

# All are Worth Words: A ViT Backbone for Diffusion Models

Fan Bao<sup>1</sup>, Shen Nie<sup>2</sup>, Kaiwen Xue<sup>2</sup>, Yue Cao<sup>3</sup>, Chongxuan Li<sup>2\*</sup>, Hang Su<sup>1</sup>, Jun Zhu<sup>1\*</sup>

<sup>1</sup>Dept. of Comp. Sci. & Tech., Institute for AI, BNRist Center

<sup>1</sup>Tsinghua-Bosch Joint ML Center, THBI Lab, Tsinghua University, Beijing, 100084 China

<sup>2</sup>Gaoling School of Artificial Intelligence, Renmin University of China,

<sup>2</sup>Beijing Key Laboratory of Big Data Management and Analysis Methods, Beijing, China

<sup>3</sup>Beijing Academy of Artificial Intelligence

bf19@mails.tsinghua.edu.cn; nieshen@ruc.edu.cn; {kevin.kaiwenxue, caoyue10}@gmail.com

chongxuanli@ruc.edu.cn; {suhangss, dcszj}@tsinghua.edu.cn

## Abstract

Vision transformers (ViT) have shown promise in various vision tasks while the U-Net based on a convolutional neural network (CNN) remains dominant in diffusion models. We design a simple and general ViT-based architecture (named U-ViT) for image generation with diffusion models. U-ViT is characterized by treating all inputs including the time, condition and noisy image patches as tokens and employing long skip connections between shallow and deep layers. We evaluate U-ViT in unconditional and class-conditional image generation, as well as text-to-image generation tasks, where U-ViT is comparable if not superior to a CNN-based U-Net of a similar size. In particular, latent diffusion models with U-ViT achieve record-breaking FID scores of 2.29 in class-conditional image generation on ImageNet 256×256, and 5.48 in text-to-image generation on MS-COCO, among methods without accessing large external datasets during the training of generative models.

Our results suggest that, for diffusion-based image modeling, the long skip connection is crucial while the down-sampling and up-sampling operators in CNN-based U-Net are not always necessary. We believe that U-ViT can provide insights for future research on backbones in diffusion models and benefit generative modeling on large scale cross-modality datasets.

## 1. Introduction

Diffusion models [24, 56, 61] are powerful deep generative models that emerge recently for high quality image generation [12, 25, 49]. They grow rapidly and find applications in text-to-image generation [47, 49, 51], image-to-image generation [10, 42, 74], video generation [23, 27],

\*Corresponding to C. Li and J. Zhu.

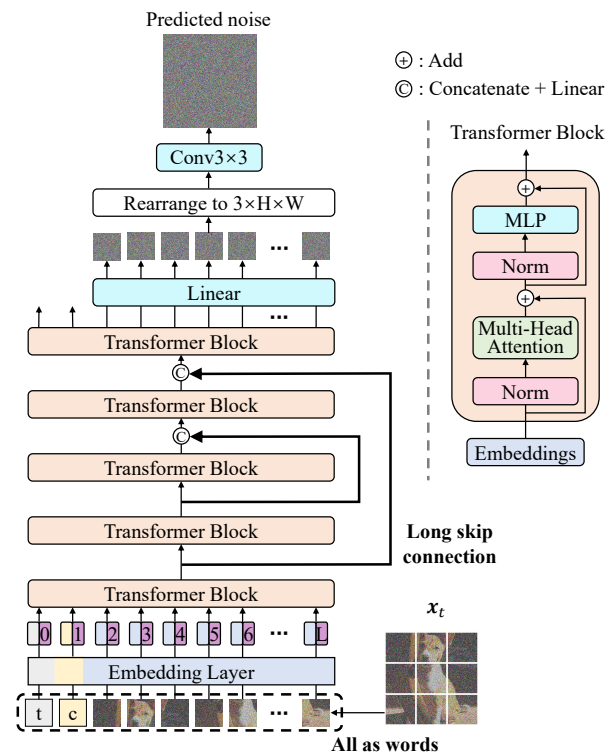


Figure 1. The U-ViT architecture for diffusion models, which is characterized by treating **all** inputs including the time, condition and noisy image patches as **tokens** and employing  $(\#Blocks-1)/2$  **long skip connections** between shallow and deep layers.

speech synthesis [6, 33], and 3D synthesis [46].

Along with the development of algorithms [2, 3, 14, 24, 32, 40, 41, 45, 57, 58, 61, 65], the revolution of backbones plays a central role in diffusion models. A representative example is U-Net based on a convolutional neural network (CNN) employed in prior work [24, 59]. The CNN-based U-

Net is characterized by a group of down-sampling blocks, a group of up-sampling blocks, and long skip connections between the two groups, which dominates diffusion models for image generation tasks [12, 47, 49, 51]. On the other hand, vision transformers (ViT) [15] have shown promise in various vision tasks, where ViT is comparable or even superior to CNN based approaches [9, 20, 35, 62, 75]. Therefore, a very natural question arises: *whether the reliance of the CNN-based U-Net is necessary in diffusion models?*

In this paper, we design a simple and general ViT-based architecture called U-ViT (Figure 1). Following the design methodology of transformers, U-ViT treats all inputs including the time, condition and noisy image patches as tokens. Crucially, U-ViT employs long skip connections between shallow and deep layers inspired by U-Net. Intuitively, low-level features are important to the pixel-level prediction objective in diffusion models and such connections can ease the training of the corresponding prediction network. Besides, U-ViT optionally adds an extra  $3 \times 3$  convolutional block before output for better visual quality. See a systematical ablation study for all elements in Figure 2.

