

Vector Quantized Diffusion Model for Text-to-Image Synthesis

Shuyang Gu¹ Dong Chen² Jianmin Bao^{2*} Fang Wen²
Bo Zhang² Dongdong Chen³ Lu Yuan³ Baining Guo²

¹University of Science and Technology of China

²Microsoft Research

³Microsoft Cloud+AI

gsy777@mail.ustc.edu.cn

{doch, jianbao, fangwen, zhanbo, dochen, luyuan, bainguo}@microsoft.com

Abstract

We present the vector quantized diffusion (VQ-Diffusion) model for text-to-image generation. This method is based on a vector quantized variational autoencoder (VQ-VAE) whose latent space is modeled by a conditional variant of the recently developed Denoising Diffusion Probabilistic Model (DDPM). We find that this latent-space method is well-suited for text-to-image generation tasks because it not only eliminates the unidirectional bias with existing methods but also allows us to incorporate a mask-and-replace diffusion strategy to avoid the accumulation of errors, which is a serious problem with existing methods. Our experiments show that the VQ-Diffusion produces significantly better text-to-image generation results when compared with conventional autoregressive (AR) models with similar numbers of parameters. Compared with previous GAN-based text-to-image methods, our VQ-Diffusion can handle more complex scenes and improve the synthesized image quality by a large margin. Finally, we show that the image generation computation in our method can be made highly efficient by reparameterization. With traditional AR methods, the text-to-image generation time increases linearly with the output image resolution and hence is quite time consuming even for normal size images. The VQ-Diffusion allows us to achieve a better trade-off between quality and speed. Our experiments indicate that the VQ-Diffusion model with the reparameterization is fifteen times faster than traditional AR methods while achieving a better image quality. The code and models are available at <https://github.com/cientgu/VQ-Diffusion>.

1. Introduction

Recent success of Transformer [11, 65] in neural language processing (NLP) has raised tremendous interest in using successful language models for computer vision tasks. Autoregressive (AR) model [4, 46, 47] is one of the most natural and popular approach to transfer from text-to-

text generation (*i.e.*, machine translation) to text-to-image generation. Based on the AR model, recent work DALL-E [48] has achieved impressive results for text-to-image generation.

Despite their success, existing text-to-image generation methods still have weaknesses that need to be improved. One issue is the unidirectional bias. Existing methods predict pixels or tokens in the reading order, from top-left to bottom-right, based on the attention to all prefix pixels/tokens and the text description. This fixed order introduces unnatural bias in the synthesized images because important contextual information may come from any part of the image, not just from left or above. Another issue is the accumulated prediction errors. Each step of the inference stage is performed based on previously sampled tokens – this is different from that of the training stage, which relies on the so-called “teacher-forcing” practice [15] and provides the ground truth for each step. This difference is important and its consequence merits careful examination. In particular, a token in the inference stage, once predicted, cannot be corrected and its errors will propagate to the subsequent tokens.

We present the vector quantized diffusion (VQ-Diffusion) model for text-to-image generation, a model that eliminates the unidirectional bias and avoids accumulated prediction errors. We start with a vector quantized variational autoencoder (VQ-VAE) and model its latent space by learning a parametric model using a conditional variant of the Denoising Diffusion Probabilistic Model (DDPM) [23, 59], which has been applied to image synthesis with compelling results [12]. We show that the latent-space model is well-suited for the task of text-to-image generation. Roughly speaking, the VQ-Diffusion model samples the data distribution by reversing a forward diffusion process that gradually corrupts the input via a fixed Markov chain. The forward process yields a sequence of increasingly noisy latent variables of the same dimensionality as the input, producing pure noise after a fixed number of timesteps. Starting from this noise result, the reverse process gradually denoises the latent variables to-

*Corresponding author.

wards the desired data distribution by learning the conditional transit distribution.

