Improving Visual Quality of Image Synthesis by A Token-based Generator with Transformers

Anonymous Author(s)
Affiliation
Address
email

Abstract

We present a new perspective of achieving image synthesis by viewing this task as a visual token generation problem. Different from existing paradigms that directly synthesize a full image from a single input (e.g., a latent code), the new formulation enables a local content-dependent manipulation of each visual token, leading to fine-grained style learning for image synthesis. Specifically, it takes a sequence of latent tokens as input to predict the visual tokens for synthesizing an image. Under this perspective, we propose a token-based generator (i.e.,TokenGAN) for the visual token generation. Particularly, the TokenGAN inputs two semantically different visual tokens, i.e., the learned constant content tokens and the style tokens from the latent space. Given a sequence of style tokens, the TokenGAN learns to control the image synthesis by assigning the styles to semantic-related content tokens by attention mechanism with a Transformer. We conduct extensive experiments and show that the proposed TokenGAN has achieved state-of-the-art results on several widely-used image synthesis benchmarks, including FFHQ and LSUN CHURCH with different resolutions. In particular, it is able to synthesize high-fidelity images with 1024×1024 image size, dispensing with convolutions entirely.

1 Introduction

Unconditional image synthesis aims to generate an image from a latent space by adversarial training [10, 16, 19, 24, 26, 31, 46]. Recent advances in this domain have been achieved by a style-based generator architecture in terms of both the visual quality and resolution of generated images [25, 26, 27, 35, 47]. In particular, the style-based generator has been widely used in many other generative tasks, including facial editing [9, 38], style transfer [1, 36], image super-resolution [17, 29], and image inpainting [2, 47].

The key to the success of the style-based generator lies in the learning of the style control based on the intermediate latent space $\mathcal{W}$ [26, 27]. Instead of feeding the input latent code $\mathbf{z} \in \mathcal{Z}$ through the input layer only (Figure1-a), the style-based generator maps the input $\mathbf{z}$ to an intermediate latent space $\mathbf{w} \in \mathcal{W}$, which then controls the "style" of the image at each layer via adaptive instance normalization (AdaIN [20]) (Figure 1-b). It has been demonstrated that such a design allows a less entangled representation learning in $\mathcal{W}$, leading to better generative image modeling [12, 22, 26, 38].

Despite the promising results, the style-based generator can suffer from the style control via AdaIN operation [27, 47]. Specifically, the style code $\mathbf{w}$ affects the entire image by scaling and biasing complete feature maps with a single value via AdaIN operation [33, 42, 49]. Such an imposed single style over multiple image regions can inevitably result in entangled representation of different image variations (e.g., hairstyle, facial expression) [27, 42, 47]. This limitation for image modeling can even lead to visible artifacts in plenty of synthesized results (e.g., droplet artifacts [26]).

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

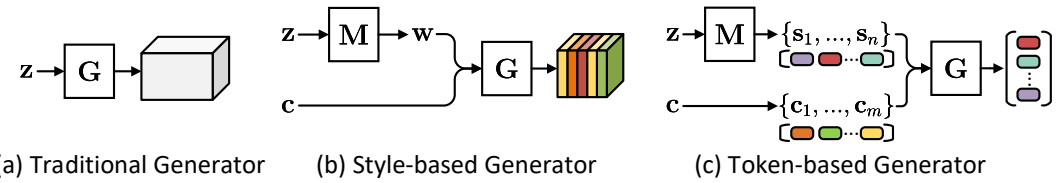

Figure 1: Overview of different generators. **(a)** A traditional generator [16] feeds a single latent $\mathbf{z}$ as input to control the image synthesis. **(b)** A style-based generator [26] maps $\mathbf{z}$ to an intermediate latent space $\mathbf{w} \in \mathcal{W}$ to control the styles of the content $\mathbf{c}$ via AdaIN [20]. **(c)** Our token-based generator starts from a sequence of content tokens $\{\mathbf{c}_1, \cdots, \mathbf{c}_m\}$ and controls each content token with a set of style tokens $\{\mathbf{s}_1, \cdots, \mathbf{s}_n\} \in \mathcal{S}$ by attention mechanism with a visual Transformer. $\mathbf{M}$ denotes the mapping network and $\mathbf{G}$ denotes the generator network.

To address the above problems, in this paper we propose a new perspective that views image synthesis as a visual token generation problem. The visual token is a popular representation of an image patch with a predefined size and position [6, 7, 13]; and has shown an increasing impact in various tasks due to its fine-grained controllability [6, 13, 41, 43]. Inspired by the appealing property of the visual token, we propose to achieve image synthesis by fine-grained visual token generation. Specifically, it takes as input a sequence of latent tokens to predict the visual tokens of an image. Such a token-based formulation enables fine-grained control for each token, which makes it possible to control the style of a certain region in image while maintain other parts.

