016/T</b>	<b>7.8T</b>	<b>20.9</b>	<b>9.56</b>	<b>2.84</b>
UViT-H [1]	Flux-VAE-f8 [18]	2	55	349/T	30.4T	54.2	30.91	12.63
	Asym-VAE-f8 [53]	2	55	351/T	30.3T	54.1	11.36	3.51
	SD-VAE-f8 [36]	2	55	351/T	30.3T	54.1	11.04	3.55
	DC-AE-f32	1	247	1622/T	8.2T	18.6	<b>9.83</b>	<b>2.53</b>
	DC-AE-f64	1	<b>984</b>	<b>6706/T</b>	<b>3.5T</b>	<b>10.6</b>	13.96	3.01
	DC-AE-f64 <sup>†</sup>	1	<b>984</b>	<b>6706/T</b>	<b>3.5T</b>	<b>10.6</b>	12.26	2.66

Table 3: **Class-Conditional Image Generation Results on ImageNet 512×512.** † represents the model is trained for 4× training iterations (i.e., 500K → 2,000K iterations). ‘T’ denotes the diffusion sampling steps.

FFHQ 1024×1024 (Unconditional) & MJHQ 1024×1024 (Class-Conditional)										
Diffusion Model	Autoencoder	Patch Size	Throughput (image/s) ↑		Latency (ms) ↓	Memory (GB) ↓	FFHQ FID ↓		MJHQ FID ↓	
			Training	Inference			w/o CFG	w/ CFG	w/o CFG	w/ CFG
DiT-S [34]	SD3-VAE-f8 [8]	2	83	814/T	14.2T	41.4	46.28	109.43	103.02	
	Flux-VAE-f8 [18]	2	83	814/T	14.2T	41.4	59.15	143.16	139.06	
	SDXL-VAE-f8 [35]	2	84	833/T	14.1T	41.2	16.82	49.00	39.21	
	Asym-VAE-f8 [53]	2	84	833/T	14.1T	41.2	17.12	48.25	38.36	
	SD-VAE-f8 [36]	2	84	833/T	14.1T	41.2	16.98	48.05	38.19	
			4	470	5566/T	2.5T	10.7	23.81	60.94	51.29
	DC-AE-f32	1	475	5575/T	2.5T	10.7	<b>13.65</b>	<b>34.35</b>	<b>27.20</b>	
	DC-AE-f64	1	<b>2085</b>	<b>25259/T</b>	<b>1.0T</b>	<b>3.1</b>	26.88	61.30	53.38	
	<b>MapillaryVistas 2048×2048 (Unconditional)</b>									
	Diffusion Model	Autoencoder	Patch Size	Throughput (image/s) ↑		Latency (ms) ↓	Memory (GB) ↓	MapillaryVistas FID ↓		
Training				Inference	w/o CFG			w/ CFG	w/ CFG	
DiT-S [34]	SD-VAE-f8 [36]	4	84	810/T	14.3T	41.4	69.50			
	DC-AE-f64	1	<b>459</b>	<b>5435/T</b>	<b>2.6T</b>	<b>11.0</b>	<b>59.55</b>			

Table 4: **1024×1024 and 2048×2048 Image Generation Results.**

### 4.3 LATENT DIFFUSION MODELS

We compare DC-AE with the widely used SD-VAE-f8 autoencoder (Rombach et al., 2022) on various diffusion transformer models. For DC-AE, we always use a patch size of 1 (denoted as p1). For SD-VAE-f8, we follow the common setting and use a patch size of 2 or 4 (denoted as p2, p4). The results are summarized in Table 3, Table 4, and Table 5.



Diffusion Model	Autoencoder	Patch Size	Throughput (image/s) $\uparrow$		Latency (ms) $\downarrow$	Memory (GB) $\downarrow$	MJHQ $512 \times 512$	
			Training	Inference			FID $\downarrow$	CLIP Score $\uparrow$
PIXART- $\alpha$ [5]	SD-VAE-f8 [36]	2	43	312/T	37.1T	60.45	6.3	26.36
	DC-AE-f32	1	173	1251/T	10.4T	23.77	6.1	26.41

Table 5: Text-to-Image Generation Results.

	f32	f64	f128				
Original							
	SD-VAE						
		DC-AE					
			Original				
				SD-VAE			
					DC-AE		

Figure 7: **Autoencoder Image Reconstruction Samples.** We select representative images to visualize the reconstruction results. The images are reconstructed at resolution  $1024 \times 1024$ . The samples are cropped for better visualization of details like human faces and small texts.

**ImageNet  $512 \times 512$ .** As shown in Table 3, DC-AE-f32p1 consistently delivers better FID than SD-VAE-f8p2 on all diffusion transformer models. In addition, it has  $4\times$  fewer tokens than SD-VAE-f8p2, leading to  $4.5\times$  higher H100 training throughput and  $4.8\times$  higher H100 inference



486  
487  
488  
489  
490  
491  
492  
493  
494  
495  
496  
497  
498  
499  
500  
501  
502  
503  
504  
505  
506  
507  
508  
509  
510  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539




Figure 8: We show selected samples generated by the diffusion models using our DC-AE.

throughput for DiT-XL. We also observe that larger diffusion transformer models seem to benefit more from our DC-AE. For example, DC-AE-f64p1 has a worse FID than SD-VAE-f8p2 on UViT-S but a better FID on UViT-H. We conjecture it is because DC-AE-f64 has a larger latent channel number than SD-VAE-f8, thus needing more model capacity (Esser et al., 2024).

