

Rethinking the Objectives of Vector-Quantized Tokenizers for Image Synthesis

Yuchao Gu¹, Xintao Wang², Yixiao Ge², Ying Shan², Mike Zheng Shou^{1*}

¹Show Lab, National University of Singapore ²ARC Lab, Tencent PCG

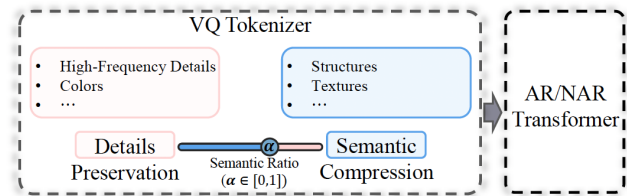
<https://github.com/TencentARC/BasicVQ-GEN>

Abstract

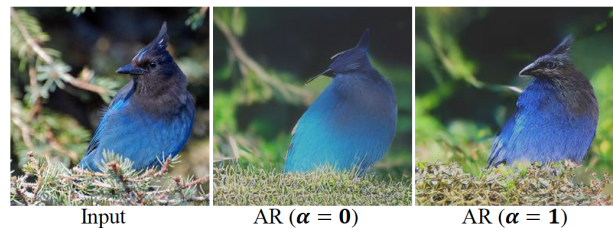
Vector-Quantized (VQ-based) generative models usually consist of two basic components, i.e., VQ tokenizers and generative transformers. Prior research focuses on improving the reconstruction fidelity of VQ tokenizers but rarely examines how the improvement in reconstruction affects the generation ability of generative transformers. In this paper, we surprisingly find that improving the reconstruction fidelity of VQ tokenizers does not necessarily improve the generation. Instead, learning to compress semantic features within VQ tokenizers significantly improves generative transformers' ability to capture textures and structures. We thus highlight two competing objectives of VQ tokenizers for image synthesis: **semantic compression** and **details preservation**. Different from previous work that prioritizes better details preservation, we propose **Semantic-Quantized GAN (SeQ-GAN)** with two learning phases to balance the two objectives. In the first phase, we propose a semantic-enhanced perceptual loss for better semantic compression. In the second phase, we fix the encoder and codebook, but enhance and finetune the decoder to achieve better details preservation. Our proposed SeQ-GAN significantly improves VQ-based generative models for both unconditional and conditional image generation. Specifically, SeQ-GAN achieves a Fréchet Inception Distance (FID) of 6.25 and Inception Score (IS) of 140.9 on 256×256 ImageNet generation, which is a remarkable improvement over VIT-VQGAN (714M), which obtains 11.2 FID and 97.2 IS.

1. Introduction

In recent years, remarkable progress has been made in image synthesis using likelihood-based generative methods, such as diffusion models [11, 45], autoregressive (AR) [15, 43, 56, 57], and non-autoregressive (NAR) [6, 18] transformers. These models offer stable training and better diversity compared to Generative Adversarial Networks (GANs) [29, 30]. However, unlike GANs, which can gener-



(a) Two competing objectives when optimizing VQ tokenizers.



(b) Influence of VQ tokenizers with different semantic ratio (α) on autoregressive (AR) transformers.

Figure 1. Visualizing impact of VQ tokenizers on generative transformers with α trade-off between details preservation and semantic compression in VQ tokenizer training.

ate high-resolution (e.g., 256² and 512²) images at one forward pass, likelihood-based methods usually require multiple forward passes by sequential decoding [15, 56] or iterative refinement [6, 18]. Consequently, early works [7, 23, 39], which maximize likelihood on pixel space, are limited in their ability to synthesize high-resolution images due to the high computational cost and slow decoding speed.

Instead of directly modeling the underlying distribution in the pixel space, recent vector-quantized (VQ-based) generative models [50] construct a discrete latent space for generative transformers. There are two basic components in VQ-based generative models, i.e., VQ tokenizers and generative transformers. VQ tokenizers learn to quantize images into discrete codes, and then decode the codes to recover the input images, which process is termed as **reconstruction**. Then, a generative transformer is trained to learn the underlying distribution in the discrete latent space constructed by the VQ tokenizer. Once trained, the generative transformer can be used to sample images from the underlying distribution, and this process is termed as **generation**. Thanks to the discrete latent space, VQ-based generative

*Corresponding Author.



Figure 2. Generation results of SeQ-GAN+NAR. 1st row: LSUN- $\{\text{cat, bedroom, church}\}$. 2nd row: FFHQ and ImageNet.

models [6, 15, 45] can easily scale up to synthesize high-resolution images without prohibitive computation cost.

The VQ tokenizer has received much attention as the core component in VQ-based generative models. Various techniques, such as factorized codes and smaller compression ratio in VIT-VQGAN [56], recursive quantization in Residual Quantization [34], and multichannel quantization with spatial modulated decoder in MoVQ [60], have been used to compress more fine-grained details into VQ tokenizers, leading to steadily improving reconstruction fidelity. However, none of the previous works have carefully examined a fundamental question, *how the improved reconstruction of VQ tokenizers affects the generation*. Lacking such analysis is due to two main reasons: 1) the underlying assumption that “*better reconstruction, better generation*”, and 2) the absence of a visualization pipeline to intuitively compare generation results of various VQ tokenizers.

In this paper, we introduce a visualization pipeline for examining how different VQ tokenizers influence generative transformers. Unlike previous works that compare randomly-sampled generation results, our approach models specific images and facilitates a straightforward comparison of the generative transformer’s ability using different VQ tokenizers. The key idea is to reduce the flexibility of the sampling process by providing ground-truth context to generative transformers, which can be easily implemented with an autoregressive (AR) transformer with causal attention.