Under this new paradigm, we design a token-based generator, i.e., TokenGAN, for the visual token generation problem. We further consider two different types of input tokens in the generator, i.e., the content tokens and the style tokens. Specifically, the content tokens are learned as the constant input in the generator network while the style tokens are projected from a learned intermediate latent space (Figure1-c). Given a sequence of style tokens, the TokenGAN aims at controlling the generated visual tokens by assigning each content token with related styles according to their semantics. In particular, since the Transformer has been verified to be effective in sequence modeling in a broad range of tasks [6, 41], we adopt a generator network architecture from a visual Transformer to model the relations between content tokens and style tokens. Through such a pair-wise interaction modeling by the attention mechanism i Transformer, the styles for the semantic-related tokens can be assigned and control the corresponding tokens, leading to a fine-grained learning of the style control for the visual token generation.

We conduct both quantitative comparisons and qualitative analysis on several unconditional image generation benchmarks. The results show that the token-based generator has achieved comparable results to the state-of-the-art in image synthesis. We summarize our contributions as below:

- We propose a new perspective of achieving image synthesis by a visual token generation. Such a formulation fully exploits the fine-grained controllability of visual tokens, leading to a better image modeling.

- We propose a token-based generator for the visual token generation. Specifically, the token-based generator decompose the visual tokens to the content and the styles, and it enables a fine-grained style control for the visual tokens by a Transformer.

- We show extensive experiments to verify the effectiveness of the token-based generator. Specifically, the token-based generator is able to synthesize high-fidelity 1024×1024 images without any convolutions in the generator.

We organize the main body of the paper as follows. We review the related work of this work in Section 2. Then, we describe the proposed visual token generation task as well as the TokenGANin Section 3. Finally, we conduct quantitative and qualitative experiments to show the effectiveness of TokenGANin Section 4.

2 Related works

2.1 Style-based generator

The distinguishing feature of the style-based generator is its unconventional generator architecture [25, 26, 27, 47]. Typically, the style-based generator consists of a *mapping network* f and a *synthesis network* g. The mapping network is used to transform the input latent code $\mathbf{z}$ to an intermediate latent code $\mathbf{w} \in \mathcal{W}$ for learning a less entangled latent space. The synthesis network follows a progressive growing design that is able to first output low-resolution images that are not affected significantly by high-resolution layers [24, 26]. Such an architecture design can lead to an automatic, unsupervised separation of high-level attributes in different layers. For example, when trained on human faces, the low-resolution layers control coarse styles (e.g., pose, identity) of images and high-resolution layers control fine styles (e.g., micro-structure, color scheme) [24].

To control the attributes of the image, the style-based generator produces *styles* from the intermediate latent code $\mathbf{w}$ by affine transforms, which then control each layer of the synthesis network via adaptive instance normalization (AdaIN) [20, 33]. The AdaIN operation removes styles from previous layers by normalizing each feature map to zero mean and unit deviation, and it assigns the new *styles* by scaling and biasing the complete normalized feature maps. However, it has been demonstrated such an AdaIN operation can destroy the magnitude information of features and thus results in the well-known droplet artifacts [27]. Besides, the scale-specific style disentanglement in StyleGAN can be limited due to the complex semantic components in each feature map [27, 47]. In this paper, we explore a new architecture by flattening the image as a sequence of tokens, which enables fine-grained control of the image by assigning token-wise semantic-aware styles based on the attention mechanism.

2.2 Transformer in vision

The Transformer typically takes as input a sequence of vectors, called tokens [41]. The Transformer is built for sequence modeling solely on attention mechanisms over the tokens, dispensing with recurrence and convolutions entirely [5, 11, 41]. Due to its great success in the field of natural language processing, an increasing number of works attempt to extend Transformer for computer vision tasks [6, 7, 8, 13, 18, 34, 50]. For example, iGPT shows promising results in image classification by pre-training a sequence Transformer with the task of auto-regressive next pixel prediction and masked pixel prediction [8]. Dosoyvitskiy et al. propose to split an image into patches and feed these patches into a standard Transformer (ViT) for image classification [13]. They show that a ViT with large-scale training can trump CNNs equipped with inductive bias. DETR is a seminal work that views object detection as a direct set prediction problem, eliminating the need for hand-crafted components and achieving impressive performance by Transformer [6]. To mitigate the issue of high computation complexity associated with long sequences caused by high-resolution images, a branch of works explores lightweight Transformer (e.g.., Deformable DETR) for vision tasks [50].