We evaluate U-ViT in three popular tasks: unconditional image generation, class-conditional image generation and text-to-image generation. In all settings, U-ViT is comparable if not superior to a CNN-based U-Net of a similar size. In particular, latent diffusion models with U-ViT achieve record-breaking FID scores of 2.29 in class-conditional image generation on ImageNet  $256 \times 256$ , and 5.48 in text-to-image generation on MS-COCO, among methods without accessing large external datasets during the training of generative models.

Our results suggest that the long skip connection is crucial while the down/up-sampling operators in CNN-based U-Net are not always necessary for image diffusion models. We believe that U-ViT can provide insights for future research on diffusion model backbones and benefit generative modeling on large scale cross-modality datasets.

## 2. Background

**Diffusion models** [24, 56, 61] gradually inject noise to data, and then reverse this process to generate data from noise. The noise-injection process, also called the forward process, is formalized as a Markov chain:

$$q(\mathbf{x}_{1:T}|\mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1}).$$

Here  $\mathbf{x}_0$  is the data,  $q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t|\sqrt{\alpha_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I})$ , and  $\alpha_t$  and  $\beta_t$  represent the noise schedule such that  $\alpha_t + \beta_t = 1$ . To reverse this process, a Gaussian model  $p(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}|\boldsymbol{\mu}_t(\mathbf{x}_t), \sigma_t^2\mathbf{I})$  is adopted to approximate the ground truth reverse transition  $q(\mathbf{x}_{t-1}|\mathbf{x}_t)$ , and

the optimal mean [3] is

$$\boldsymbol{\mu}_t^*(\mathbf{x}_t) = \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{x}_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}} \mathbb{E}[\boldsymbol{\epsilon}|\mathbf{x}_t] \right).$$

Here  $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$ , and  $\boldsymbol{\epsilon}$  is the standard Gaussian noises injected to  $\mathbf{x}_t$ . Thus, the learning is equivalent to a noise prediction task. Formally, a noise prediction network  $\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)$  is adopted to learn  $\mathbb{E}[\boldsymbol{\epsilon}|\mathbf{x}_t]$  by minimizing a noise prediction objective, i.e.,  $\min_{\theta} \mathbb{E}_{t, \mathbf{x}_0, \boldsymbol{\epsilon}} \|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)\|_2^2$ , where  $t$  is uniform between 1 and  $T$ . To learn conditional diffusion models, e.g., class-conditional [12] or text-to-image [47] models, the condition information is further fed into the noise prediction objective:

$$\min_{\theta} \mathbb{E}_{t, \mathbf{x}_0, c, \boldsymbol{\epsilon}} \|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t, c)\|_2^2, \quad (1)$$

where  $c$  is the condition or its continuous embedding. In prior work on image modeling, the success of diffusion models heavily rely on CNN-based U-Net [50, 59], which is a convolutional backbone characterized by a group of down-sampling blocks, a group of up-sampling blocks and long skip connections between the two groups, and  $c$  is fed into U-Net by mechanisms such as adaptive group normalization [12] and cross attention [49].

**Vision Transformer (ViT)** [15] is a pure transformer architecture that treats an image as a sequence of tokens (words). ViT rearranges an image into a sequence of flattened patches. Then, ViT adds learnable 1D position embeddings to linear embeddings of these patches before feeding them into a transformer encoder [66]. ViT has shown promise in various vision tasks but it is not clear whether it is suitable for diffusion-based image modeling yet.

## 3. Method

U-ViT is a simple and general backbone for diffusion models in image generation (Figure 1). In particular, U-ViT parameterizes the noise prediction network<sup>1</sup>  $\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t, c)$  in Eq. (1). It takes the time  $t$ , the condition  $c$  and the noisy image  $\mathbf{x}_t$  as inputs and predicts the noise injected into  $\mathbf{x}_t$ . Following the design methodology of ViT, the image is split into patches, and U-ViT treats all inputs including the time, condition and image patches as tokens (words).

Inspired by the success of the CNN-based U-Net in diffusion models [59], U-ViT also employs similar long skip connections between shallow and deep layers. Intuitively, the objective in Eq. (1) is a pixel-level prediction task and is sensitive to low-level features. The long skip connections provide shortcuts for the low-level features and therefore ease the training of the noise prediction network.

<sup>1</sup>U-ViT can also parameterize other types of prediction, e.g.,  $\mathbf{x}_0$ -prediction [24].