Our proposed visualization pipeline leads us to two important observations. 1) Improving the reconstruction fidelity of VQ tokenizers does not necessarily improve the generation. 2) Learning to compress semantic features within VQ tokenizers significantly improves generative transformers’ ability to capture textures and structures. As shown in Fig. 1, increasing the semantic ratio ($\alpha=1$) improves the AR transformer’s ability to capture texture and structure, while decreasing it ($\alpha=0$) results in the transformer modeling rough colors instead. These observations arise due to the competing objectives of reconstruction and

generation optimization. Reconstruction aims to retain variation in the dataset by favoring latent spaces with larger variance (*i.e.*, weaker separability), whereas generation optimization favors latent spaces with smaller variance (*i.e.*, better separability) to optimize a classification objective.

Our observations reveal that there are two competing objectives for VQ tokenizers: *semantic compression* and *details preservation*, but recent VQ tokenizers [34, 44, 56, 60] have primarily focused on the latter. To balance the two objectives for better generation, we propose **Semantic-Quantized GAN (SeQ-GAN)**, which consists of two learning phases. The first phase utilizes a semantic-enhanced perceptual loss to achieve semantic compression, while the second phase finetunes the decoder to restore fine-grained details while preserving structures and textures. Compared to previous VQ tokenizers, SeQ-GAN compresses semantic features rather than fine-grained details (*e.g.*, high-frequency details, colors) into codebook and finetunes the decoder to restore those details, which does not affect transformer learning but improves local details generation.

Our main contributions are summarized as follows. (1) We rethink the common assumption “*better reconstruction, better generation*” in recent VQ tokenizers, and propose a visualization pipeline to explore the impact of different VQ tokenizers on generative transformers. (2) We identify two competing objectives in optimizing VQ tokenizers: *semantic compression* and *details preservation*, and introduce SeQ-GAN as a solution that balances these objectives to achieve better generation quality. (3) Our SeQ-GAN achieves significant improvements over prior VQ tokenizers in both conditional and unconditional image generation, as demonstrated through experiments with both AR and NAR transformers. (Generation results are shown in Fig. 2).

2. Related Work

VQ-based Generative Models. The VQ-based generative model is first introduced by VQ-VAE [50], which constructs a discrete latent space by VQ tokenizers and learns

the underlying latent distribution by prior models [8, 49]. VQGAN [15] improves upon this by utilizing perceptual loss [28, 58] and adversarial learning [17] in training VQ tokenizers, and using autoregressive transformers [41] as the prior model, leading to significant improvements in generation quality. VQ-based generative models have been applied in various generation tasks, such as image generation [6, 15, 56], video generation [16, 25, 54], text-to-image generation [12, 43, 45, 57], and face restoration [19, 51, 62].

Building on the success of VQGAN [15], recent works have focused on improving the two fundamental components of VQ-based generative models: VQ tokenizers and generative transformers. To enhance VQ tokenizers, VIT-VQGAN [56] proposes quantizing image features into factorized and L2-normed codes with a larger codebook and small compression ratio, achieving finer reconstruction results. Residual Quantization [34] recursively quantizes feature maps using a shared codebook to precisely approximate image features. MoVQ [60] enhances the VQ tokenizer’s decoder with modulation [27] and proposes multi-channel quantization with a shared codebook, resulting in state-of-the-art reconstruction results. Different from previous works, we argue that improving reconstruction fidelity does not necessarily lead to better generation quality.

Another line orthogonal to our work is improving generative transformers. Early works adopt autoregressive (AR) transformers [15, 44, 56]. However, AR transformers suffer from low sampling speed and ignore bidirectional contexts. To overcome these limitations, non-autoregressive (NAR) transformers are introduced based on different theories, like mask image modeling [2, 21] (*i.e.*, MaskGIT [6]) and discrete diffusion [1, 26] (*i.e.*, VQ-diffusion [18, 48]). In this paper, we demonstrate that integrating our SeQ-GAN as the VQ tokenizer consistently enhances the generation quality of both AR and NAR transformers.

Visual Tokenizers for Generative Pretraining. Recent works in large-scale generative visual pretraining also explore the potential of the visual tokenizer. Instead of directly performing mask image modeling on pixels [21, 53], the pioneer BEiT [2] reconstructs masked patches quantized by a discrete VAE [43]. Follow-up works further strengthen the semantics of the visual tokenizer, such as PeCo [13], which adopts contrastive perceptual loss [9, 20] during tokenizer training, and mc-BEiT [35], which softens and reweights the masked prediction target during visual pretraining. To further reduce the low-level representation in the visual tokenizer, iBOT [61] abandons reconstructing pixels, but updates the tokenizer online during the pretraining. BEiT-v2 [40] formulates the training objective of the visual tokenizer by reconstructing semantic features extracted by CLIP [42]. Unlike prior attempts to remove low-level representation interference in visual pretraining, we highlight the importance of semantic compression and details preser-

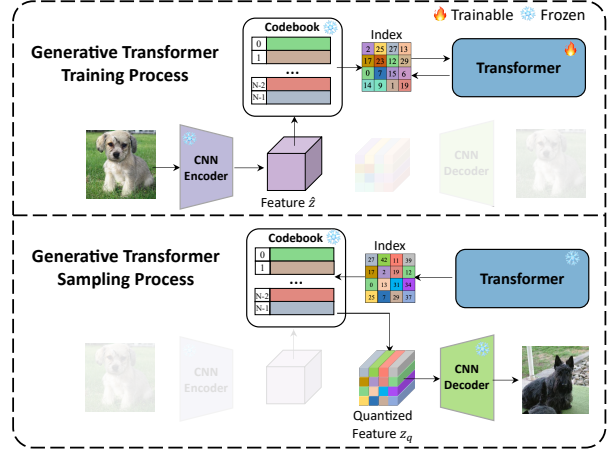


Figure 3. The influence of VQ tokenizers on the training and sampling process of generative transformers.

vation in training VQ tokenizers for image synthesis.

3. Methodology

In this section, we first review how VQ tokenizers affect generation in VQ-based generative models in Sec. 3.1. Then, we present a visualization pipeline in Sec. 3.2 to examine the impact of different VQ tokenizers on generative transformers. Based on this pipeline, we make two critical observations in Sec. 3.3, highlighting the competing objectives in designing VQ tokenizers. Finally, we propose SeQ-GAN in Sec. 3.4 as a solution that balances these objectives to improve generation quality.