Transformer is also attracting increasing attention in low-level vision tasks. Parmar et al. cast image generation as an autoregressive sequence generation problem and propose Image Transformer with a local self-attention mechanism [14, 32, 34, 40]. However, Image Transformer can suffer from its quadratic computation cost and its long inference time due to auto-regressive prediction. Most existing works adopt a hybrid CNN-Transformer architecture, which consists of a CNN head for feature extraction, a Transformer encoder-decoder backbone, and a CNN tail for feature decoding [7, 43, 45]. For example, IPT fully utilizes the Transformer architecture by a large-scale pre-training and achieves promising results in several image restoration tasks [7]. Our token-based generator inherits the network architecture from Transformer without any convolution layers and is able to yield surprising promising results for high-resolution image synthesis.

3 Approach

In this section, we introduce the details of the proposed token-based generator(TokenGAN). As depicted in Figure 2, the TokenGAN takes as input two kinds of latent tokens to generate the visual tokens for the image by a Transformer. We introduce the input tokens in Section 3.1 and the token-based generator architecture in Section 3.2, following the overall optimization in Section 3.3.

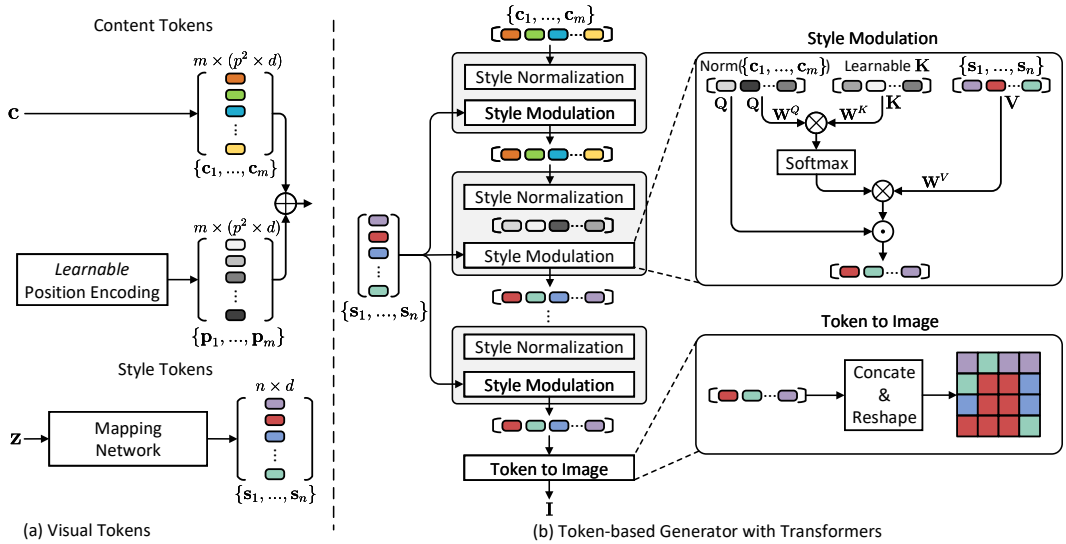

Figure 2: The overview of the TokenGAN for the visual token generation task for image synthesis. The TokenGAN takes as input two kinds of visual tokens, i.e., the style tokens and the content tokens, to generate visual tokens of the image. Specifically, the style of each content token is controlled by related style tokens by attention mechanism with a Transformer in a fine-grained way.

3.1 Visual tokens

It's intuitive that a visual token typically contains the information of both content (i.e., semantics) and related styles. In our token-based generator, we would like to disentangle these two information by a sequence of content tokens and a sequence of style tokens, so that the interactions between them can be modeled and controlled for the image synthesis. We introduce the details of each as below.

Content tokens. As depicted in Figure 2-(a), the TokenGAN starts from the sequence of learned constant content tokens $\mathbf{c} \in \mathbb{R}^{m \times (p^2 \times d)}$ through the input layer of the TokenGAN. Specifically, d is the dimension in terms of channels, $(p; p)$ is the patch size of the content token, and m is the length of the content token sequence. To maintain the position information of each token, we add learnable position encodings $\mathbf{p} \in \mathbb{R}^{m \times d}$ to each content token following a commonly used paradigm [7, 13]. For simplicity, we rewrite the notation for the content tokens as $\{\mathbf{c}_1, \cdots, \mathbf{c}_m\}$, where $\mathbf{c}_i \in \mathbb{R}^d$.

Style tokens. Given a latent code $\mathbf{z}$ in the input latent space $\mathcal{Z}$, the mapping network adopts several MLPs to map the input $\mathbf{z}$ to set of different style tokens $\{\mathbf{s}_1, \cdots, \mathbf{s}_n\} \in \mathcal{S}$, where $\mathbf{s}_i \in \mathbb{R}^d$. The style tokens control the content tokens in each layer of the TokenGANby the attention mechanism with a Transformer.