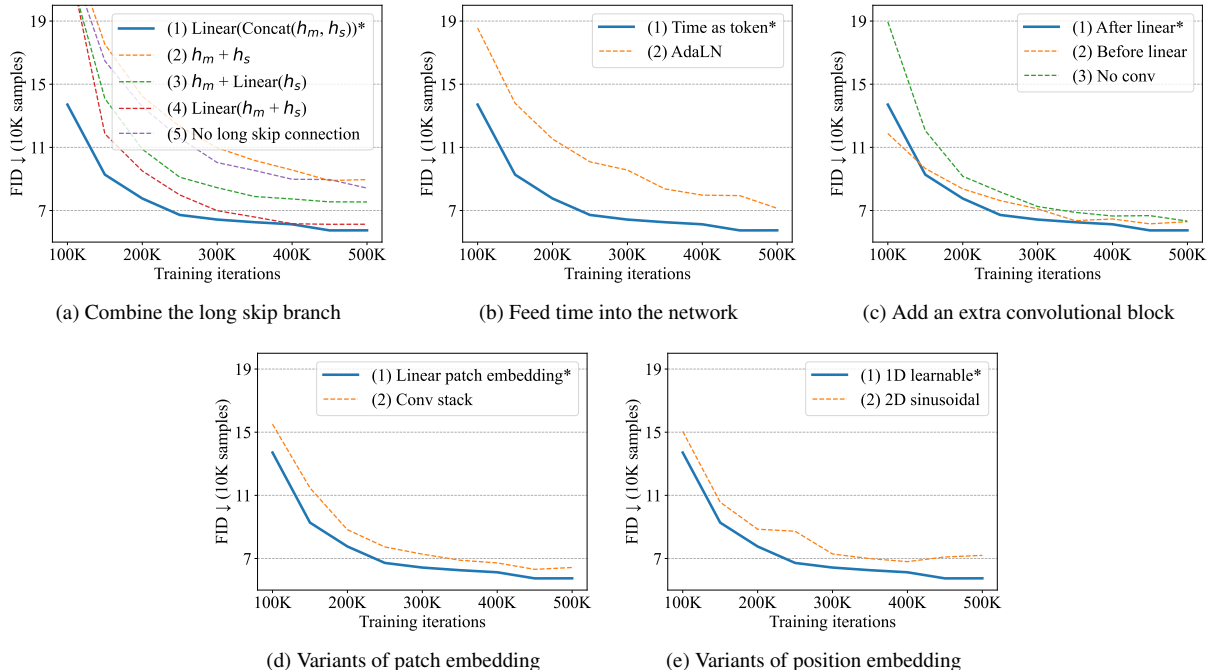


Figure 2. Ablate design choices. The one marked with \* is the adopted choice of U-ViT illustrated in Figure 1. Since this ablation aims to determine implementation details, we evaluate FID on 10K generated samples (instead of 50K samples for efficiency).

Additionally, U-ViT optionally adds a  $3 \times 3$  convolutional block before output. This is intended to prevent the potential artifacts in images produced by transformers [72]. The block improves the visual quality of the samples generated by U-ViT according to our experiments.

In Section 3.1, we present the implementation details of U-ViT. In Section 3.2, we present the scaling properties of U-ViT by studying the effect of depth, width and patch size.

### 3.1. Implementation Details

Although U-ViT is conceptually simple, we carefully design its implementation. To this end, we perform a systematical empirical study on key elements in U-ViT. In particular, we ablate on CIFAR10 [34], evaluate the FID score [22] every 50K training iterations on 10K generated samples (instead of 50K samples for efficiency), and determine default implementation details.

**The way to combine the long skip branch.** Let  $\mathbf{h}_m, \mathbf{h}_s \in \mathbb{R}^{L \times D}$  be the embeddings from the main branch and the long skip branch respectively. We consider several ways to combine them before feeding them to the next transformer block: (1) concatenating them and then performing a linear projection as illustrated in Figure 1, i.e.,  $\text{Linear}(\text{Concat}(\mathbf{h}_m, \mathbf{h}_s))$ ; (2) directly adding them, i.e.,  $\mathbf{h}_m + \mathbf{h}_s$ ; (3) performing a linear projection to  $\mathbf{h}_s$  and then adding them, i.e.,  $\mathbf{h}_m + \text{Linear}(\mathbf{h}_s)$ ; (4) adding them and then performing a linear projection, i.e.,  $\text{Linear}(\mathbf{h}_m + \mathbf{h}_s)$ . (5) We also compare with the case where the long skip

connection is dropped. As shown in Figure 2 (a), directly adding  $\mathbf{h}_m, \mathbf{h}_s$  does not provide benefits. Since a transformer block has skip connections via the adding operator inside it,  $\mathbf{h}_m$  already contains information of  $\mathbf{h}_s$  in a linear form. As a result, the only effect of  $\mathbf{h}_m + \mathbf{h}_s$  is to increase the coefficient of  $\mathbf{h}_s$  in the linear form, which does not change the nature of the network. In contrast, all other ways to combine  $\mathbf{h}_s$  perform a linear projection on  $\mathbf{h}_s$  and improve the performance compared to no long skip connection. Among them, the first way with concatenation performs the best. In Appendix E, we visualize the similarity between representations in a network, and we find the first way with concatenation significantly changes the representations, which validates its effectiveness.

**The way to feed the time into the network.** We consider two ways to feed  $t$  into the network. (1) The first way is to treat it as a token as illustrated in Figure 1. (2) The second way is to incorporate the time after the layer normalization in the transformer block [19], which is similar to the adaptive group normalization [12] used in U-Net. The second way is referred to as adaptive layer normalization (AdaLN). Formally,  $\text{AdaLN}(h, y) = y_s \text{LayerNorm}(h) + y_b$ , where  $h$  is an embedding inside a transformer block, and  $y_s, y_b$  are obtained from a linear projection of the time embedding. As shown in Figure 2 (b), while simple, the first way that treats time as a token performs better than AdaLN.

**The way to add an extra convolutional block after the transformer.** We consider two ways to add an extra convo-

lutional block after the transformer. (1) The first way is to add a  $3\times 3$  convolutional block after the linear projection that maps the token embeddings to image patches, as illustrated in Figure 1. (2) The second way is to add a  $3\times 3$  convolutional block before this linear projection, which needs to first rearrange the 1D sequence of token embeddings  $\mathbf{h} \in \mathbb{R}^{L\times D}$  to a 2D feature of shape  $H/P \times W/P \times D$ , where  $P$  is the patch size. (3) We also compare with the case where we drop the extra convolutional block. As shown in Figure 2 (c), the first way that adds a  $3\times 3$  convolutional block after the linear projection performs slightly better than other two choices.