3.1. Preliminaries

In this section, we cover the fundamental process of VQ-based generative models and highlight the potential impact of VQ tokenizers on generation results.

Reconstruction: training VQ tokenizers. The role of VQ tokenizers is to compress the image into discrete indices. Specifically, a VQ tokenizer is comprised of an encoder E , a decoder G and a codebook $\mathcal{Z} = \{z_k\}_{k=1}^K$ with K discrete codes. Given an input image $x \in \mathbb{R}^{H \times W \times 3}$, a latent feature $\hat{z} \in \mathbb{R}^{\frac{H}{f} \times \frac{W}{f} \times n_z}$ is first extracted, where n_z and f represent the dimension of the latent features and the spatial compression ratio, respectively. Then, the feature vector at each spatial position (i, j) is quantized to the nearest code in the codebook by

$$z_q = \mathbf{q}(\hat{z}) := \left(\arg \min_{z_k \in \mathcal{Z}} \|\hat{z}_{ij} - z_k\| \right) \in \mathbb{R}^{\frac{H}{f} \times \frac{W}{f} \times n_z}. \quad (1)$$

The decoder G is responsible for decoding the quantized features back to the image space, *i.e.*, $\hat{x} = G(z_q)$.

The training objective of the VQ tokenizer is to minimize the reconstruction error with respect to the input image. Following VQGAN [15] to use adversarial loss (\mathcal{L}_{adv}) [17] and perceptual loss (\mathcal{L}_{per}) [28, 58], the reconstruction objective

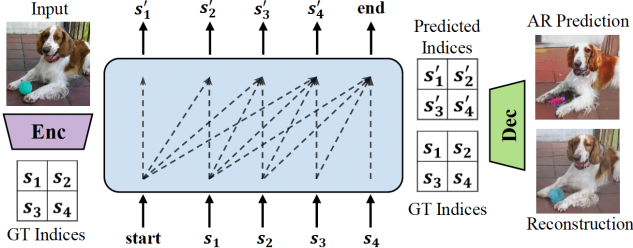


Figure 4. Visualization pipeline to examine the influence of VQ tokenizers on generative transformers.

can be formulated as

$$\mathcal{L}(E, G, Z) = \mathcal{L}_{vq} + \mathcal{L}_{per} + \mathcal{L}_{adv}, \text{ where}$$

$$\mathcal{L}_{vq} = \|x - \hat{x}\|_1 + \|\text{sg}[E(x)] - z_q\|_2^2 + \beta \|\text{sg}[z_q] - E(x)\|_2^2. \quad (2)$$

In Eq. 2, $\text{sg}[\cdot]$ means stop-gradient and $\beta \|\text{sg}[z_q] - E(x)\|_2^2$ is known as the commitment loss [50], where the commitment weight β is set to 0.25 following [15, 50, 56].

Training generative transformers. As shown in Fig. 3, the encoder and codebook of a trained VQ tokenizer define a discrete latent space that quantizes an image into a sequence of discrete indices for generative transformer training. This sequence serves as input and label in training the generative transformer with token classification loss. In this paper, we use the autoregressive (AR) transformer in VQGAN [15] and the non-autoregressive (NAR) transformer in MaskGIT [6]. Therefore, the quality of the discrete latent space defined by the encoder and codebook of VQ tokenizers will influence the generative transformer training.

Generation: sampling from generative transformers. After training a generative transformer, we can sample discrete index sequences from it through either autoregressive decoding [15] or iterative refinement [6]. To map the discrete indices back to visual details, we retrieve the corresponding feature from the codebook and decode it into image space using the VQ tokenizer’s decoder. Therefore, the decoder will affect the generation quality by influencing the index-to-visual-details mapping.

3.2. Pipeline for Visualizing VQ Generative Models

Recent VQ-based generative models examine their designs by looking into the random sampled generation results, where different sampling techniques are adopted (*e.g.*, top- p top- k sampling [24], classifier-free guidance [22], or rejection sampling [43]). However, instead of examining random samples, we are more curious about how generative transformers model specific images, enabling us to check the influence of different VQ tokenizers on generative transformers side by side. To achieve that goal, we propose to reduce the flexibility of the sampling process by providing ground-truth (GT) contexts for predicting each index, which can be easily implemented by AR transformers.

The pipeline is shown in Fig. 4. First, we train a VQ

Model	Params	rFID↓	Generation FID↓			
			AR	AR-L	AR-L-2×	NAR
baselineVQ	54.5M	3.45	16.97	13.86	11.49	13.26
+Conv×2	70.0M	3.22	17.19	14.50	12.03	13.51
+Attention×2	61.4M	2.90	17.42	14.91	12.04	14.02

Table 1. Comparison of the baseline and decoder-enhanced VQ tokenizers on the reconstruction FID (rFID) and generation FID, evaluated on different transformer configurations.

tokenizer along with its corresponding AR transformer. To analyze a specific image, we obtain the GT index sequence $s = [s_i]_{i=1}^N$ from the VQ tokenizer and feed it to the trained AR transformer, similar to the teacher forcing strategy [3, 52] used in training AR transformers. Because the AR transformer adopts casual attention [41], it does not directly access the GT indices, but can access all GT context indices for predicting each index. Given the same context (*i.e.*, preceding GT indices), the next index prediction task is well-controlled and thus we can get the top-1 predicted index sequence s' within one forward pass. Finally, we decode the GT sequence s and the AR predicted sequence s' back to the image space by the decoder of the VQ tokenizer. Following this approach, we are able to visualize both the reconstruction of VQ tokenizers and the upper limit prediction of AR transformers for specific images.