3.2 Token-based generator

As shown in Figure 2, the TokenGAN inherits the network architecture from the Transformer, which enables modeling pairwise interactions between the styles and the contents based on attention mechanism. After calculating the new style by the Transformer, we impose the styles to the content tokens by a paradigm of style "normalization-and-modulation" following previous works [26]. Specifically, the style normalization aims at removing the styles from previous layers by relieving the dependence on the original statistics of input token features. At the same time, the style modulation focuses on assigning the new styles in the current layer to the input content tokens. In particular, the new styles in our token-based generator are calculated based on the sequence of style tokens via attention mechanism. We introduce more details about the style normalization, semantic-aware style, and style modulation in our generator as below.

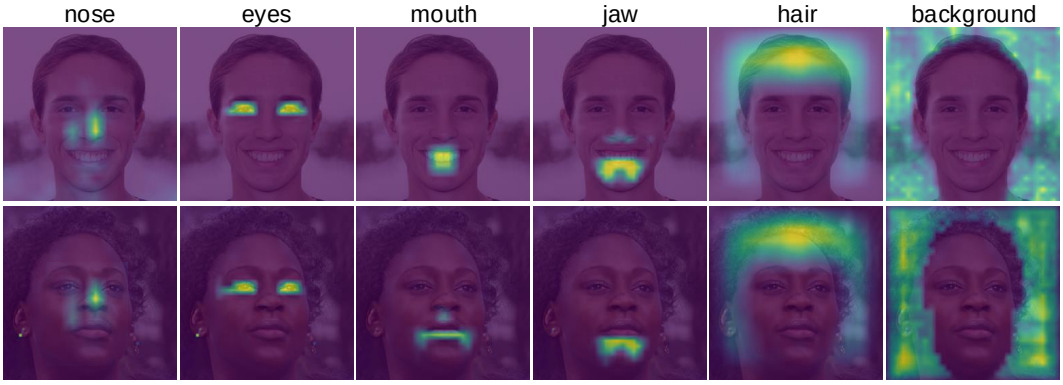

Figure 3: Visualization of the attention maps obtained by Eq. (2). Each column highlights the image regions that are affected most by the corresponding semantic-aware style. It shows that TokenGAN is able to attend to meaningful semantics for different persons in an unsupervised way. The tags are associated by human based on the attention responses for clear presentation.

Style normalization. We apply style normalization on the input content tokens to remove the styles from previous layers, which is denoted as:

$$\mathrm{Norm}(\mathbf{c}_i) = (\mathbf{c}_i - \mu(\mathbf{c}_i)) \,/\, \sigma(\mathbf{c}_i), \quad i = 1, \cdots, m, \tag{1}$$

where each $\mathbf{c}_i \in \mathbb{R}^d$ denotes a content token, d denotes the dimension of each content token, and m is the length of the token sequence. Specifically, the style-based generator typically removes the styles from previous layers by *Instance Normalization*, which can destroy the information conveyed by the magnitude of the feature map relative to each other [26, 27, 39]. To avoid the above issue, we adopt *LayerNorm* to relive the style dependence while maintaining the information of the relative magnitude following previous works [27, 28, 41].

Semantic-aware style. After style normalization, the token-based generator calculates the new styles to control the input content tokens. In practice, given a sequence of content tokens $\{\mathbf{c}_1, \cdots, \mathbf{c}_m\}$ from the input and a sequence of style tokens from the mapping network $\{\mathbf{s}_1, \cdots, \mathbf{s}_n\} \in \mathcal{S}$, the new styles for each content token are calculated by attention mechanism:

$$\mathbf{S}' = Attention(\mathbf{C}, \mathbf{K}, \mathbf{S}) = softmax(\frac{\mathbf{C}\mathbf{K}^T}{\sqrt{d}})\mathbf{S}, \tag{2}$$

where $\mathbf{S}' \in \mathbb{R}^{m \times d}$ denotes the new styles for all the content tokens, the set of content tokens are packed together into a matrix $\mathbf{C} \in \mathbb{R}^{m \times d}$, the style tokens are also packed together into matrix $\mathbf{S} \in \mathbb{R}^{n \times d}$, the matrix $\mathbf{K} \in \mathbb{R}^{n \times d}$ is the learnable semantic embedding in where each row vector indicates a kind of learned semantic for corresponding styles and contents. We oimit the linear projection in the standard attention calculation for simplicity [41]. Through such an attention mechanism, for each token, related styles are fetched and summed with weights to obtain the final styles used for style modulation. As shown in Figure 3, the learned embedding of $\mathbf{K}$ attends to meaningful image regions.