**Variants of the patch embedding.** We consider two variants of the patch embedding. (1) The original patch embedding adopts a linear projection that maps a patch to a token embedding, as illustrated in Figure 1. (2) Alternatively, [67] use a stack of  $3\times 3$  convolutional blocks followed by a  $1\times 1$  convolutional block to map an image to token embeddings. We compare them in Figure 2 (d), and the original patch embedding performs better.

**Variants of the position embedding.** We consider two variants of the position embedding. (1) The first one is the 1-dimensional learnable position embedding proposed in the original ViT [15], which is the default setting in this paper. (2) The second one is the 2-dimensional sinusoidal position embedding, which is obtained by concatenating the sinusoidal embeddings [66] of  $i$  and  $j$  for a patch at position  $(i, j)$ . As shown in Figure 2 (e), the 1-dimensional learnable position embedding performs better. We also try not use any position embedding, and find the model fails to generate meaningful images, which implies the position information is critical in image generation.

### 3.2. Effect of Depth, Width and Patch Size

We evaluate scaling properties of U-ViT by studying the effect of the depth (i.e., the number of layers), width (i.e., the hidden size  $D$ ) and patch size on CIFAR10. We present the curve of FID during training in Appendix B. The performance improves as the depth (i.e., the number of layers) increases from 9 to 13. Nevertheless, U-ViT does not gain from a larger depth like 17 in 50K training iterations. Similarly, increasing the width (i.e., the hidden size) from 256 to 512 improves the performance, and further increase to 768 brings no gain; decreasing the patch size from 8 to 2 improves the performance, and further decrease to 1 brings no gain. Note that a small patch size like 2 is required for a good performance. We hypothesize it is because that the noise prediction task in diffusion models is low-level and requires small patches, differing from high-level tasks (e.g., classification). Since using a small patch size is costly for high resolution images, we firstly convert them to low-dimensional latent representations [49] and model these latent representations using U-ViT.

## 4. Related Work

**Transformers in diffusion models.** A related work is GenViT [70]. GenViT employs a smaller ViT that does not employ long skip connections and the  $3\times 3$  convolutional block, and incorporates time before normalization layers for image diffusion models. Empirically, our U-ViT performs much better than GenViT (see Table 1) by carefully designing implementation details. Another related work is VQ-Diffusion [19] and its variants [55, 63]. VQ-Diffusion firstly obtains a sequence of discrete image tokens via a VQ-GAN [16], and then models these tokens using a discrete diffusion model [1, 56] with a transformer as its backbone. Time and condition are fed into the transformer through cross attention or adaptive layer normalization. In contrast, our U-ViT simply treats all inputs as tokens, and employs long skip connections between shallow and deep layers, which achieves a better FID (see Table 1 and Table 4). In addition to images, transformers in diffusion models are also employed to encode texts [44, 47, 49, 51], decode texts [7, 28, 36, 43] and generate CLIP embeddings [47].

**U-Net in diffusion models.** [59, 60] initially introduce CNN-based U-Net to model the gradient of log-likelihood function for continuous image data. After that, improvements on the CNN-based U-Net for (continuous) image diffusion models are made, including using group normalization [24], multi-head attention [45], improved residual block [12] and cross attention [49]. In contrast, our U-ViT is a ViT-based backbone with conceptually simple design, and meanwhile has a comparable performance if not superior to a CNN-based U-Net of a similar size (see Table 1 and Table 4).

**Improvements of diffusion models.** In addition to the backbone, there are also improvements on other aspects, such as fast sampling [3, 41, 52, 57, 68], improved training methodology [2, 14, 29, 31, 32, 40, 45, 58, 65] and controllable generation [4, 10, 12, 21, 26, 42, 54, 74].

## 5. Experiments

We evaluate the proposed U-ViT in unconditional and class-conditional image generation (Section 5.2), as well as text-to-image generation (Section 5.3). Before presenting these results, we list main experimental setup below, and more details such as the sampling hyperparameters are provided in Appendix A.

### 5.1. Experimental Setup

**Datasets.** For unconditional learning, we consider CIFAR10 [34], which contain 50K training images, and CelebA  $64\times 64$  [38], which contain 162,770 training images of human faces. For class-conditional learning, we consider ImageNet [11] at  $64\times 64$ ,  $256\times 256$  and  $512\times 512$  resolutions, which contains 1,281,167 training images from



(a) ImageNet 512×512

(b) ImageNet 256×256

(c) CIFAR10

(d) CelebA 64×64

(e) ImageNet 64×64

Figure 3. Image generation results of U-ViT: selected samples on ImageNet 512×512 and ImageNet 256×256, and random samples on CIFAR10, CelebA 64×64, and ImageNet 64×64.

1K different classes. For text-to-image learning, we consider MS-COCO [37] at 256×256 resolution, which contains 82,783 training images and 40,504 validation images. Each image is annotated with 5 captions.

**High resolution image generation.** We follow latent diffusion models (LDM) [49] for images at 256×256 and 512×512 resolutions. We firstly convert them to latent representations at 32×32 and 64×64 resolutions respectively, using a pretrained image autoencoder provided by Stable Diffusion<sup>2</sup> [49]. Then we model these latent representations using the proposed U-ViT.

**Text-to-image learning.** On MS-COCO, we convert discrete texts to a sequence of embeddings using a CLIP text encoder following Stable Diffusion. Then these embeddings are fed into U-ViT as a sequence of tokens.

**U-ViT configurations.** We identify several configurations of U-ViT in Table 2. In the rest of the paper, we use brief notation to represent the U-ViT configuration and the input patch size (for instance, U-ViT-H/2 means the U-ViT-

Huge configuration with an input patch size of 2×2).