**1024×1024 and 2048×2048 Image Generation.** Apart from ImageNet 512×512, we also test our models for higher-resolution image generation. As shown in Table 4, we have a similar finding where DC-AE-f32p1 achieves better FID than SD-VAE-f8p2 for all cases.

**Text-to-Image Generation.** Table 5 reports our text-to-image generation results on PIXART- $\alpha$  (Chen et al., 2024b). All models are trained for 100K iterations from scratch. Similar to prior cases, we observe DC-AE-f32p1 provides a better FID and a better CLIP Score than SD-VAE-f8p2. Figure 8 demonstrates samples generated by the diffusion models with our DC-AE, showing the capacity to synthesize high-quality images while being significantly more efficient than prior models.

## 5 CONCLUSION

We accelerate high-resolution diffusion models by designing deep compression autoencoders to reduce the number of tokens. We proposed two techniques: *residual autoencoding* and *decoupled high-resolution adaptation* to address the challenges brought by the high compression ratio. The resulting new autoencoder model family DC-AE demonstrated satisfactory reconstruction accuracy with a spatial compression ratio of up to 128. DC-AE also demonstrated significant training and inference efficiency improvements when applied to latent diffusion models.

## REFERENCES

- 540  
541  
542 Fan Bao, Shen Nie, Kaiwen Xue, Yue Cao, Chongxuan Li, Hang Su, and Jun Zhu. All are worth  
543 words: A vit backbone for diffusion models. In *Proceedings of the IEEE/CVF conference on*  
544 *computer vision and pattern recognition*, pp. 22669–22679, 2023. 1, 2, 3, 6, 7, 8, 15, 18
- 545 Han Cai, Junyan Li, Muyan Hu, Chuang Gan, and Song Han. Efficientvit: Lightweight multi-scale  
546 attention for high-resolution dense prediction. In *Proceedings of the IEEE/CVF International*  
547 *Conference on Computer Vision*, pp. 17302–17313, 2023. 7
- 548  
549 Han Cai, MUYANG LI, QINSHENG ZHANG, MING-YU LIU, and SONG HAN. Condition-aware neural net-  
550 work for controlled image generation. In *Proceedings of the IEEE/CVF Conference on Computer*  
551 *Vision and Pattern Recognition*, pp. 7194–7203, 2024. 3
- 552 Junsong Chen, Chongjian Ge, Enze Xie, Yue Wu, Lewei Yao, Xiaozhe Ren, Zhongdao Wang, Ping  
553 Luo, Huchuan Lu, and Zhenguo Li. Pixart- $\sigma$ : Weak-to-strong training of diffusion transformer  
554 for 4k text-to-image generation. *arXiv preprint arXiv:2403.04692*, 2024a. 1, 3, 5, 15
- 555  
556 Junsong Chen, YU Jincheng, GE Chongjian, Lewei Yao, Enze Xie, Zhongdao Wang, James Kwok,  
557 Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart- $\alpha$ : Fast training of diffusion transformer for  
558 photorealistic text-to-image synthesis. In *International Conference on Learning Representations*,  
559 2024b. 1, 3, 5, 9, 10, 15
- 560 Xiaoliang Dai, Ji Hou, Chih-Yao Ma, Sam Tsai, Jialiang Wang, Rui Wang, Peizhao Zhang, Simon  
561 Vandenhende, Xiaofang Wang, Abhimanyu Dubey, et al. Emu: Enhancing image generation  
562 models using photogenic needles in a haystack. *arXiv preprint arXiv:2309.15807*, 2023. 3
- 563  
564 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hi-  
565 erarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,  
566 pp. 248–255. Ieee, 2009. 6, 7
- 567 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam  
568 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for  
569 high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*,  
570 2024. 1, 3, 8, 10, 15, 18
- 571  
572 Gongfan Fang, Xinyin Ma, and Xinchao Wang. Structural pruning for diffusion models. *Advances*  
573 *in Neural Information Processing Systems*, 36, 2024. 3
- 574 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-  
575 nition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.  
576 770–778, 2016. 4
- 577  
578 Yefei He, Luping Liu, Jing Liu, Weijia Wu, Hong Zhou, and Bohan Zhuang. Ptqd: Accurate post-  
579 training quantization for diffusion models. *Advances in Neural Information Processing Systems*,  
580 36, 2024. 3
- 581 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*  
582 *neural information processing systems*, 33:6840–6851, 2020. 15
- 583  
584 Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with  
585 conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and*  
586 *pattern recognition*, pp. 1125–1134, 2017. 16
- 587 Sadeep Jayasumana, Srikumar Ramalingam, Andreas Veit, Daniel Glasner, Ayan Chakrabarti, and  
588 Sanjiv Kumar. Rethinking fid: Towards a better evaluation metric for image generation. In *Pro-*  
589 *ceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9307–  
590 9315, 2024. 17
- 591  
592 Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative  
593 adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*  
*recognition*, pp. 4401–4410, 2019. 6, 7