The VQ-Diffusion model eliminates the unidirectional bias. It consists of an independent text encoder and a diffusion image decoder, which performs denoising diffusion on discrete image tokens. At the beginning of the inference stage, all image tokens are either masked or random. Here the masked token serves the same function as those in mask-based generative models [11]. The denoising diffusion process gradually estimates the probability density of image tokens step-by-step based on the input text. In each step, the diffusion image decoder leverages the contextual information of all tokens of the entire image predicted in the previous step to estimate a new probability density distribution and use this distribution to predict the tokens in the current step. This bidirectional attention provides global context for each token prediction and eliminates the unidirectional bias.

The VQ-Diffusion model, with its mask-and-replace diffusion strategy, also avoids the accumulation of errors. In the training stage, we do not use the “teacher-forcing” strategy. Instead, we deliberately introduce both masked tokens and random tokens and let the network learn to predict the masked token and modify incorrect tokens. In the inference stage, we update the density distribution of all tokens in each step and resample all tokens according to the new distribution. Thus we can modify the wrong tokens and prevent error accumulation. Comparing to the conventional replace-only diffusion strategy for unconditional image generation [1], the masked tokens effectively direct the network’s attention to the masked areas and thus greatly reduce the number of token combinations to be examined by the network. This mask-and-replace diffusion strategy significantly accelerates the convergence of the network.

To assess the performance of the VQ-Diffusion method, we conduct text-to-image generation experiments with a wide variety of datasets, including CUB-200 [66], Oxford-102 [40], and MSCOCO [36]. Compared with AR model with similar numbers of model parameters, our method achieves significantly better results, as measured by both image quality metrics and visual examination, and is much faster. Compared with previous GAN-based text-to-image methods [67, 70, 71, 73], our method can handle more complex scenes and the synthesized image quality is improved by a large margin. Compared with extremely large models (models with ten times more parameters than ours), including DALL-E [48] and CogView [13], our model achieves comparable or better results for specific types of images, *i.e.*, the types of images that our model has seen during the training stage. Furthermore, our method is general and produces strong results in our experiments on both unconditional and conditional image generation with FFHQ [28] and ImageNet [10] datasets.

The VQ-Diffusion model also provides important benefits for the inference speed. With traditional AR methods, the inference time increases linearly with the output image resolution and the image generation is quite time consuming even for normal-size images (*e.g.*, images larger than small thumbnail images of 64×64 pixels). The VQ-Diffusion provides the global context for each token prediction and makes it independent of the image resolution. This allows us to provide an effective way to achieve a better tradeoff between the inference speed and the image quality by a simple reparameterization of the diffusion image decoder. Specifically, in each step, we ask the decoder to predict the original noise-free image instead of the noise-reduced image in the next denoising diffusion step. Through experiments we have found that the VQ-Diffusion method with reparameterization can be fifteen times faster than AR methods while achieving a better image quality.

2. Related Work

GAN-based Text-to-image generation. In the past few years, Generative Adversarial Networks (GANs) [18] have shown promising results on many tasks [19–21], especially text-to-image generation [5, 8, 9, 14, 17, 25, 27, 31–34, 38, 43, 44, 50, 51, 58, 60–63, 67–73]. GAN-INT-CLS [50] was the first to use a conditional GAN formulation for text-to-image generation. Based on this formulation, some approaches [34, 44, 67–71, 73] were proposed to further improve the generation quality. These models generate high fidelity images on single domain datasets, *e.g.*, birds [66] and flowers [40]. However, due to the inductive bias on the locality of convolutional neural networks, they struggle on complex scenes with multiple objects, such as those in the MS-COCO dataset [36].

Other works [25, 33] adopt a two-step process which first infer the semantic layout then generate different objects, but this kind of method requires fine-grained object labels, *e.g.*, object bounding boxes or segmentation maps.

Autoregressive Models. AR models [4, 46, 47] have shown powerful capability of density estimation and have been applied for image generation [7, 16, 41, 42, 49, 53, 64] recently. PixelRNN [53, 64], Image Transformer [42] and ImageGPT [7] factorized the probability density on an image over raw pixels. Thus, they only generate low-resolution images, like 64×64 , due to the unaffordable amount of computation for large images.

VQ-VAE [41, 49], VQGAN [16] and ImageBART [15] train an encoder to compress the image into a low-dimensional discrete latent space and fit the density of the hidden variables. It greatly improves the performance of image generation.