Model	#Layers	Hidden size $D$	MLP size	#Heads	#Params
U-ViT-Small	13	512	2048	8	44M
U-ViT-Small (Deep)	17	512	2048	8	58M
U-ViT-Mid	17	768	3072	12	131M
U-ViT-Large	21	1024	4096	16	287M
U-ViT-Huge	29	1152	4608	16	501M

Table 2. Configurations of U-ViT.

**Training.** We use the AdamW optimizer [39] with a weight decay of 0.3 for all datasets. We use a learning rate of 2e-4 for most datasets, except ImageNet 64×64 where we use 3e-4. We train 500K iterations on CIFAR10 and CelebA 64×64 with a batch size of 128. We train 300K iterations on ImageNet 64×64 and ImageNet 256×256, and 500K iterations on ImageNet 512×512, with a batch size of 1024. We train 1M iterations on MS-COCO with a batch size of 256. On ImageNet 256×256, ImageNet 512×512 and MS-COCO, we adopt classifier-free guidance [26] following [49]. We provide more details, such as the training time and how we choose hyperparameters in Appendix A.

<sup>2</sup><https://github.com/CompVis/stable-diffusion>

Model on <b>CIFAR10</b>			FID ↓
GAN			
StyleGAN2-ADA [30]			2.92
Diff. based on U-Net		#Params	
DDPM [24]		36M	3.17
IDDPM [45]		53M	2.90
DDPM++ cont. [61]		62M	2.55
EDM <sup>†</sup> [29]		56M	<b>1.97</b>
Diff. based on ViT		#Params	
GenViT [70]		11M	20.20
U-ViT-S/2		44M	3.11
Model on <b>CelebA 64×64</b>			FID ↓
GAN			
Diff. based on U-Net		#Params	
DDIM [57]		79M	3.26
Soft Truncation <sup>†</sup> [31]		62M	<b>1.90</b>
Diff. model based on ViT		#Params	
U-ViT-S/4		44M	2.87
Model on <b>ImageNet 64×64</b>			FID ↓
GAN			
BigGAN-deep [5]			4.06
StyleGAN-XL [53]			1.51
Diff. based on U-Net		#Params	
IDDPM (small) [45]		100M	6.92
IDDPM (large) [45]		270M	2.92
CDM [25]		Unknown	1.48
ADM [12]		296M	2.07
EDM <sup>†</sup> [29]		296M	<b>1.36</b>
Diff. based on ViT		#Params	
U-ViT-M/4		131M	5.85
U-ViT-L/4		287M	4.26
Model on <b>ImageNet 256×256</b>			FID ↓
GAN			
BigGAN-deep [5]			6.95
StyleGAN-XL [53]			2.30
Discrete diff. based on transformer			
VQ-Diffusion [19]			11.89
VQ-Diffusion (acc0.05) [19]			5.32
Diff. based on U-Net		#Params	
IDDPM [45]		270M + 280M (SR)	12.26
CDM [25]		Unknown	4.88
ADM [12]		554M	10.94
ADM-U [12]		296M + 312M (SR)	7.49
ADM-G [12]		554M + 54M (Cls)	4.59
ADM-G, ADM-U [12]		296M + 65M (Cls) + 312M (SR)	3.94
LDM <sup>‡</sup> [49]		400M + 55M (AE)	3.60
Diff. based on ViT		#Params	
U-ViT-H/2 <sup>‡</sup>		501M + 84M (AE)	<b>2.29</b>
Model on <b>ImageNet 512×512</b>			FID ↓
GAN			
BigGAN-deep [5]			8.43
StyleGAN-XL [53]			2.41
Diff. based on U-Net		#Params	
ADM [12]		559M	23.24
ADM-U [12]		422M + 309M (SR)	9.96
ADM-G [12]		559M + 54M (Cls)	7.72
ADM-G, ADM-U [12]		422M + 43M (Cls) + 309M (SR)	<b>3.85</b>
Diff. based on ViT		#Params	
U-ViT-H/4 <sup>‡</sup>		501M + 84M (AE)	4.05

Table 1. FID results of unconditional image generation on CIFAR10 and CelebA 64×64, and class-conditional image generation on ImageNet 64×64, 256×256 and 512×512. We mark the best results *among diffusion models* in **bold**. We also include GAN results (gray) for completeness. Methods marked with <sup>†</sup> use advanced training techniques for diffusion models. Methods marked with <sup>‡</sup> model latent representations of images [49] and use classifier-free guidance [26]. We also present the number of parameters of auxiliary components for diffusion models, where SR represents a super-resolution module, AE represents an image autoencoder, and Cls represents a classifier.

	4	5	10	15	20
LDM (trained 178K)	34.48	12.73	4.51	3.87	3.68
U-ViT-H/2 (trained 200K)	16.48	4.94	3.87	3.54	2.91
U-ViT-H/2 (trained 500K)	15.44	4.64	3.18	2.92	2.53

Table 3. FID results on ImageNet 256×256 under different number of sampling steps using the DPM-Solver sampler [41].

## 5.2. Unconditional and Class-Conditional Image Generation

We compare U-ViT with prior diffusion models based on U-Net. We also compare with GenViT [70], a smaller ViT

which does not employ long skip connections, and incorporates time before normalization layers. Consistent with previous literature, we report the FID score [22] on 50K generated samples to measure the image quality.

As shown in Table 1, U-ViT is comparable to U-Net on unconditional CIFAR10 and CelebA 64×64, and meanwhile performs much better than GenViT.

On class-conditional ImageNet 64×64, we initially try the U-ViT-M configuration with 131M parameters. As shown in Table 1, it gets a FID of 5.85, which is better than 6.92 of IDDPM that employs a U-Net with 100M parameters. To further improve the performance, we employ the U-ViT-L configuration with 287M parameters, and the FID improves from 5.85 to 4.26.