- 594 Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete  
595 Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Pro-*  
596 *ceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4015–4026, 2023.  
597 6
- 598 Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved  
599 precision and recall metric for assessing generative models. *Advances in neural information*  
600 *processing systems*, 32, 2019. 17
- 602 Black Forest Labs. Flux. *Online*, 2024. URL [https://github.com/](https://github.com/black-forest-labs/flux)  
603 [black-forest-labs/flux](https://github.com/black-forest-labs/flux). 1, 3, 8, 15, 18
- 605 Daiqing Li, Aleks Kamko, Ehsan Akhgari, Ali Sabet, Linmiao Xu, and Suhail Doshi. Playground  
606 v2. 5: Three insights towards enhancing aesthetic quality in text-to-image generation. *arXiv*  
607 *preprint arXiv:2402.17245*, 2024a. 7, 21
- 608 Muyang Li, Ji Lin, Chenlin Meng, Stefano Ermon, Song Han, and Jun-Yan Zhu. Efficient spa-  
609 tially sparse inference for conditional gans and diffusion models. *Advances in neural information*  
610 *processing systems*, 35:28858–28873, 2022. 3
- 612 Muyang Li, Tianle Cai, Jiaxin Cao, Qinsheng Zhang, Han Cai, Junjie Bai, Yangqing Jia, Kai Li,  
613 and Song Han. Distrifusion: Distributed parallel inference for high-resolution diffusion models.  
614 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.  
615 7183–7193, 2024b. 3
- 616 Xiuyu Li, Yijiang Liu, Long Lian, Huanrui Yang, Zhen Dong, Daniel Kang, Shanghang Zhang,  
617 and Kurt Keutzer. Q-diffusion: Quantizing diffusion models. In *Proceedings of the IEEE/CVF*  
618 *International Conference on Computer Vision*, pp. 17535–17545, 2023. 3
- 620 Yanyu Li, Huan Wang, Qing Jin, Ju Hu, Pavlo Chemerys, Yun Fu, Yanzhi Wang, Sergey Tulyakov,  
621 and Jian Ren. Snapfusion: Text-to-image diffusion model on mobile devices within two seconds.  
622 *Advances in Neural Information Processing Systems*, 36, 2024c. 3
- 623 Songhua Liu, Weihao Yu, Zhenxiong Tan, and Xinchao Wang. Linfusion: 1 gpu, 1 minute, 16k  
624 image. *arXiv preprint arXiv:2409.02097*, 2024. 3
- 626 Xingchao Liu, Xiwen Zhang, Jianzhu Ma, Jian Peng, et al. InstafLOW: One step is enough for  
627 high-quality diffusion-based text-to-image generation. In *The Twelfth International Conference*  
628 *on Learning Representations*, 2023. 3
- 629 I Loshchilov. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 16
- 631 Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast  
632 ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural*  
633 *Information Processing Systems*, 35:5775–5787, 2022a. 3, 15
- 634 Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast  
635 solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*,  
636 2022b. 3
- 637  
638 Simian Luo, Yiqin Tan, Longbo Huang, Jian Li, and Hang Zhao. Latent consistency models: Synthe-  
639 sizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*, 2023.  
640 3
- 642 Xinyin Ma, Gongfan Fang, and Xinchao Wang. Deepcache: Accelerating diffusion models for free.  
643 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.  
644 15762–15772, 2024. 3
- 645  
646 Heusel Martin, Ramsauer Hubert, Unterthiner Thomas, Nessler Bernhard, and Hochreiter Sepp.  
647 Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in*  
*neural information processing systems*, 30:6626–6637, 2017. 17

- 648 Chenlin Meng, Robin Rombach, Ruiqi Gao, Diederik Kingma, Stefano Ermon, Jonathan Ho, and  
649 Tim Salimans. On distillation of guided diffusion models. In *Proceedings of the IEEE/CVF*  
650 *Conference on Computer Vision and Pattern Recognition*, pp. 14297–14306, 2023. 3
- 651
- 652 Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kotschieder. The mapillary vistas  
653 dataset for semantic understanding of street scenes. In *Proceedings of the IEEE international*  
654 *conference on computer vision*, pp. 4990–4999, 2017. 6, 7
- 655 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of*  
656 *the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023. 1, 3, 6, 7, 8,  
657 15, 18
- 658 Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe  
659 Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image  
660 synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 8, 18
- 661
- 662 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-  
663 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-*  
664 *ence on computer vision and pattern recognition*, pp. 10684–10695, 2022. 1, 2, 6, 7, 8, 9, 15, 17,  
665 18
- 666 Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In  
667 *International Conference on Learning Representations*, 2022. 3
- 668
- 669 Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen.  
670 Improved techniques for training gans. *Advances in neural information processing systems*, 29,  
671 2016. 17
- 672 Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel  
673 Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient  
674 sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer*  
675 *vision and pattern recognition*, pp. 1874–1883, 2016. 4
- 676 Andy Shih, Suneel Belkhale, Stefano Ermon, Dorsa Sadigh, and Nima Anari. Parallel sampling of  
677 diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- 678
- 679 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Interna-*  
680 *tional Conference on Learning Representations*, 2021. 3
- 681 Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *International*  
682 *Conference on Machine Learning*, pp. 32211–32252. PMLR, 2023. 3
- 683
- 684 Zhiwei Tang, Jiasheng Tang, Hao Luo, Fan Wang, and Tsung-Hui Chang. Accelerating parallel  
685 sampling of diffusion models. In *Forty-first International Conference on Machine Learning*, 2024.  
686 3
- 687 Jiannan Wang, Jiarui Fang, Aoyu Li, and PengCheng Yang. Pipefusion: Displaced patch pipeline  
688 parallelism for inference of diffusion transformer models. *arXiv preprint arXiv:2405.14430*,  
689 2024. 3
- 690
- 691 Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and  
692 William T Freeman. Improved distribution matching distillation for fast image synthesis. *arXiv*  
693 *preprint arXiv:2405.14867*, 2024a. 3
- 694 Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman,  
695 and Taesung Park. One-step diffusion with distribution matching distillation. In *Proceedings of*  
696 *the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6613–6623, 2024b.  
697 3
- 698 Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator.  
699 In *The Eleventh International Conference on Learning Representations*, 2023. 3
- 700
- 701 Qinsheng Zhang, Molei Tao, and Yongxin Chen. gddim: Generalized denoising diffusion implicit  
models. In *International Conference on Learning Representations*, 2023. 3