Style modulation. Different from the traditional Transformer that adds the attention results back to the input features as residual features, our attention results are used as a new style, which is then used to control the style by amplifying each channel of the content tokens:

$$\mathbf{C}' = \mathbf{C} \odot \mathbf{S}', \tag{3}$$

where each row vector in $\mathbf{C} \in \mathbb{R}^{m \times d}$ indicates a content token, and $\odot$ indicates element-wise multiplication. Such a style modulation affects the operation of subsequent embedding layers (which we implemented by fully-connected layers) and thus control the style of generated images.

Implementation details. We translate the generated visual tokens to the image by concatenation and reshaping. In practice, we incorporate the skip connections in each layer of the TokenGAN to upsample and sum the RGB outputs corresponding to different resolutions [26]. Such a design makes it possible for the generator to first output low resolution images that are not affected by the higher-resolution layers significantly [23, 24].

Table 1: Detailed statistics of the datasets used in this paper.

Dataset	Resolution	#Images	#Images viewed
FFHQ-256 [26]	256×256	70 K	25 M
FFHQ-1024 [26]	1024×1024	70 K	15 M
LSUN CHURCH [44]	256×256	126 K	15 M

3.3 Overall optimization

Style mixing. To further encourage the styles to localize, we employ the mixing regularization technique during training [26]. To be specific, a given percentage of images are generated using two random latent codes. We run two latent codes $\mathbf{z}_1$, $\mathbf{z}_2$ through the mapping network, and have the corresponding $\{\mathbf{s}_1^1, \cdots, \mathbf{s}_n^1\}$, $\{\mathbf{s}_1^2, \cdots, \mathbf{s}_n^2\}$. We randomly choose an inject point to mix the styles, so that a part of the final style tokens from $\mathbf{z}_1$ and another part from $\mathbf{z}_2$. This regularization technique prevents the network from assuming that a set of style tokens are correlated. A similar strategy is adopts in terms of different layers following the style-based generator [26].

Optimization objectives. We denote the generated image output by the token-based generator as:

$$G(\mathbf{z}) = g(f(\mathbf{z}), \mathbf{C}), \tag{4}$$

where $\mathbf{z} \in \mathcal{Z} \sim \mathcal{N}(0, 1)$ indicates the input latent code, $\mathbf{C} = \{\mathbf{c}_1, \cdots, \mathbf{c}_m\}$ denotes the sequence of learned constant content tokens through the input layer, and f, g, G indicate the mapping network, the generator network and the token-based generator, respectively. Our token-based generator follows the same optimization objectives used by the style-based generator [26, 27], i.e., an non-saturating logistic adversarial loss with R1 regularization. To be specific, the generator loss is:

$$L_G = \mathbb{E}_{\mathbf{z}\sim\mathcal{N}(0,1)}[log(1 + exp(-D(G(\mathbf{z}))))], \tag{5}$$

where D is the discriminator and the discriminator loss is

$$L_D = \mathbb{E}_{\mathbf{x}\sim P_d}[log(1 + exp(-D(\mathbf{x})))] + \mathbb{E}_{\mathbf{z}\sim\mathcal{N}(0,1)}[log(1 + exp(D(G(\mathbf{z}))))] + R_1(\psi), \tag{6}$$

where $x \sim P_d$ denotes images from the real data, and $R_1(\psi)$ is the regularization term calculated by:

$$R_1(\psi) = \frac{\gamma}{2}\mathbb{E}_{\mathbf{x}\sim P_d}\left[\|\nabla D_\psi(x)\|^2\right]. \tag{7}$$

Specifically, the R_1 regularization term [30] is computed less frequently than the main loss function following the commonly used lazy regularization strategy, thus greatly diminishing the computational cost and the overall memory usage [26, 27, 47].

4 Experiments

4.1 Experiment Setup

Dataset To evaluate the token-based generator and make fair comparisons with the SOTA approach (i.e., the style-based generator [27]), we conduct experiments on the most commonly-used public research datasets, i.e., Flickr-Faces-HQ (FFHQ) [26] and Large-scale Scene Understanding (LSUN) [44]. The detailed statistics of the datasets can be found in Table 1. Specifically, FFHQ consists of 70,000 high-quality images of human face, which is able to evaluate the model's ability to synthesize high-frequency details. We conduct experiments on FFHQ with both 256×256 and 1024×1024 image size following previous works [15, 21]. Besides, we choose the church set of LSUN [44] to evaluate the synthesis in terms of complex scenes.

Evaluation. In this section, we conduct both quantitative and qualitative experiments on the token-based generator. Specifically, we report quantitative results by three commonly-used metrics, i.e., Fréchet Inception Distance (FID) [19], Precision and Recall [28, 37]. We use FID as it has been widely used in many generation tasks as an effective perceptual metric [4, 25, 26]. In addition to the image quality assessment, we use the metrics of Precision and Recall to evaluate the distribution learned by GAN. Specifically, the Precision metric intuitively measures the quality of samples from the generated images, while recall measures the proportion of real images that is covered by the learned distribution by GAN [37].