Model	FID	Type	Training datasets	#Params
Generative model trained on external large dataset (zero-shot)				
DALL-E [48]	~ 28	Autoregressive	DALL-E dataset (250M)	12B
CogView [13]	27.1	Autoregressive	Internal dataset (30M)	4B
LAFITE [76]	26.94	GAN	CC3M (3M)	75M + 151M (TE)
GLIDE [44]	12.24	Diffusion	DALL-E dataset (250M)	3.5B + 1.5B (SR)
Make-A-Scene [18]	11.84	Autoregressive	Union datasets (without MS-COCO) (35M)	4B
DALL-E 2 [47]	10.39	Diffusion	DALL-E dataset (250M)	4.5B + 700M (SR)
Imagen [51]	7.27	Diffusion	Internal dataset (460M) + LAION (400M)	2B + 4.6B (TE) + 600M (SR)
Parti [71]	7.23	Autoregressive	LAION (400M) + FIT (400M) + JFT (4B)	20B + 630M (AE)
Re-Imagen [8]	<b>6.88</b>	Diffusion	KNN-ImageText (50M)	2.5B + 750M (SR)
Generative model trained on external large dataset with access to MS-COCO				
VQ-Diffusion <sup>†</sup> [19]	13.86	Discrete diffusion	Conceptual Caption Subset (7M)	370M
Make-A-Scene [18]	7.55	Autoregressive	Union datasets (with MS-COCO) (35M)	4B
Re-Imagen <sup>‡</sup> [8]	5.25	Diffusion	KNN-ImageText (50M)	2.5B + 750M (SR)
Parti <sup>†</sup> [71]	<b>3.22</b>	Autoregressive	LAION (400M) + FIT (400M) + JFT (4B)	20B + 630M (AE)
Generative model trained on MS-COCO				
AttnGAN [69]	35.49	GAN	MS-COCO (83K)	230M
DM-GAN [77]	32.64	GAN	MS-COCO (83K)	46M
VQ-Diffusion [19]	19.75	Discrete diffusion	MS-COCO (83K)	370M
DF-GAN [64]	19.32	GAN	MS-COCO (83K)	19M
XMC-GAN [73]	9.33	GAN	MS-COCO (83K)	166M
Fririo [17]	8.97	Diffusion	MS-COCO (83K)	512M + 186M (TE) + 68M (AE)
LAFITE [76]	8.12	GAN	MS-COCO (83K)	75M + 151M (TE)
U-Net*	7.32	Latent diffusion	MS-COCO (83K)	53M + 123M (TE) + 84M (AE)
U-ViT-S/2	5.95	Latent diffusion	MS-COCO (83K)	45M + 123M (TE) + 84M (AE)
U-ViT-S/2 (Deep)	<b>5.48</b>	Latent diffusion	MS-COCO (83K)	58M + 123M (TE) + 84M (AE)

Table 4. FID results of different models on MS-COCO validation ( $256 \times 256$ ). U-ViT-S (Deep) increases the number of layers from 13 to 17 compared to U-ViT-S. We also present the number of parameters of auxiliary components for a model when it is reported in the corresponding paper, where SR represents a super-resolution module, AE represents an image autoencoder, and TE represents a text encoder. Methods marked with <sup>†</sup> finetune on MS-COCO. Methods marked with <sup>‡</sup> use MS-COCO as a knowledge base for retrieval. The U-Net\* is trained by ourselves to serve as a direct baseline of U-ViT, where we leave other parts unchanged except for the backbone.




Figure 4. Ablate the long skip connection on ImageNet  $256 \times 256$  (w/o classifier-free guidance).

Meanwhile, we find that our U-ViT performs especially well in the latent space [49], where images are firstly converted to their latent representations before applying diffusion models. On class-conditional ImageNet  $256 \times 256$ , our U-ViT obtains a state-of-the-art FID of 2.29, which outperforms all prior diffusion models. Table 3 further demon-

strates that our U-ViT outperforms LDM under different number of sampling steps using the same sampler. Note that our U-ViT also outperforms VQ-Diffusion, which is a discrete diffusion model [1] that employs a transformer as its backbone. We also try replace our U-ViT with a U-Net with a similar amount of parameters and computational cost, where our U-ViT still outperforms U-Net (see details in Appendix F). On class-conditional ImageNet  $512 \times 512$ , our U-ViT outperforms ADM-G that directly models the pixels of images. In Figure 3, we provide selected samples on ImageNet  $256 \times 256$  and ImageNet  $512 \times 512$ , and random samples on other datasets, which have good quality and clear semantics. We provide more generated samples including class-conditional and random ones in Appendix G.

In Section 3.1 we have demonstrated the importance of long skip connection on small-scale dataset (i.e., CIFAR10). Figure 4 further shows it is also critical for large-scale dataset such as ImageNet.

In Appendix D, we present results of other metrics (e.g., sFID, inception score, precision and recall) as well as the computational cost (GFLOPs) with more U-ViT configura-



Figure 5. Text-to-image generation on MS-COCO. All the other settings except the backbone are the same. U-Net and U-ViT generate samples using the same random seed for a fair comparison. The random seed is unselected.

tions on ImageNet. Our U-ViT is still comparable to state-of-the-art diffusion models on other metrics, and meanwhile has comparable if not smaller GFLOPs.

### 5.3. Text-to-Image Generation on MS-COCO

We evaluate U-ViT for text-to-image generation on the standard benchmark dataset MS-COCO. We train our U-ViT in the latent space of images [49] as detailed in Section 5.1. We also train another latent diffusion model that employs a U-Net of comparable model size to U-ViT-S, and leave other parts unchanged. Its hyperparameters and training details are provided in Appendix C. We report the FID score [22] to measure the image quality. Consistent with previous literature, we randomly draw 30K prompts from the MS-COCO validation set, and generate samples on these prompts to compute FID.