DALL-E [48], CogView [13] and M6 [35] propose AR-based text-to-image frameworks. They model the joint distribution of text and image tokens. With powerful large

transformer structure and massive text-image pairs, they greatly advance the quality of text-to-image generation, but still have weaknesses of unidirectional bias and accumulated prediction errors due to the limitation of AR models.

Denoising Diffusion Probabilistic Models. Diffusion generative models were first proposed in [59] and achieved strong results on image generation [12,23,24,39] and image super super-resolution [52] recently. However, most previous works only considered continuous diffusion models on the raw image pixels. Discrete diffusion models were also first described in [59], and then applied to text generation in Argmax Flow [26]. D3PMs [1] applies discrete diffusion to image generation. However, it also estimates the density of raw image pixels and can only generate low-resolution (e.g., 32×32) images.

3. Background: Learning Discrete Latent Space of Images Via VQ-VAE

Transformer architectures have shown great promise in image synthesis due to their outstanding expressivity [7, 16, 48]. In this work, we aim to leverage the transformer to learn the mapping from text to image. Since the computation cost is quadratic to the sequence length, it is computationally prohibitive to directly model raw pixels using transformers. To address this issue, recent works [16, 41] propose to represent an image by discrete image tokens with reduced sequence length. Hereafter a transformer can be effectively trained upon this reduced context length and learn the translation from the text to image tokens.

Formally, a *vector quantized variational autoencoder* (VQ-VAE) [41] is employed. The model consists of an encoder E , a decoder G and a codebook $\mathcal{Z} = \{z_k\}_{k=1}^K \in \mathbb{R}^{K \times d}$ containing a finite number of embedding vectors, where K is the size of the codebook and d is the dimension of codes. Given an image $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$, we obtain a spatial collection of image tokens z_q with the encoder $z = E(\mathbf{x}) \in \mathbb{R}^{h \times w \times d}$ and a subsequent spatial-wise quantizer $Q(\cdot)$ which maps each spatial feature z_{ij} into its closest codebook entry z_k :

$$z_q = Q(z) = \left(\underset{z_k \in \mathcal{Z}}{\operatorname{argmin}} \|z_{ij} - z_k\|_2^2 \right) \in \mathbb{R}^{h \times w \times d} \quad (1)$$

Where $h \times w$ represents the encoded sequence length and is usually much smaller than $H \times W$. Then the image can be faithfully reconstructed via the decoder, *i.e.*, $\tilde{\mathbf{x}} = G(z_q)$. Hence, image synthesis is equivalent to sampling image tokens from the latent distribution. Note that the image tokens are quantized latent variables in the sense that they take discrete values. The encoder E , the decoder G and the codebook \mathcal{Z} can be trained end-to-end via the following loss function:

$$\mathcal{L}_{\text{VQVAE}} = \|\mathbf{x} - \tilde{\mathbf{x}}\|_1 + \|\operatorname{sg}[E(\mathbf{x})] - z_q\|_2^2 + \beta \|\operatorname{sg}[z_q] - E(\mathbf{x})\|_2^2 \quad (2)$$

Where, $\operatorname{sg}[\cdot]$ stands for the stop-gradient operation. In practice, we replace the second term of Equation 2 with exponential moving averages (EMA) [41] to update the codebook entries which is proven to work better than directly using the loss function.

4. Vector Quantized Diffusion Model

Given the text-image pairs, we obtain the discrete image tokens $\mathbf{x} \in \mathbb{Z}^N$ with a pretrained VQ-VAE, where $N = hw$ represents the sequence length of tokens. Suppose the size of the VQ-VAE codebook is K , the image token x^i at location i takes the index that specifies the entries in the codebook, *i.e.*, $x^i \in \{1, 2, \dots, K\}$. On the other hand, the text tokens $\mathbf{y} \in \mathbb{Z}^M$ can be obtained through BPE-encoding [56]. The overall text-to-image framework can be viewed as maximizing the conditional transition distribution $q(\mathbf{x}|\mathbf{y})$.