3.3. Rethinking the Objectives of VQ Tokenizers

3.3.1 Reconstruction vs. Generation

Motivation. Recent advancements in VQ tokenizers have led to improved reconstruction results, with MoVQ [60] in particular enhancing their decoder with modulation to add variation to quantized code and achieve the highest reconstruction fidelity. However, few studies investigate whether improvements of reconstruction fidelity of VQ tokenizers benefit generation quality. To address this gap, we conduct the following experiment to answer this question.

Experimental Settings. In Sec. 3.1, we identify two key factors in VQ tokenizers that affect generation: 1) the quality of the discrete latent space defined by the encoder/codebook, and 2) the index-to-visual-details mapping defined by the decoder. Inspired by MoVQ [60], we keep the configuration of encoder/codebook the same, and enhance the decoder to strengthen the index-to-visual-detail mapping. Our baseline is a convolution-only VQGAN [15], and we add two extra convolution blocks or two interleaved regional and dilated attention blocks [59] at each resolution level to enhance the decoder. Based on each tokenizer, we train the generative transformer with different configurations, including different parameter sizes (AR and AR-Large), different types (AR and NAR), and different training iterations (AR-Large and AR-Large-2×). Additional experimental settings can be found in the supplementary.

Results. The results presented in Table. 1 show that en-

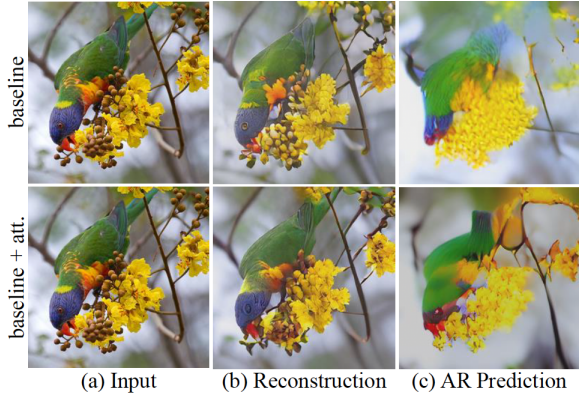


Figure 5. Visualizing the reconstruction and AR prediction of the baseline tokenizer and its attention-enhanced variant.

enhancing the decoder improves reconstruction fidelity, but it does not necessarily lead to better generation quality. Surprisingly, the baseline tokenizer achieves the best generation quality. Assuming that the quality of the discrete latent space (defined by encoder/codebook) remains unchanged, enhancing the decoder should improve generation quality by improving the index-visual-details mapping. However, in reality, enhancing the decoder leads to a degradation in generation quality. This suggests that *jointly learning* the encoder/codebook with an enhanced decoder actually degrades the quality of the discrete latent space.

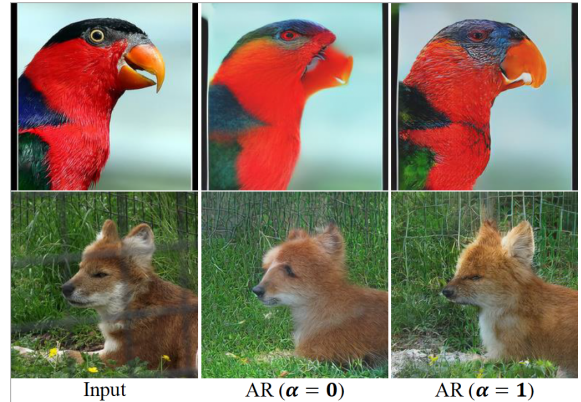
Using the proposed visualization pipeline in Sec. 3.2, we visualize the reconstruction and AR prediction results of the baseline tokenizer and its attention-enhanced variant in Fig. 5. Although the attention-enhanced variant leads to a more consistent reconstruction, the AR transformer faces challenges in capturing the details and can only predict a rough color for the main object, even given ground-truth contexts. This highlights the generative transformers’ difficulties in modeling the discrete latent space.

The discrepancy between reconstruction and generation is due to the conflicting optimization objectives. In reconstruction training, VQ tokenizers prefer a latent space with larger variance (*i.e.*, weaker separability) to retain the variation of datasets, while generative transformer training prefers smaller variance (*i.e.*, better separability), because it optimizes the classification objective (*i.e.*, cross-entropy). Therefore, in Table. 1 and Fig. 5, a powerful decoder promotes encoding more variation in the codebook, which hinders the separability of the discrete latent space and thus results in suboptimal generation performance. Through the result, we arrive at the following observation.

Observation 1. *Improving the reconstruction fidelity of VQ tokenizers does not necessarily improve the generation.*

3.3.2 Details Preservation vs. Semantic Compression

Motivation. Observation 1 suggests that compressing more fine-grained details within the tokenizer in reconstruction



(a) Visualization of the AR predicted results when trained using VQ tokenizers with different semantic ratios (α).
 (b) Reconstruction FID and Generation FID with different semantic ratio (α) in optimizing VQ tokenizers.

Figure 6. Influence of the semantic ratio α in VQ tokenizers on generation quality.

does not always improve generation. Therefore, we shift our focus towards exploring the role of semantics in VQ tokenizers for better generation quality.

Semantic-Enhanced Perceptual Loss. Unlike generative pretraining [40, 61] that uses fully semantic tokenizers, image synthesis requires consideration of low-level details. To balance the trade-off between low-level details and semantics in VQ tokenizers, we introduce a semantic-enhanced perceptual loss that controls the details/semantic ratio.

Specifically, given an input and a reference image, we extract their activation features \hat{y}^l and y^l from a pre-trained VGG [47] network. For each layer l , the feature is of shape $H_l \times W_l \times C_l$. Then, the perceptual loss can be calculated as $\mathcal{L}_{per} = \sum_l \frac{1}{H_l W_l C_l} \|\hat{y}^l - y^l\|_2^2$. To preserve details, perceptual loss [58] used in previous VQ tokenizers adopts the features from both the shallow and high layers, which we denote as L_{per}^{low} in this paper. To better compress semantic information during reconstruction, we propose a semantic-enhanced perceptual loss L_{per}^{sem} , which removes the features from shallow layers and further includes the logit feature (*i.e.*, feature before softmax classifier). The layers l to extract features can be summarized as

- \mathcal{L}_{per}^{low} : $l \in \{\text{relu}\{-1.2, 2.2, 3.3, 4.3, 5.3\}\}$,
- \mathcal{L}_{per}^{sem} : $l \in \{\text{relu}5.3, \text{logit}\}$.