Figure 4: Uncurated results for FFHQ-1024 [26] and LSUN CHURCH [44]. The results show that token-based generator is able to synthesize both high-frequency details (e.g., hair and beard on the human face) and plausible structures for complex scenes (e.g., the church).

Table 2: Quantitative comparisons with the state-of-the-art model (i.e., StyleGAN2 [27]) on FFHQ-256 [26], FFHQ-1024[26] and LSUN CHURCH [44]. We compute each metric 10 times with different random seeds and report their average. The result shows that the token-based generator achieves comparable performance with the SOTA. $\uparrow$ the higher, the better. $\downarrow$ the lower, the better.

Model	Metrics	FFHQ-256 [26]	FFHQ-1024 [26]	LSUN CHURCH [44]
StyleGAN2 [26]	FID$\downarrow$	6.09 ± 0.051	5.20 ± 0.049	5.88 ± 0.043
	Precesion$\uparrow$	0.643 ± 0.002	0.653 ± 0.002	$0.587 \pm .002$
	Recall $\uparrow$	0.401 ± 0.002	0.411 ± 0.002	0.358 ± 0.002
TokenGAN(Ours)	FID$\downarrow$	5.41 ± 0.050	5.21 ± 0.032	5.56 ± 0.037
	Precesion$\uparrow$	0.660 ± 0.002	0.651 ± 0.001	0.577 ± 0.002
	Recall$\uparrow$	0.447 ± 0.003	0.442 ± 0.004	0.376 ± 0.002

Training details We use 8 NVIDIA V100 GPUs for training each model. We build upon the public Pytorch implementation of StyleGAN2[1] [27]. For fair comparisons, we inherit all of the training details and parameter settings from the configuration E setup. All models are trained and report the best results under the same training iterations. For each iteration, we set the size of mini-batch as 32. We use the Adam solver [3] with the same momentum parameters β_1 = 0, β_2 = 0.99 to train both the generator and discriminator. We apply R_1 regularization for every 16 iterations.

4.2 Unconditional image synthesis

Quantitative comparisons. We report quantitative comparisons with StyleGAN2 [27] due to its state-of-the-art performance. For fair comparisons, we report the results with the same training iterations on FFHQ-256 [26], FFHQ-1024 [26] and LSUN CHURCH [44] by three objective metrics, i.e., FID, Precision and Recall in Table 2. The results show that the token-based generator has achieved the state-of-the-art results in terms of both perceived quality and distribution modeling.

Qualitative results. To demonstrate the image quality of the synthesis results for the token-based generator, we randomly sample noise $\mathbf{z} \in \mathcal{Z} \sim \mathcal{N}(0, 1)$ to generate images for high-resolution images (i.e., with resolution 1024×1024) of human face and the images of complex church scenes. The uncurated results can be found in Figure 4. It shows that the token-based generator is able to synthesize both high-frequency details of the human face and plausible structures for complex scenes.

[1] https://github.com/rosinality/stylegan2-pytorch

Figure 5: Visual results of image inversion by StyleGAN2 [27] and the token-based generator. We adopt the same inversion technique following the common paradigm [1, 2, 26]. It shows that the token-based generator is able to reconstruct fine-grained details with the dense style control over image regions (e.g., the facial expressions in the third case).

Image inversion. To better apply well-trained GANs to real-world applications, the technique of GAN inversion has attracted an increasing attention [38, 48, 1, 2]. Such an inversion technique enables real image editing by searching the most accurate latent code in the learned latent space to recover the real image [48]. It has been demonstrated that the model with better inversion results tend to learn a better image modeling of the real data [27]. Specifically, we randomly sample several real images from FFHQ-256 [27] and adopt the optimization-based inversion technique used by most works [26, 38, 48]. For fair comparisons, we use the same setting following StyleGAN2 [26]. The inversion results by the style-based generator and the token-based generator are shown in Figure 5. It shows that the token-based generator is able to reconstruct real images better. For example, StyleGAN2 tends to reconstruct the lips and cloths in the fifth case with the same red color, while the token-based generator is able to reconstruct accurate colors for the lips and cloths respectively.