702 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable  
703 effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on*  
704 *computer vision and pattern recognition*, pp. 586–595, 2018. 16

705  
706 Tianchen Zhao, Tongcheng Fang, Enshu Liu, Wan Rui, Widyadewi Soedarmadji, Shiyao Li, Zinan  
707 Lin, Guohao Dai, Shengen Yan, Huazhong Yang, et al. Vedit-q: Efficient and accurate quantiza-  
708 tion of diffusion transformers for image and video generation. *arXiv preprint arXiv:2406.02540*,  
709 2024a. 3

710 Wenliang Zhao, Lujia Bai, Yongming Rao, Jie Zhou, and Jiwen Lu. Unipc: A unified predictor-  
711 corrector framework for fast sampling of diffusion models. *Advances in Neural Information*  
712 *Processing Systems*, 36, 2024b. 3

713 Kaiwen Zheng, Cheng Lu, Jianfei Chen, and Jun Zhu. Dpm-solver-v3: Improved diffusion ode  
714 solver with empirical model statistics. *Advances in Neural Information Processing Systems*, 36:  
715 55502–55542, 2023. 3

716  
717 Zixin Zhu, Xuelu Feng, Dongdong Chen, Jianmin Bao, Le Wang, Yinpeng Chen, Lu Yuan, and Gang  
718 Hua. Designing a better asymmetric vqgan for stablediffusion. *arXiv preprint arXiv:2306.04632*,  
719 2023. 3, 8, 18

720  
721  
722  
723  
724  
725  
726  
727  
728  
729  
730  
731  
732  
733  
734  
735  
736  
737  
738  
739  
740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755



## A ADDITIONAL SAMPLES

In Figure 14 and 15, we provide additional image reconstruction samples produced by SD-VAE and DC-AE. Reconstructed images by DC-AE demonstrate better visual qualities than SD-VAE’s reconstructed images, especially for the f64 and f128 autoencoders. Some samples are cropped for better visualization of details like human faces and small texts.

In Figure 16 and Figure 17, we show randomly generated samples on ImageNet 512×512 and MJHQ-30K 512×512 by the diffusion models using our DC-AE.

## B DIFFUSION SAMPLING HYPERPARAMETERS

For the DiT models, we use the DDPM (Ho et al., 2020) sampler from the DiT (Peebles & Xie, 2023) codebase with 250 sampling steps and a guidance scale of 1.3.

For the UViT models, we use the DPMSolver (Lu et al., 2022a) sampler with 30 sampling steps and a guidance scale of 1.5.




Figure 9: Ablation Study on Diffusion Sampling Hyperparameters. We use the DPMSolver sampler for both DiT-XL and UViT-H. DC-AE provides significant speedup over the baseline models while maintaining the generation performance under different diffusion sampling hyperparameters.

## C ABLATION STUDY EXPERIMENTS ON DC-AE

ImageNet 512×512	Latent Shape	Autoencoder	rFID ↓	PSNR ↑	SSIM ↑	LPIPS ↓
f64c128	8×8×128	SD-VAE [36]	16.84	19.49	0.48	0.282
		SD-VAE [36] w/ DHRA	5.54	21.13	0.54	0.228
		DC-AE	<b>0.22</b>	<b>26.15</b>	<b>0.71</b>	<b>0.080</b>

Table 6: Ablation Study Experiments on DC-AE. ‘DHRA’ represents the decoupled high-resolution adaptation.

Table 6 reports ablation study results on DC-AE. We can see that both residual autoencoding and decoupled high-resolution adaptation contribute significantly to DC-AE’s superior performances in high spatial-compression settings.

## D LATENT SCALING AND SHIFTING FACTORS

Following the common practice (Rombach et al., 2022; Peebles & Xie, 2023; Bao et al., 2023; Esser et al., 2024; Labs, 2024; Chen et al., 2024b;a), we normalize the latent space of our autoencoders to apply to latent diffusion models. Given a dataset, we compute the root mean square of the latent features and use its multiplicative inverse as the scaling factor for our autoencoders. We do not use the shifting factor for our autoencoders.

## E DC-AE ARCHITECTURE AND TRAINING DETAILS

We present the detailed architecture of DC-AE encoder and decoder stages in Figure 10 to complement Figure 4 (b).