We re-weight the two perceptual losses to control the proportion between the details and semantic information by

$$\mathcal{L}_{per}^{\alpha} = \alpha \mathcal{L}_{per}^{sem} + (1 - \alpha) \mathcal{L}_{per}^{low}, \quad (3)$$

where $\alpha \in [0, 1]$ is the semantic ratio.

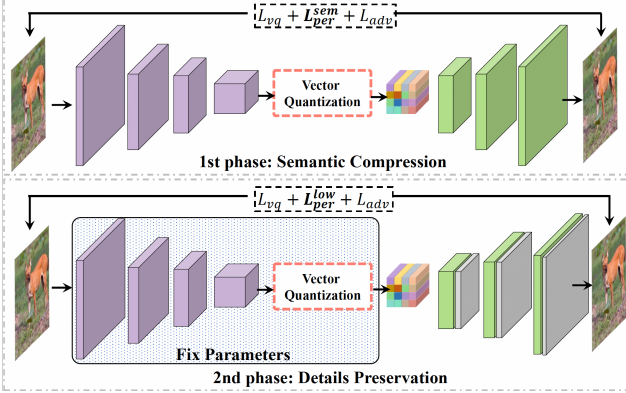


Figure 7. Pipeline of the two-phase learning in SeQ-GAN.

Results. Using the proposed semantic-enhanced perceptual loss, we examine the impact of different semantic ratios (α) during VQ tokenizer training on generation quality. In particular, we find that increasing α initially improves the reconstruction FID (rFID), with the best rFID achieved at $\alpha=0.4$ before decreasing. However, increasing α consistently improves the generation FID. Our visualizations in Fig. 6(a) and Fig. 1(b) demonstrate that the semantic-enhanced VQ tokenizer ($\alpha=1$) enables the AR transformer to capture more overall structures and textures than the baseline tokenizer ($\alpha=0$). We provide additional visualizations *in the supplementary*. These results help us arrive at the following observation.

Observation 2. *Semantic compression within VQ tokenizers benefits the generative transformer.*

3.3.3 Discussion

Tokenizers [32, 46] in Natural Language Processing (NLP) are naturally discrete and semantically meaningful, and in large-scale generative visual pretraining [40, 61], fully semantic visual tokenizers that abandon low-level information are preferred. However, VQ tokenizers in VQ-based generative models should consider low-level details. Previous works [34, 44, 56, 60] prioritize preserving details to achieve better reconstruction fidelity, but we find solely compressing fine-grained details within VQ tokenizers will degrade the discrete latent space and hinder transformer training. We argue that both semantic compression and detail preservation should be considered when designing VQ tokenizers for image synthesis.

3.4. Our Solution: SeQ-GAN

To achieve better generation quality, we propose the **Semantic-Quantized GAN (SeQ-GAN)** as the VQ tokenizer in VQ-based generative models, balancing the objectives of semantic compression and details preservation.

Fig. 7 illustrates the two-phase approach of SeQ-GAN for tokenizer learning. In the first phase, we prioritize semantic compression by applying the proposed semantic-

Model	Datset	Latent Size	Codebook K	Usage	rFID
VQGAN [15]	FFHQ	16×16	1024	42%	4.42
VIT-VQGAN [56]		32×32	8192	-	3.13
RQ-VAE [34]		16×16×4	2048	-	3.88
MoVQ [60]		16×16×4	1024	-	2.26
SeQ-GAN (Ours)		16×16	1024	100%	3.12
VQGAN [15]	ImageNet	16×16	1024	44%	7.94
VQGAN [15]		16×16	16384	5.9%	4.98
VIT-VQGAN [56]		32×32	8192	96%	1.28
RQ-VAE [34]		8×8×16	16384	-	1.83
MoVQ [60]		16×16×4	1024	-	1.12
SeQ-GAN (Ours)		16×16	1024	100%	1.99

Table 2. Reconstruction results on ImageNet and FFHQ validation set, with K representing codebook size.

enhanced perceptual loss $\mathcal{L}_{per}^{\alpha=1}$ in Eq. 3. However, semantic compression with VQ tokenizers may cause some loss of color fidelity and high-frequency details. To address this, we enhance the decoder in the second phase using interleaved block regional and dilated attention [59]. We fix the encoder and codebook of the tokenizer and finetune the enhanced decoder with $\mathcal{L}_{per}^{\alpha=0}$ to achieve better detail preservation. Note that in the second phase of tokenizer learning, we fix the discrete latent space by fixing the encoder and codebook. Therefore, our decoder-only finetuning enhances the generation quality of local details without affecting the transformer learning of structures and textures.

During the training of our SeQ-GAN, we observe the issue of low codebook usage, which has also been reported in prior VQGAN [15]. To address this issue, we incorporate entropy regularization techniques that are commonly used in self-supervised representation learning [5, 36] to mitigate the problem of empty clusters. Specifically, given the feature before quantization $\hat{z} \in \mathbb{R}^{N \times n_z}$, we aim to map \hat{z} to the codebook feature $\mathcal{Z} = \{z_k\}_{k=1}^K$. Denote the matrix $\mathcal{D} \in \mathbb{R}^{N \times K}$ as the L2 distance between each feature \hat{z}_i and each code entry z_k , we normalize it by softmax $\mathcal{D}_{i,k} = \frac{\exp(-\mathcal{D}_{i,k})}{\sum_{k=1}^K \exp(-\mathcal{D}_{i,k})}$. Then we average \mathcal{D} along the spatial size by $\bar{\mathcal{D}}_k = \frac{1}{N} \sum_{i=1}^N \mathcal{D}_{i,k}$, where $\bar{\mathcal{D}} \in \mathbb{R}^K$ can be interpreted as a soft codebook usage. To increase the codebook usage, we encourage a smoother $\bar{\mathcal{D}}$, which achieved by penalizing the entropy $H(\bar{\mathcal{D}}) = -\sum_k \bar{\mathcal{D}}_k \log \bar{\mathcal{D}}_k$. And we update the optimization objective in Eq. 2 to $\mathcal{L}_{vq'} = \mathcal{L}_{vq} + \gamma H(\bar{\mathcal{D}})$, where we fix $\gamma = 0.01$ in our experiments.