Image interpolation. For further exploring the property of the learned latent space by the token-based generator, we perform image interpolation by a linear interpolation in the learned intermediate latent space $\mathcal{S}$. In practice, we randomly sample two noises $\mathbf{z}_1, \mathbf{z}_2$ and run them through the mapping network to have their corresponding sequence of style tokens $\{\mathbf{s}_1^1, \cdots, \mathbf{s}_n^1\}$, $\{\mathbf{s}_1^2, \cdots, \mathbf{s}_n^2\}$. After that we perform linear interpolation for each token by $\mathbf{s}_i^3 = \alpha \times \mathbf{s}_i^1 + (1-\alpha) \times \mathbf{s}_i^2, \alpha \in (0, 1)$ and use the new style $\mathbf{s}^3$ to synthesize a new image. As shown in Figure 6, we show the sampled styles $\mathbf{s}^1, \mathbf{s}^2$ in the first column and the last column, and the interpolated styles between them. The results show that the token-based generator is able to perform perceptually smooth transition between two different styles (e.g., glasses, age, hair length, coloring, etc.)

4.3 Ablation study

To study the effectiveness of each component of the token-based generator, this section presents ablation studies on FFHQ-256 [26] after 0.4 million training iterations in terms of FID [19]. Specifically, we calculate the metric 10 times with different random seed and report the average results.

The number of style tokens. To study the effectiveness of the style tokens, we conduct ablation studies by using different number of style tokens. As shown in Table 3, with the growing number of the style tokens, the quality of images are improved. However, when the number of style tokens grows

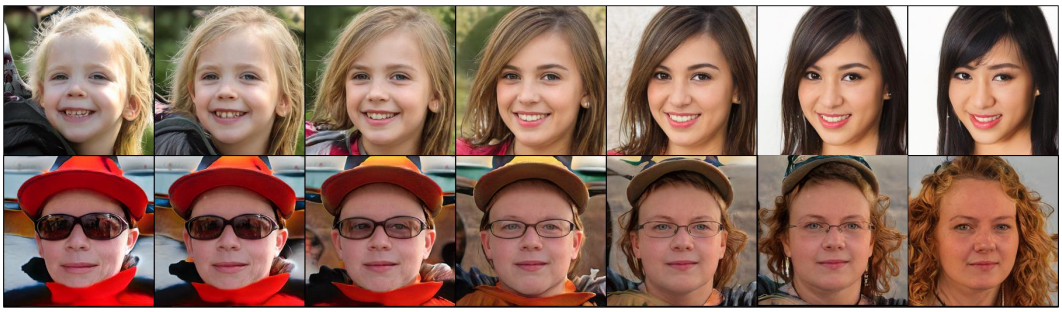

Figure 6: Results of image interpolation by the token-based generator. In practice, we sample two latent codes (the 1st and the 7th column) and perform linear interpolation of them in the learned intermediate latent space (interpolated results are from the 2th to the 6th column). The token-based generator shows smooth transition of high-level attributes in results (e.g., glasses in the second row).

Table 3: Comparisons on the style tokens.

style tokens	FID
8	7.66 ± 0.060
16	7.02 ± 0.060
32	$\mathbf{6.81 \pm 0.044}$
64	7.60 ± 0.071

Table 4: Comparisons on the style normalization.

Model	FID
InstanceNorm [39]	13.7 ± 0.051
PixelNorm [24]	6.96 ± 0.050
LayerNorm [41]	$\mathbf{6.81 \pm 0.044}$

to 64, the performance drops. We speculate that, with the growing number of style tokens, there are more style tokens to model the distribution of styles for different semantics, leading to a better result. However, when the number becomes large, the learning will be difficult and performance would drop.

Style normalization layers. We compare different style normalization layers in Table 4. The results show that, PixelNorm [24] and LayerNorm [41] significantly outperforms the style normalizatioin by instance normalization in the token-based generator [39]. We choose LayerNorm as the style normalization layer in the token-based generator.

5 Conclusion and Discussion

In this paper, we propose a novel perspective for image synthesis by view the task as a visual token generation problem. Specifically, we propose a token-based generator under this formulation for image synthesis. The token-based generator takes the content tokens as input and manipulates them using the style tokens generated from the latent space. Such a process is achieved by the attention mechanisms of Transformers, in particular, the cross-attention mechanism. Coupled with the adversarial training, our model can learn disentangled image representations and semantic-aware correspondence between content and style tokens. Both quantitative and qualitative results demonstrate the effectiveness of our model in generating high-quality images.

We notice that the proposed token-based generator is built by a cross-attention modeling between the style tokens and the content tokens, and it does well in modeling correspondence between style and content tokens. However, it has been widely demonstrate the self-attention module is able to improve the model capacity significantly by a global modeling. In the future, we would like to improve the token-based generator by incorporating the self-attention module, which may help in the visual token generation for image synthesis.

References

[1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *ICCV*, pages 4432–4441, 2019.

[2] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan++: How to edit the embedded images? In *CVPR*, pages 8296–8305, 2020.

[3] Diederik P. Kingma andJimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *ICLR*, 2015.

[4] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *ICLR*, 2018.

[5] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.

[6] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229, 2020.