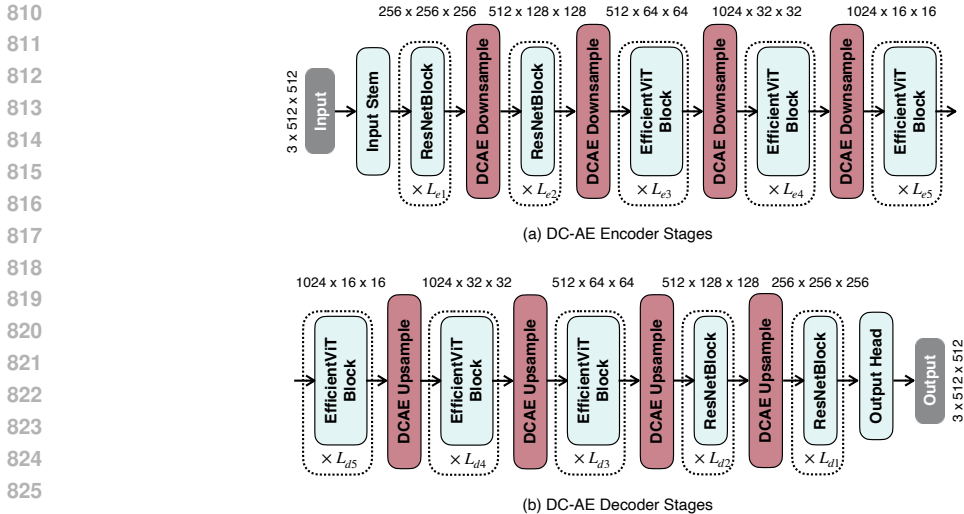


Figure 10: Detailed Architecture of DC-AE Encoder and Decoder Stages.

We use the AdamW optimizer (Loshchilov, 2017) for all training phases. In phase 1 (low-resolution full training), we use a constant learning rate of  $6.4e-5$  with a weight decay of 0.1, and AdamW betas of (0.9, 0.999). We use L1 loss and LPIPS loss (Zhang et al., 2018). In phase 2 (high-resolution latent adaptation), we use a constant learning rate of  $1.6e-5$ , a weight decay of 0.001, and AdamW betas of (0.9, 0.999). We use the same loss as phase 1. In phase 3 (low-resolution local refinement), we use a constant learning rate of  $5.4e-5$ , and AdamW betas of (0.5, 0.9). We use L1 loss, LPIPS loss (Zhang et al., 2018), and PatchGAN loss (Isola et al., 2017).

### F INVESTIGATIONS ON THE GENERALIZATION GAP

Table 7 demonstrates additional investigations on the generalization gap of high spatial-compression autoencoders. Training with image crops leads to worse results than training with downscale images in our case. Training with 50% downscale images and 50% image crops can improve the rFID from 7.4 to 2.7. It shows that this strategy can partially address the generalization gap. However, this strategy is still inferior to our decoupled high-resolution adaptation.

Method	Downscale	Crop	50% Downscale, 50% Crop	Decoupled High-Resolution Adaptation
rFID ↓	7.40	8.32	2.70	<b>0.18</b>

Table 7: Investigations on the Generalization Gap.

### G ABLATION STUDY ON TRAINING DIFFERENT NUMBERS OF LAYERS

Figure 11 presents the ablation study on training different numbers of layers in phase 2 (high-resolution latent adaptation) and phase 3 (low-resolution local refinement).




Figure 11: Ablation Study on Training Different Numbers of Layers in Phase 2 (Left) and Phase 3 (Right).

## H LOSS LANDSCAPE COMPARISON

Figure 12 demonstrates the loss landscape comparison between DC-AE and SD-VAE. We can see that DC-AE’s loss landscape is flatter than SD-VAE’s, indicating that it is easier for DC-AE to reach low-loss regions than SD-VAE.




Figure 12: **Loss Landscape Comparison.** DC-AE’s loss landscape is flatter than SD-VAE’s, indicating that it is easier for DC-AE to reach low-loss regions.

## I ADDITIONAL IMAGE RECONSTRUCTION RESULTS

Table 8 reports the reconstruction results under the low spatial-compression ratio setting. DC-AE delivers slightly better results than SD-VAE under this setting.

ImageNet 256×256	Latent Shape	Autoencoder	rFID ↓	PSNR ↑	SSIM ↑	LPIPS ↓
f8c4	32×32×4	SD-VAE [36]	0.63	24.99	0.71	0.063
		DC-AE	<b>0.46</b>	<b>25.46</b>	<b>0.73</b>	<b>0.057</b>

Table 8: **Image Reconstruction Results under the Low Spatial-Compression Ratio Setting.**

## J IMAGE GENERATION RESULTS WITH OTHER EVALUATION METRICS

Table 9 presents a comprehensive evaluation of different diffusion models and autoencoders on ImageNet 512×512. The evaluation metrics include FID (Martin et al., 2017), inception score (IS) (Salimans et al., 2016), precision, recall (Kynkäänniemi et al., 2019), and CMMD (Jayasumana et al., 2024). Our DC-AE consistently delivers significant efficiency improvements while maintaining the generation performance under different evaluation metrics.

## K MODEL SCALING RESULTS




Figure 13: **Model Scaling Results on ImageNet 512×512 (UViT-S to UViT-2B).** DC-AE-f64 benefits more from scaling up than SD-VAE-f8.