4. Experiments

4.1. Image Quantization

We train the SeQ-GAN on ImageNet [10], FFHQ [29] and LSUN [55], separately. In the first phase, we train the SeQ-GAN on ImageNet and FFHQ using the Adam [31] opti-

Model	Params	steps	FFHQ	Church	Cat	Bedroom
BigGAN [4]	164M	1	12.4	-	-	-
StyleGAN2 [30]	30M	1	3.8	3.86	7.25	2.35
ADM [11]	552M	1000	-	-	5.57	1.90
DDPM [23]	114M [†] /256M [‡]	1000	-	7.89 [†]	19.75 [†]	4.90 [‡]
DCT [37]	473M [†] /448M [‡]	>1024	13.06 [†]	7.56 [‡]	-	6.40 [‡]
VQGAN [15]	72.1M + 801M	256	11.4	7.81	17.31	6.35
ImageBART [14]	-	-	9.57	7.32	15.09	5.51
VIT-VQGAN [56]	64M + 1697M	1024	5.3	-	-	-
RQ-VAE [34]	100M + 370M [†] /650M [‡]	256	10.38 [†]	7.45 [†]	8.64 [‡]	3.04 [‡]
MoVQ + AR [60]	82.7M + 307M	1024	8.52	-	-	-
MoVQ + NAR [60]	82.7M + 307M	12	8.78	-	-	-
SeQ-GAN + AR (Ours)	57.9M + 171M	256	-	2.45	3.61	1.44
SeQ-GAN + NAR (Ours)	57.9M + 171M	12	3.62	2.25	4.60	2.05

Table 3. Quantitative comparison of unconditional image generation on FFHQ [29] and LSUN [55]-{Church, Cat, Bedroom}. The AR transformer result on FFHQ is omitted due to severe overfitting, consistent with findings in RQ-VAE [34].

mizer with a learning rate of 1e-4 for 500,000 iterations. For LSUN-{cat, bedroom, church}, we follow RQ-VAE [34] to use the pretrained SeQ-GAN on ImageNet and finetune for one epoch on each dataset. In the second phase, we finetune the enhanced decoder of SeQ-GAN on three datasets for 200,000 iterations with a learning rate of 5e-5. Detailed settings are provided *in the supplementary*.

The results are summarized in Table. 2. Since VIT-VQGAN [56], RQ-VAE [34] and MoVQ [60] prioritize the reconstruction fidelity by compressing more fine-grained details within the tokenizer, they usually require a larger latent size. Our SeQ-GAN does not pursue the reconstruction fidelity, but optimizes for better generation quality. Therefore, SeQ-GAN does not achieve the best reconstruction fidelity. However, compared to VQGAN [15], with the same latent size and codebook size, SeQ-GAN still has a large improvement in rFID and codebook usage.

4.2. Unconditional Image Generation

We train AR and NAR transformers on top of SeQ-GAN for unconditional image generation on FFHQ [29] and LSUN [55] datasets. All models are trained for 500,000 iterations with the Adam optimizer, using a learning rate of 1e-4. Detailed hyperparameters are *in the supplementary*.

From the results in Table. 3, previous state-of-the-art results are achieved by continuous diffusion model ADM [11] and StyleGAN2 [30], while VQ-based generative models typically lag behind. Using SeQ-GAN as the VQ tokenizer enables our AR/NAR transformers with 171M parameters to surpass VIT-VQGAN [56], RQ-VAE [34], and MoVQ [60], despite having fewer parameters. Our method achieves comparable performance to ADM and StyleGAN2 on both FFHQ and LSUN datasets.

Model	Params	Steps	FID	IS
BigGAN-Deep [4]	160M	1	6.95	198.2
DCT [37]	738M	>1024	36.51	-
Improved DDPM [38]	280M	250	12.26	-
ADM [11]	554M	250	10.94	101.0
VQ-VAE-2 [44]	13.5B	5120	31.11	~45
VQGAN [15]	1.4B	256	15.78	78.3
VIT-VQGAN [56]	714M	256	11.20	97.2
VIT-VQGAN [56]	1.7B	1024	4.17	175.1
RQ-VAE [34]	1.4B	1024	8.71	119.0
MoVQ + AR [60]	389M	1024	7.13	138.3
SeQ-GAN + AR	229M	256	7.55	121.3
SeQ-GAN + AR-L	364M	256	6.25	140.9
VQ-Diffusion [18]	370M	100	11.89	-
MaskGIT [6]	227M	8	6.18	182.1
MoVQ + NAR [60]	389M	12	7.22	130.1
SeQ-GAN + NAR	229M	12	4.99	189.1
SeQ-GAN + NAR-L	364M	12	4.55	200.4

Table 4. FID and Inception Score (IS) comparison of conditional image generation on ImageNet [10].

4.3. Conditional Image Generation

We train AR and NAR transformers with our SeQ-GAN tokenizer on 256×256 ImageNet generation. The model is trained with a learning rate of 1e-4 for 300 epochs to enable direct comparison with VIT-VQGAN [56] and MaskGIT [6]. Further training settings can be found *in the supplementary material*.