As shown in Table 4, our U-ViT-S already achieves a state-of-the-art FID among methods without accessing large external datasets during the training of generative models. By further increasing the number of layers from 13 to 17, our U-ViT-S (Deep) can even achieve a better FID of 5.48. Figure 5 shows generated samples of U-Net and U-ViT using the same random seed for a fair comparison. We find U-ViT generates more high quality samples, and meanwhile the semantics matches the text better. For example, given the text ‘‘a baseball player swinging a bat at a ball’’, U-Net generates neither the bat nor the ball. In contrast, our U-ViT-S generates the ball with even a smaller number of parameters, and our U-ViT-S (Deep) further generates the bat. We hypothesize this is because texts and images interact at every layer in our U-ViT, which is

more frequent than U-Net that only interact at cross attention layer. We provide more samples in Appendix G.

## 6. Conclusion

This work presents U-ViT, a simple and general ViT-based architecture for image generation with diffusion models. U-ViT treats all inputs including the time, condition and noisy image patches as tokens and employs long skip connections between shallow and deep layers. We evaluate U-ViT in tasks including unconditional and class-conditional image generation, as well as text-to-image generation. Experiments demonstrate U-ViT is comparable if not superior to a CNN-based U-Net of a similar size. These results suggest that, for diffusion-based image modeling, the long skip connection is crucial while the down/up-sampling operators in CNN-based U-Net are not always necessary. We believe that U-ViT can provide insights for future research on backbones in diffusion models and benefit generative modeling on large scale cross-modality datasets.

**Acknowledgments.** This work was supported by NSF of China Projects (Nos. 62061136001, 61620106010, 62076145, U19B2034, U1811461, U19A2081, 6197222); Beijing Outstanding Young Scientist Program NO. BJJWZYJH012019100020098; a grant from Tsinghua Institute for Guo Qiang; the High Performance Computing Center, Tsinghua University; the Fundamental Research Funds for the Central Universities, and the Research Funds of Renmin University of China (22XNKJ13). C. Li was also sponsored by Beijing Nova Program. J.Z was also supported by the XPlorer Prize.



## References

- [1] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. In *Advances in Neural Information Processing Systems*, 2021. 4, 7
- [2] Fan Bao, Chongxuan Li, Jiacheng Sun, Jun Zhu, and Bo Zhang. Estimating the optimal covariance with imperfect mean in diffusion probabilistic models. *ArXiv preprint*, 2022. 1, 4
- [3] Fan Bao, Chongxuan Li, Jun Zhu, and Bo Zhang. Analytic-dpm: an analytic estimate of the optimal reverse variance in diffusion probabilistic models. *ArXiv preprint*, 2022. 1, 2, 4
- [4] Fan Bao, Min Zhao, Zhongkai Hao, Peiyao Li, Chongxuan Li, and Jun Zhu. Equivariant energy-guided sde for inverse molecular design. *ArXiv preprint*, 2022. 4
- [5] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *7th International Conference on Learning Representations*, 2019. 6
- [6] Nanxin Chen, Yu Zhang, Heiga Zen, Ron J. Weiss, Mohammad Norouzi, and William Chan. Wavegrad: Estimating gradients for waveform generation. In *9th International Conference on Learning Representations*, 2021. 1
- [7] Ting Chen, Ruixiang Zhang, and Geoffrey Hinton. Analog bits: Generating discrete data using diffusion models with self-conditioning. *ArXiv preprint*, 2022. 4
- [8] Wenhu Chen, Hexiang Hu, Chitwan Saharia, and William W Cohen. Re-imagen: Retrieval-augmented text-to-image generator. *ArXiv preprint*, 2022. 7
- [9] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 2
- [10] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. *ArXiv preprint*, 2021. 1, 4
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. Imagenet: A large-scale hierarchical image database. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2009. 4
- [12] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis. *ArXiv preprint*, 2021. 1, 2, 3, 4, 6
- [13] Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, et al. Cogview: Mastering text-to-image generation via transformers. In *Advances in Neural Information Processing Systems*, 2021. 7
- [14] Tim Dockhorn, Arash Vahdat, and Karsten Kreis. Score-based generative modeling with critically-damped langevin diffusion. *ArXiv preprint*, 2021. 1, 4
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations*, 2021. 2, 4
- [16] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021. 4
- [17] Wan-Cyuan Fan, Yen-Chun Chen, DongDong Chen, Yu Cheng, Lu Yuan, and Yu-Chiang Frank Wang. Frido: Feature pyramid diffusion for complex scene image synthesis. *ArXiv preprint*, 2022. 7
- [18] Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. Make-a-scene: Scene-based text-to-image generation with human priors. *ArXiv preprint*, 2022. 7
- [19] Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 3, 4, 6, 7
- [20] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [21] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *ArXiv preprint*, 2022. 4
- [22] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, 2017. 3, 6, 8
- [23] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *ArXiv preprint*, 2022. 1
- [24] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, 2020. 1, 2, 4, 6
- [25] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *J. Mach. Learn. Res.*, 2022. 1, 6
- [26] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *ArXiv preprint*, 2022. 4, 5, 6
- [27] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *ArXiv preprint*, 2022. 1
- [28] Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. In *Advances in Neural Information Processing Systems*, 2021. 4
- [29] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *ArXiv preprint*, 2022. 4, 6
- [30] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative ad-