918  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939  
940  
941  
942  
943  
944  
945  
946  
947  
948  
949  
950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971

Diffusion Model	Autoencoder	Patch Size	Inference Throughput	FID ↓		Inception Score ↑		Precision ↑		Recall ↑		CMMD ↓	
				w/o CFG	w/ CFG	w/o CFG	w/ CFG	w/o CFG	w/ CFG	w/o CFG	w/ CFG		
UViT-S [1]	SD3-VAE-f8 [8]	2	2984/T	164.34	143.82	6.07	7.53	0.06	0.09	0.31	0.39	3.13	2.94
	Flux-VAE-f8 [18]	2	2984/T	106.07	84.73	13.39	17.71	0.28	0.37	0.39	0.42	1.90	1.67
	SDXL-VAE-f8 [35]	2	2991/T	51.03	26.38	27.58	56.72	0.57	0.74	0.58	0.50	1.35	1.05
	Asym-VAE-f8 [53]	2	2991/T	52.68	25.14	30.22	65.27	0.58	0.74	0.62	0.51	1.09	0.80
	SD-VAE-f8 [36]	2	2991/T	51.96	24.57	30.37	65.73	0.57	0.74	0.64	0.52	1.23	0.91
	SD-VAE-f16 [36]	2	12881/T	76.86	44.22	21.38	43.35	0.43	0.62	0.60	0.55	1.83	1.46
	SD-VAE-f32 [36]	1	12883/T	70.23	38.63	23.07	47.72	0.46	0.64	0.58	0.56	1.71	1.36
	DC-AE-f32	1	12850/T	46.12	18.08	34.82	84.73	0.59	0.76	0.66	0.56	1.00	0.70
	DC-AE-f64	1	53774/T	67.30	35.96	24.55	52.86	0.44	0.64	0.60	0.56	1.44	1.14
DC-AE-f64 <sup>†</sup>	1	53774/T	61.84	30.63	27.28	61.76	0.47	0.67	0.63	0.56	1.35	1.04	
DiT-XL [34]	Flux-VAE-f8 [18]	2	416/T	27.35	8.72	53.09	130.20	0.68	0.83	0.61	0.48	0.54	0.30
	Asym-VAE-f8 [53]	2	424/T	11.39	2.97	108.70	241.10	0.75	0.83	0.65	0.53	0.37	0.20
	SD-VAE-f8 [36]	2	424/T	12.03	3.04	105.25	240.82	0.75	0.84	0.64	0.54	0.43	0.25
	DC-AE-f32	1	2016/T	9.56	2.84	117.49	226.98	0.75	0.82	0.64	0.55	0.34	0.22
UViT-H [1]	Flux-VAE-f8 [18]	2	349/T	30.91	12.63	56.72	127.93	0.64	0.76	0.59	0.49	0.50	0.31
	Asym-VAE-f8 [53]	2	351/T	11.36	3.51	124.24	249.21	0.75	0.82	0.61	0.53	0.32	0.20
	SD-VAE-f8 [36]	2	351/T	11.04	3.55	125.08	250.66	0.75	0.82	0.61	0.53	0.39	0.26
	DC-AE-f32	1	1622/T	9.83	2.53	121.91	255.07	0.76	0.83	0.65	0.54	0.34	0.20
	DC-AE-f64	1	6706/T	13.96	3.01	99.20	229.16	0.73	0.83	0.64	0.53	0.50	0.31
DC-AE-f64 <sup>†</sup>	1	6706/T	12.26	2.66	109.20	239.82	0.73	0.82	0.67	0.57	0.43	0.27	
UViT-2B [1]	Flux-VAE-f8 [18]	2	155/T	25.03	10.12	74.04	161.29	0.67	0.78	0.58	0.51	0.38	0.24
	Asym-VAE-f8 [53]	2	157/T	9.87	3.62	131.95	258.63	0.76	0.83	0.59	0.52	0.30	0.19
	SD-VAE-f8 [36]	2	157/T	9.73	3.57	132.86	260.50	0.76	0.83	0.59	0.52	0.37	0.24
	DC-AE-f32	1	665/T	8.13	2.30	135.44	272.73	0.76	0.82	0.66	0.56	0.30	0.17
	DC-AE-f64	1	2733/T	7.78	2.47	138.11	280.49	0.77	0.84	0.63	0.54	0.35	0.22
DC-AE-f64 <sup>†</sup>	1	2733/T	6.50	2.25	152.35	293.45	0.77	0.83	0.65	0.56	0.31	0.19	

Table 9: **Class-Conditional Image Generation Results on ImageNet 512×512 with More Evaluation Metrics.** <sup>†</sup> represents the model is trained for 4× training iterations (i.e., 500K → 2,000K iterations). ‘T’ denotes the diffusion sampling steps.

In addition to existing UViT models, we scaled the model up to 1.6B parameters, with a depth of 28, a hidden dimension of 2048, and 32 heads. We denote this model as UViT-2B. Figure 13 demonstrates that DC-AE-f64 benefits more from scaling up than SD-VAE-f8.





Figure 14: More Autoencoder Image Reconstruction Samples.