[7] Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. Pre-trained image processing transformer. *CVPR*, 2021.

[8] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pretraining from pixels. In *ICML*, pages 1691–1703, 2020.

[9] Edo Collins, Raja Bala, Bob Price, and Sabine Susstrunk. Editing in style: Uncovering the local semantics of gans. In *CVPR*, pages 5771–5780, 2020.

[10] Emily Denton, Soumith Chintala, Arthur Szlam, and Robert Fergus. Deep generative image models using a laplacian pyramid of adversarial networks. In *NeurIPS*, pages 1486–1494, 2015.

[11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186, 2019.

[12] Chris Donahue, Zachary C Lipton, Akshay Balsubramani, and Julian McAuley. Semantically decomposing the latent spaces of generative adversarial networks. In *ICLR*, 2018.

[13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.

[14] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis. *CVPR*, 2021.

[15] Rinon Gal, Dana Cohen, Amit Bermano, and Daniel Cohen-Or. Swagan: A style-based wavelet-driven generative model. *arXiv preprint arXiv:2102.06108*, 2021.

[16] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014.

[17] Jinjin Gu, Yujun Shen, and Bolei Zhou. Image processing using multi-code gan prior. In *CVPR*, pages 3012–3021, 2020.

[18] Kai Han, Yunhe Wang, Hanting Chen, Xinghao Chen, Jianyuan Guo, Zhenhua Liu, Yehui Tang, An Xiao, Chunjing Xu, Yixing Xu, et al. A survey on visual transformer. *arXiv preprint arXiv:2012.12556*, 2020.

[19] 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 *NeurIPS*, pages 6629–6640, 2017.

[20] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, pages 1501–1510, 2017.

[21] Drew A Hudson and C Lawrence Zitnick. Generative adversarial transformers. *arXiv preprint arXiv:2103.01209*, 2021.

[22] Ali Jahanian, Lucy Chai, and Phillip Isola. On the" steerability" of generative adversarial networks. In *ICLR*, 2019.

[23] Animesh Karnewar and Oliver Wang. Msg-gan: Multi-scale gradients for generative adversarial networks. In *CVPR*, pages 7799–7808, 2020.

[24] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *ICLR*, 2018.

[25] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *NeurIPS*, 2020.

[26] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4401–4410, 2019.

[27] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, pages 8110–8119, 2020.

[28] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. In *NeurIPS*, 2019.

[29] Sachit Menon, Alexandru Damian, Shijia Hu, Nikhil Ravi, and Cynthia Rudin. Pulse: Self-supervised photo upsampling via latent space exploration of generative models. In *CVPR*, pages 2437–2445, 2020.

[30] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *ICML*, 2018.

[31] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018.

[32] Aäron van den Oord, Nal Kalchbrenner, Oriol Vinyals, Lasse Espeholt, Alex Graves, and Koray Kavukcuoglu. Conditional image generation with pixelcnn decoders.

[33] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, pages 2337–2346, 2019.

[34] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *ICML*, pages 4055–4064, 2018.

[35] Stanislav Pidhorskyi, Donald A Adjeroh, and Gianfranco Doretto. Adversarial latent autoencoders. In *CVPR*, pages 14104–14113, 2020.

[36] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. *CVPR*, 2021.

[37] Mehdi SM Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. In *NeurIPS*, pages 5234–5243, 2018.

[38] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *CVPR*, pages 9243–9252, 2020.

[39] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *CVPR*, pages 6924–6932, 2017.

[40] Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *ICML*, pages 1747–1756, 2016.

[41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.

[42] Yi Wang, Ying-Cong Chen, Xiangyu Zhang, Jian Sun, and Jiaya Jia. Attentive normalization for conditional image generation. In *CVPR*, pages 5094–5103, 2020.

[43] Fuzhi Yang, Huan Yang, Jianlong Fu, Hongtao Lu, and Baining Guo. Learning texture transformer network for image super-resolution. In *CVPR*, pages 5791–5800, 2020.

[44] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.

[45] Yanhong Zeng, Jianlong Fu, and Hongyang Chao. Learning joint spatial-temporal transformations for video inpainting. In *ECCV*, pages 528–543, 2020.

[46] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *ICML*, pages 7354–7363, 2019.

[47] Heliang Zheng, Jianlong Fu, Yanhong Zeng, Jiebo Luo, and Zheng-Jun Zha. Learning semantic-aware normalization for generative adversarial networks. *NeurIPS*, 33:21853–21864, 2020.

[48] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *ECCV*, pages 592–608, 2020.

[49] Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. Sean: Image synthesis with semantic region-adaptive normalization. In *CVPR*, pages 5104–5113, 2020.

[50] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *ICLR*, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 5
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 4
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Table 2
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section 4
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] See Section 4
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]