Results are summarized in Table. 4. Our SeQ-GAN+AR (364M, 256 sample steps) achieves FID of 6.25 and IS of 140.9, a remarkable improvement over VIT-VQGAN [56] (714M, 256 sample steps), which obtains 11.2 FID and 97.2 IS. Compared to MaskGIT [6], which obtains 6.18 FID, our SeQ-GAN+NAR achieves a better 4.99 FID with a similar sampling step and model size. Compared to MoVQ+NAR

Model	Dim \mathcal{Z}	K	Usage	rFID	AR	NAR
VQGAN	256	1024	43.5%	4.07	17.19	14.58
+ F&N	32	8192	99.7%	2.93	24.91	-
+ K-means	256	1024	100%	3.54	16.84	15.02
+ $H(\hat{D})$	256	1024	100%	3.45	16.97	13.26

Table 5. Ablation study of codebook regularization compared to VQGAN [15] baseline and VIT-VQGAN [56] with factorized and L2-normed code (F&N). K denotes the codebook size.

Loss	relu1.2	relu2.2	relu3.3	relu4.3	relu5.3	logit	rFID	AR	NAR
\mathcal{L}_{per}^{low}	✓	✓	✓	✓	✓		3.45	16.97	13.26
A	✓	✓	✓	✓	✓	✓	3.01	15.19	11.78
B		✓	✓	✓	✓	✓	2.93	14.47	11.58
C			✓	✓	✓	✓	2.81	14.08	10.52
D				✓	✓	✓	2.62	13.34	9.56
\mathcal{L}_{per}^{sem}						✓	2.77	12.07	8.84
E						✓	4.65	17.88	14.00

Table 6. Ablation of semantic-enhanced perceptual loss.

(389M, 12 sample steps), obtaining 7.22 FID and 130.1 IS, our SeQ-GAN+NAR-L (364M, 12 sample steps) achieves much better performance of 4.55 FID and 200.4 IS.

4.4. Ablation

Codebook regularization. We ablate the strategy for increasing codebook usage in the baseline setting (one-phase training with $\mathcal{L}_{per}^{\alpha=0}$). As shown in Table. 5, although the factorized and L2-normed codebook in VIT-VQGAN can largely enhance the reconstruction fidelity, its large codebook size results in a suboptimal performance on the AR transformer. Moreover, optimizing the NAR transformer on a large codebook size is unstable. Compared to the offline K-means clustering used in previous codebook learning [33], the entropy regularization used in our paper achieves a better reconstruction and generation performance.

Design of semantic-enhanced perceptual loss. Our baseline, \mathcal{L}_{per}^{low} , utilizes all five layers to compute perceptual loss. As shown in Table. 6, adding the logit feature improves both rFID and generation FID. Variant-D achieves the best rFID, while \mathcal{L}_{per}^{sem} achieves the best generation FID. This demonstrates that reconstruction fidelity does not necessarily correlate with generation performance. Removing more shallow layers consistently improves generation quality, highlighting the importance of semantics when optimizing VQ tokenizers for generation quality. However, adopting the logit feature (variant-E) without the spatial feature results in significantly worse performance. While adjusting the balance between details and semantics by removing different perceptual layers is possible, it usually requires extensive parameter tuning to match the loss scale. Instead, we fix \mathcal{L}_{per}^{low} and \mathcal{L}_{per}^{sem} and simply tune the semantic ratio α in Eq. 3 to achieve our goal.



(a) Effect of 2nd-phase tokenizer learning on generation quality.

Model	SeQ-GAN	ImageNet	FFHQ	Church	Cat	Bedroom
AR	1st phase	7.83	-	3.49	4.73	2.15
	2nd phase	7.55	-	2.45	3.61	1.44
NAR	1st phase	5.31	3.89	3.41	5.22	2.88
	2nd phase	4.99	3.62	2.25	4.60	2.05

(b) Effect of 2nd-phase tokenizer learning on generation FID.

Figure 8. Ablation study on the impact of 2nd-phase tokenizer learning on generation.

Influence of the second phase tokenizer learning. SeQ-GAN is trained with semantic-enhanced perceptual loss $\mathcal{L}_{per}^{\alpha=1}$ in the first phase, which can result in some loss of color fidelity and high-frequency details. However, by finetuning the enhanced decoder in the second phase, those details can be preserved for the generation. As shown in Fig. 8(a), the second phase learning can restore color distortion (e.g., windows). Furthermore, Fig. 8(b) shows that second phase learning consistently improves generation FID. It’s worth noting that joint learning the encoder/codebook and an enhanced decoder degrades the generation performance in our observation 1 (see Sec. 3.3.1). Therefore, the decoder-only finetuning is an effective way to promote details preservation without degrading discrete latent space.

5. Conclusion

This work examines a fundamental question in VQ-based generative models, “how the improved reconstruction of VQ tokenizers affects the generation”. To answer this question, we introduce a visualization pipeline to examine the influence of different tokenizers on AR transformers. Based on this pipeline, we find both semantic compression and details preservation should be considered in optimizing VQ tokenizers, in which previous works prioritize the latter. Based on this finding, we propose a simple solution SeQ-GAN, which achieves remarkable improvement over existing VQ-based generative models on image synthesis.

Acknowledgement This project is supported by the National Research Foundation, Singapore under its NRFF Award NRF-NRFF13-2021-0008.