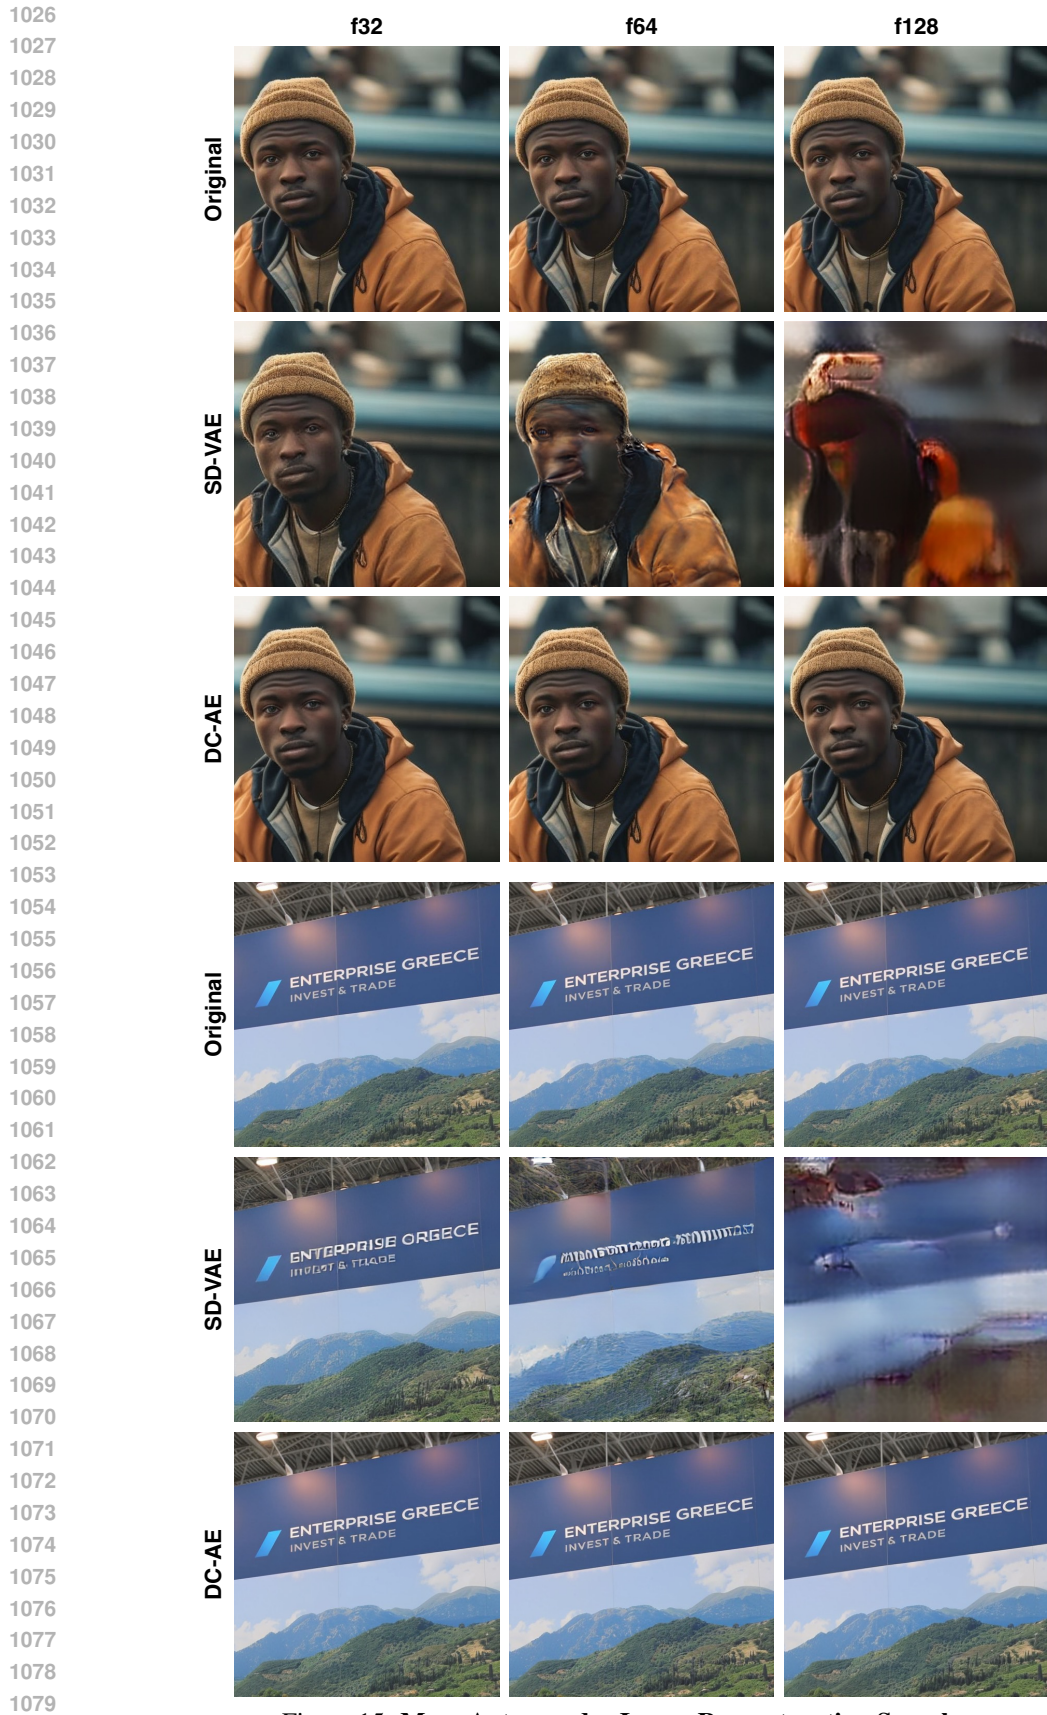


Figure 15: More Autoencoder Image Reconstruction Samples.