- versarial networks with limited data. In *Advances in Neural Information Processing Systems*, 2020. 6
- [31] Dongjun Kim, Seungjae Shin, Kyungwoo Song, Wanmo Kang, and Il-Chul Moon. Soft truncation: A universal training technique of score-based diffusion model for high precision score estimation. In *International Conference on Machine Learning*, 2022. 4, 6
- [32] Diederik P Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *ArXiv preprint*, 2021. 1, 4
- [33] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. In *9th International Conference on Learning Representations*, 2021. 1
- [34] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 3, 4
- [35] Kwonjoon Lee, Huiwen Chang, Lu Jiang, Han Zhang, Zhuowen Tu, and Ce Liu. Vitgan: Training gans with vision transformers. *ArXiv preprint*, 2021. 2
- [36] Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B Hashimoto. Diffusion-lm improves controllable text generation. *ArXiv preprint*, 2022. 4
- [37] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, 2014. 5
- [38] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *IEEE International Conference on Computer Vision*, 2015. 4
- [39] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *7th International Conference on Learning Representations*, 2019. 5
- [40] Cheng Lu, Kaiwen Zheng, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Maximum likelihood training for score-based diffusion odes by high order denoising score matching. In *International Conference on Machine Learning*, 2022. 1, 4
- [41] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *ArXiv preprint*, 2022. 1, 4, 6
- [42] Chenlin Meng, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Image synthesis and editing with stochastic differential equations. *ArXiv preprint*, 2021. 1, 4
- [43] Eliya Nachmani and Shaked Dovrat. Zero-shot translation using diffusion models. *ArXiv preprint*, 2021. 4
- [44] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *ArXiv preprint*, 2021. 4, 7
- [45] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *Proceedings of the 38th International Conference on Machine Learning*, 2021. 1, 4, 6
- [46] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *ArXiv preprint*, 2022. 1
- [47] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *ArXiv preprint*, 2022. 1, 2, 4, 7
- [48] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *Proceedings of the 38th International Conference on Machine Learning*, 2021. 7
- [49] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1, 2, 4, 5, 6, 7, 8
- [50] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 2015. 2
- [51] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *ArXiv preprint*, 2022. 1, 2, 4, 7
- [52] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *ArXiv preprint*, 2022. 4
- [53] Axel Sauer, Katja Schwarz, and Andreas Geiger. Stylegan-xl: Scaling stylegan to large diverse datasets. In *ACM SIGGRAPH 2022 Conference Proceedings*, 2022. 6
- [54] Vikash Sehwal, Caner Hazirbas, Albert Gordo, Firat Ozgenel, and Cristian Canton. Generating high fidelity data from low-density regions using diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 4
- [55] Shelly Sheynin, Oron Ashual, Adam Polyak, Uriel Singer, Oran Gafni, Eliya Nachmani, and Yaniv Taigman. Knn-diffusion: Image generation via large-scale retrieval. *ArXiv preprint*, 2022. 4
- [56] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning*, 2015. 1, 2, 4
- [57] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *9th International Conference on Learning Representations*, 2021. 1, 4, 6
- [58] Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *arXiv e-prints*, 2021. 1, 4
- [59] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems*, 2019. 1, 2, 4
- [60] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. In *Advances in Neural Information Processing Systems*, 2020. 4
- [61] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based

- generative modeling through stochastic differential equations. In *9th International Conference on Learning Representations*, 2021. 1, 2, 6
- [62] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 2
- [63] Zhicong Tang, Shuyang Gu, Jianmin Bao, Dong Chen, and Fang Wen. Improved vector quantized diffusion models. *ArXiv preprint*, 2022. 4
- [64] Ming Tao, Hao Tang, Fei Wu, Xiao-Yuan Jing, Bing-Kun Bao, and Changsheng Xu. Df-gan: A simple and effective baseline for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 7
- [65] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. *ArXiv preprint*, 2021. 1, 4
- [66] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017. 2, 4
- [67] Tete Xiao, Mannat Singh, Eric Mintun, Trevor Darrell, Piotr Dollár, and Ross Girshick. Early convolutions help transformers see better. In *Advances in Neural Information Processing Systems*, 2021. 4
- [68] Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with denoising diffusion gans. *ArXiv preprint*, 2021. 4
- [69] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 7
- [70] Xiulong Yang, Sheng-Min Shih, Yinlin Fu, Xiaoting Zhao, and Shihao Ji. Your vit is secretly a hybrid discriminative-generative diffusion model. *ArXiv preprint*, 2022. 4, 6
- [71] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gungjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-rich text-to-image generation. *ArXiv preprint*, 2022. 7
- [72] Bowen Zhang, Shuyang Gu, Bo Zhang, Jianmin Bao, Dong Chen, Fang Wen, Yong Wang, and Baining Guo. Styleswin: Transformer-based gan for high-resolution image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 3
- [73] Han Zhang, Jing Yu Koh, Jason Baldridge, Honglak Lee, and Yinfei Yang. Cross-modal contrastive learning for text-to-image generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021. 7
- [74] Min Zhao, Fan Bao, Chongxuan Li, and Jun Zhu. Egsde: Unpaired image-to-image translation via energy-guided stochastic differential equations. *ArXiv preprint*, 2022. 1, 4
- [75] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021. 2
- [76] Yufan Zhou, Ruiyi Zhang, Changyou Chen, Chunyuan Li, Chris Tensmeyer, Tong Yu, Jiuxiang Gu, Jinhui Xu, and Tong Sun. Lafite: Towards language-free training for text-to-image generation. *ArXiv preprint*, 2021. 7
- [77] Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. DM-GAN: dynamic memory generative adversarial networks for text-to-image synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 7