Previous autoregressive models, *e.g.*, DALL-E [48] and CogView [13], sequentially predict each image token depends on the text tokens as well as the previously predicted image tokens, *i.e.*, $q(\mathbf{x}|\mathbf{y}) = \prod_{i=1}^N q(x^i|x^1, \dots, x^{i-1}, \mathbf{y})$. While achieving remarkable quality in text-to-image synthesis, there exist several limitations of autoregressive modeling. First, image tokens are predicted in a unidirectional ordering, *e.g.*, raster scan, which neglects the structure of 2D data and restricts the expressivity for image modeling since the prediction of a specific location should not merely attend to the context on the left or the above. Second, there is a train-test discrepancy as the training employs ground truth whereas the inference relies on the prediction as previous tokens. The so-called “teacher-forcing” practice [15] or exposure bias [54] leads to error accumulation due to the mistakes in the earlier sampling. Moreover, it requires a forward pass of the network to predict each token, which consumes an inordinate amount of time even for the sampling in the latent space of low resolution (*i.e.*, 32×32), making the AR model impractical for real usage.

We aim to model the VQ-VAE latent space in a *non-autoregressive* manner. The proposed VQ-Diffusion method maximizes the probability $q(\mathbf{x}|\mathbf{y})$ with the diffusion model [23, 59], an emerging approach that produces compelling quality on image synthesis [12]. While the majority of recent works focus on continuous diffusion models, using them for categorical distribution is much less researched [1, 26]. In this work, we propose to use its conditional variant discrete diffusion process for text-to-image generation. We will subsequently introduce the discrete diffusion process inspired by the masked language modeling (MLM) [11], and then discuss how to train a neural network to reverse this process.

4.1. Discrete diffusion process

On a high level, the *forward* diffusion process gradually corrupts the image data \mathbf{x}_0 via a fixed Markov chain

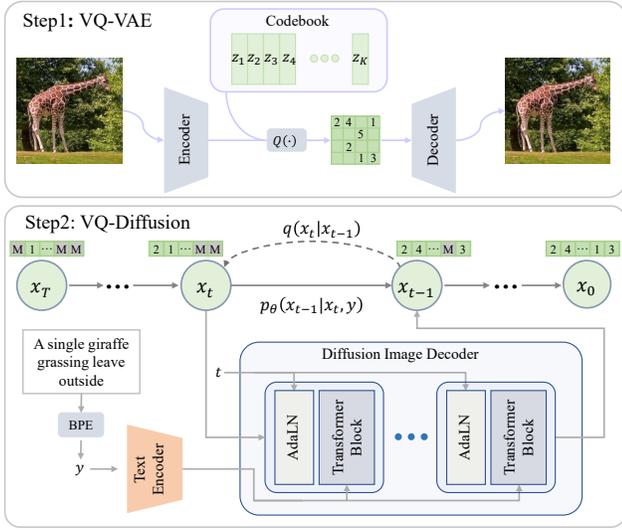


Figure 1. Overall framework of our method. It starts with the VQ-VAE. Then, the VQ-Diffusion models the discrete latent space by reversing a forward diffusion process that gradually corrupts the input via a fixed Markov chain.

$q(x_t|x_{t-1})$, e.g., random replace some tokens of x_{t-1} . After a fixed number of T timesteps, the forward process yields a sequence of increasingly noisy latent variables z_1, \dots, z_T of the same dimensionality as z_0 , and z_T becomes pure noise tokens. Starting from the noise z_T , the *reverse* process gradually denoises the latent variables and restore the real data x_0 by sampling from the reverse distribution $q(x_{t-1}|x_t, x_0)$ sequentially. However, since x_0 is unknown in the inference stage, we train a transformer network to approximate the conditional transit distribution $p_\theta(x_{t-1}|x_t, y)$ depends on the entire data distribution.