1080  
1081  
1082  
1083  
1084  
1085  
1086  
1087  
1088  
1089  
1090  
1091  
1092  
1093  
1094  
1095  
1096  
1097  
1098  
1099  
1100  
1101  
1102  
1103  
1104  
1105  
1106  
1107  
1108  
1109  
1110  
1111  
1112  
1113  
1114  
1115  
1116  
1117  
1118  
1119  
1120  
1121  
1122  
1123  
1124  
1125  
1126  
1127  
1128  
1129  
1130  
1131  
1132  
1133




Prompt: logo of a crow's head, looking in the camera, symmetrical, colorful, 4k



Prompt: ultra photorealistic, super ultra hdr quality, full figure, National Geographic professional photograph of a Giant recycled wooden sculpture of a grizzly bear, sculpture in the artistic style of Thomas Dambo, Mandurah estuary backdrop,



Prompt: Content Highly realistic Neanderthal face integrated onto a giant Bigfoot with flint black skin Medium Hyperrealistic digital painting Style Realism with a focus on intricate details and textures Lighting Soft, diffused natural lighting to enhance the facial features, fur texture, and overall form Colors Rich earthy tones for the fur, first black skin tone for the Neanderthal face, and subtle variations to convey depth and realism Composition Wideangle lens capturing the full figure of the giant Bigfoot with a Neanderthal face, standing in a natural environment, showcasing the seamless integration of facial features, expressive eyes, and distinctive bone structure Create a hyperrealistic digital painting depicting a highly realistic Neanderthal face integrated onto a giant Bigfoot with flint black skin, as if you were seeing it in person. Employ a realism style with a focus on intricate details and textures. Use soft, diffused natural lighting to enhance the facial features, fur texture, and overall form. Choose rich earthy tones for the fur and a flint black skin tone for the Neanderthal face, adding subtle variations to convey depth and realism. Compose the image with a wideangle lens capturing the full figure of the giant Bigfoot with a Neanderthal face, standing in a natural environment, showcasing the seamless integration of facial features, expressive eyes, and distinctive bone structure.



Prompt: a beautiful and colorful Monstera, dark, hyper realistic, highly detailed, intricate, volumetric light, natural lighting,cinematic 4k




Prompt: clever barkers cenobites from hellraizer working in an evil snack bar, demonic fast food restaurant, hyper realistic, UHD quality, scene from film



Prompt: Hooded faceless Hacker PLAYing Heavy Metal in Keyboards, Hd cinematic lighting, realistic, photorealistic, Heal hyper realistic, Ultra, hd 8k, Lelivka Ultra, Realistic, Shot on a Canon EOS 5D Mark IV with a 200mm f1. 4L IS USM lens 64 megapixels Zoomed out octane render shading and bokeh



Prompt: soulful woman in stunning fashionable yellow winter dress, hair pinned up, with adorable little baby duck and yellow tulips, editorial photography, Vogue, fashion and beauty, love, hyperrealistic, hyperdetailed




Prompt: an angel hiding deep in the dark forest behind the bushes, hiding from a drone, full body, in the style of renaissance painting, photorealistic, mysterious, cinematic, 4k




Prompt: a heart shaped kratom leaf is displayed on its own, in the style of high resolution, that art, light white and light emerald, realist Melike accuracy, moche art, rounded, nabis, logo



Prompt: The Kiss painting of Gustav Klimt in Cloud Monet style blue white gold




Prompt: Detailed portrait of cute smiling girl, cyberpunk futuristic, reflective puffy coat, decorated with hearts, by lomali incoglu dragan bbin hans thoma greg rutkowski alexandros pyromallis neko rene maritte illustrated, perfect face, fine details, realistic black and white lineart coloring page



Prompt: stock image popular fig leaf trend.



Prompt: hurt woman walks down a neighborhood road alone at night a car is crashed in a ditch off to the side, wide shot



Prompt: waterfall in a green Forest



Prompt: black and white image, portrait of a tiger, love. A National Geographic award winning stock image popular no text prompt trend, pinterest contest winner



Prompt: camping sticker white background

Figure 16: Random 512×512 Text-to-Image Samples. Prompts are randomly drawn from MJHQ-30K (Li et al., 2024a).



1134  
1135  
1136  
1137  
1138  
1139  
1140  
1141  
1142  
1143  
1144  
1145  
1146  
1147  
1148  
1149  
1150  
1151  
1152  
1153  
1154  
1155  
1156  
1157  
1158  
1159  
1160  
1161  
1162  
1163  
1164  
1165  
1166  
1167  
1168  
1169  
1170  
1171  
1172  
1173  
1174  
1175  
1176  
1177  
1178  
1179  
1180  
1181  
1182  
1183  
1184  
1185  
1186  
1187



Figure 17: Random Generated Samples on ImageNet 512x512.