References

- [1] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. *NeurIPS*, 34:17981–17993, 2021. 3
- [2] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. 3
- [3] Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. Scheduled sampling for sequence prediction with recurrent neural networks. *NeurIPS*, 28, 2015. 4
- [4] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 7
- [5] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *NeurIPS*, 33:9912–9924, 2020. 6
- [6] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative image transformer. In *CVPR*, pages 11315–11325, 2022. 1, 2, 3, 4, 7
- [7] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pre-training from pixels. In *ICML*, pages 1691–1703. PMLR, 2020. 1
- [8] Xi Chen, Nikhil Mishra, Mostafa Rohaninejad, and Pieter Abbeel. Pixelsnail: An improved autoregressive generative model. In *ICML*, pages 864–872. PMLR, 2018. 3
- [9] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 3
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. Ieee, 2009. 6, 7
- [11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *NeurIPS*, 34:8780–8794, 2021. 1, 7
- [12] 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. *NeurIPS*, 34:19822–19835, 2021. 3
- [13] Xiaoyi Dong, Jianmin Bao, Ting Zhang, Dongdong Chen, Weiming Zhang, Lu Yuan, Dong Chen, Fang Wen, and Nenghai Yu. Peco: Perceptual codebook for bert pre-training of vision transformers. *arXiv preprint arXiv:2111.12710*, 2021. 3
- [14] Patrick Esser, Robin Rombach, Andreas Blattmann, and Bjorn Ommer. Imagebart: Bidirectional context with multinomial diffusion for autoregressive image synthesis. *NeurIPS*, 34:3518–3532, 2021. 7
- [15] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, pages 12873–12883, 2021. 1, 2, 3, 4, 6, 7, 8
- [16] Songwei Ge, Thomas Hayes, Harry Yang, Xi Yin, Guan Pang, David Jacobs, Jia-Bin Huang, and Devi Parikh. Long video generation with time-agnostic vqgan and time-sensitive transformer. *arXiv preprint arXiv:2204.03638*, 2022. 3
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 3
- [18] 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 *CVPR*, pages 10696–10706, 2022. 1, 3, 7
- [19] Yuchao Gu, Xintao Wang, Liangbin Xie, Chao Dong, Gen Li, Ying Shan, and Ming-Ming Cheng. Vqfr: Blind face restoration with vector-quantized dictionary and parallel decoder. *arXiv preprint arXiv:2205.06803*, 2022. 3
- [20] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020. 3
- [21] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 16000–16009, 2022. 3
- [22] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 4
- [23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 33:6840–6851, 2020. 1, 7
- [24] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*, 2019. 4
- [25] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *arXiv preprint arXiv:2205.15868*, 2022. 3
- [26] Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Towards non-autoregressive language models. 2021. 3
- [27] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, pages 1501–1510, 2017. 3
- [28] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, pages 694–711. Springer, 2016. 3
- [29] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4401–4410, 2019. 1, 6, 7
- [30] 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. 1, 7
- [31] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [32] Taku Kudo and John Richardson. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*, 2018. 6

- [33] Adrian Łańcucki, Jan Chorowski, Guillaume Sanchez, Richard Marxer, Nanxin Chen, Hans JGA Dolfing, Sameer Khurana, Tanel Alumäe, and Antoine Laurent. Robust training of vector quantized bottleneck models. In *IEEE IJCNN*, pages 1–7. IEEE, 2020. 8
- [34] Doyup Lee, Chihyeon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Autoregressive image generation using residual quantization. In *CVPR*, pages 11523–11532, 2022. 2, 3, 6, 7
- [35] Xiaotong Li, Yixiao Ge, Kun Yi, Zixuan Hu, Ying Shan, and Ling-Yu Duan. mc-beit: Multi-choice discretization for image bert pre-training. *arXiv preprint arXiv:2203.15371*, 2022. 3
- [36] Yunfan Li, Peng Hu, Zitao Liu, Dezhong Peng, Joey Tianyi Zhou, and Xi Peng. Contrastive clustering. In *AAAI*, pages 8547–8555, 2021. 6
- [37] Charlie Nash, Jacob Menick, Sander Dieleman, and Peter W Battaglia. Generating images with sparse representations. *arXiv preprint arXiv:2103.03841*, 2021. 7
- [38] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *ICML*, pages 8162–8171. PMLR, 2021. 7
- [39] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *ICML*, pages 4055–4064. PMLR, 2018. 1
- [40] Zhiliang Peng, Li Dong, Hangbo Bao, Qixiang Ye, and Furu Wei. Beit v2: Masked image modeling with vector-quantized visual tokenizers. *arXiv preprint arXiv:2208.06366*, 2022. 3, 5, 6
- [41] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 3, 4
- [42] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. 3
- [43] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, pages 8821–8831. PMLR, 2021. 1, 3, 4
- [44] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *NeurIPS*, 32, 2019. 2, 3, 6, 7
- [45] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022. 1, 2, 3
- [46] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*, 2015. 6
- [47] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 5
- [48] Zhicong Tang, Shuyang Gu, Jianmin Bao, Dong Chen, and Fang Wen. Improved vector quantized diffusion models. *arXiv preprint arXiv:2205.16007*, 2022. 3
- [49] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. *NeurIPS*, 29, 2016. 3
- [50] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *NeurIPS*, 30, 2017. 1, 2, 4
- [51] Zhouxia Wang, Jiawei Zhang, Runjian Chen, Wenping Wang, and Ping Luo. Restoreformer: High-quality blind face restoration from undegraded key-value pairs. In *CVPR*, pages 17512–17521, 2022. 3
- [52] Ronald J Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280, 1989. 4
- [53] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *CVPR*, pages 9653–9663, 2022. 3
- [54] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. Videogpt: Video generation using vq-vae and transformers. *arXiv preprint arXiv:2104.10157*, 2021. 3
- [55] 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. 6, 7
- [56] Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong Xu, Jason Baldridge, and Yonghui Wu. Vector-quantized image modeling with improved vqgan. *arXiv preprint arXiv:2110.04627*, 2021. 1, 2, 3, 4, 6, 7, 8
- [57] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gungun 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 arXiv:2206.10789*, 2022. 1, 3
- [58] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018. 3, 5
- [59] Long Zhao, Zizhao Zhang, Ting Chen, Dimitris Metaxas, and Han Zhang. Improved transformer for high-resolution gans. *NeurIPS*, 34:18367–18380, 2021. 4, 6
- [60] Chuanxia Zheng, Long Tung Vuong, Jianfei Cai, and Dinh Phung. Movq: Modulating quantized vectors for high-fidelity image generation. *arXiv preprint arXiv:2209.09002*, 2022. 2, 3, 4, 6, 7
- [61] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021. 3, 5, 6
- [62] Shangchen Zhou, Kelvin CK Chan, Chongyi Li, and Chen Change Loy. Towards robust blind face restoration with codebook lookup transformer. *arXiv preprint arXiv:2206.11253*, 2022. 3