To be more specific, consider a single image token x_0^i of x_0 at location i , which takes the index that specifies the entries in the codebook, i.e., $x_0^i \in \{1, 2, \dots, K\}$. Without introducing confusion, we omit superscripts i in the following description. We define the probabilities that x_{t-1} transits to x_t using the matrices $[Q_t]_{mn} = q(x_t = m|x_{t-1} = n) \in \mathbb{R}^{K \times K}$. Then the *forward* Markov diffusion process for the whole token sequence can be written as,

$$q(x_t|x_{t-1}) = \mathbf{v}^\top(x_t) \mathbf{Q}_t \mathbf{v}(x_{t-1}) \quad (3)$$

where $\mathbf{v}(x)$ is a one-hot column vector, the length is number of tokens and only the entry x is 1. The categorical distribution over x_t is given by the vector $\mathbf{Q}_t \mathbf{v}(x_{t-1})$.

Importantly, due to the property of Markov chain, one can marginalize out the intermediate steps and derive the probability of x_t at arbitrary timestep directly from x_0 as,

$$q(x_t|x_0) = \mathbf{v}^\top(x_t) \overline{\mathbf{Q}}_t \mathbf{v}(x_0), \text{ with } \overline{\mathbf{Q}}_t = \mathbf{Q}_t \cdots \mathbf{Q}_1. \quad (4)$$

Besides, another notable characteristic is that by conditioning on z_0 , the posterior of this diffusion process is

tractable, i.e.,

$$\begin{aligned} q(x_{t-1}|x_t, x_0) &= \frac{q(x_t|x_{t-1}, x_0)q(x_{t-1}|x_0)}{q(x_t|x_0)} \\ &= \frac{(\mathbf{v}^\top(x_t) \mathbf{Q}_t \mathbf{v}(x_{t-1})) (\mathbf{v}^\top(x_{t-1}) \overline{\mathbf{Q}}_{t-1} \mathbf{v}(x_0))}{\mathbf{v}^\top(x_t) \overline{\mathbf{Q}}_t \mathbf{v}(x_0)}. \end{aligned} \quad (5)$$

The transition matrix \mathbf{Q}_t is crucial to the discrete diffusion model and should be carefully designed such that it is not too difficult for the reverse network to recover the signal from noises.

Previous works [1, 26] propose to introduce a small amount of uniform noises to the categorical distribution and the transition matrix can be formulated as,

$$\mathbf{Q}_t = \begin{bmatrix} \alpha_t + \beta_t & \beta_t & \cdots & \beta_t \\ \beta_t & \alpha_t + \beta_t & \cdots & \beta_t \\ \vdots & \vdots & \ddots & \vdots \\ \beta_t & \beta_t & \cdots & \alpha_t + \beta_t \end{bmatrix} \quad (6)$$

with $\alpha_t \in [0, 1]$ and $\beta_t = (1 - \alpha_t)/K$. Each token has a probability of $(\alpha_t + \beta_t)$ to remain the previous value at the current step while with a probability of $K\beta_t$ to be resampled uniformly over all the K categories.

Nonetheless, the data corruption using uniform diffusion is a somewhat aggressive process that may pose challenge for the reverse estimation. First, as opposed to the Gaussian diffusion process for ordinal data, an image token may be replaced to an utterly uncorrelated category, which leads to an abrupt semantic change for that token. Second, the network has to take extra efforts to figure out the tokens that have been replaced prior to fixing them. In fact, due to the semantic conflict within the local context, the reverse estimation for different image tokens may form a competition and run into the dilemma of identifying the reliable tokens.

Mask-and-replace diffusion strategy. To solve the above issues of uniform diffusion, we draw inspiration from mask language modeling [11] and propose to corrupt the tokens by stochastically masking some of them so that the corrupted locations can be explicitly known by the reverse network. Specifically, we introduce an additional special token, [MASK] token, so each token now has $(K + 1)$ discrete states. We define the mask diffusion as follows: each ordinary token has a probability of γ_t to be replaced by the [MASK] token and has a chance of $K\beta_t$ to be uniformly diffused, leaving the probability of $\alpha_t = 1 - K\beta_t - \gamma_t$ to be unchanged, whereas the [MASK] token always keeps its own state. Hence, we can formulate the transition matrix $\mathbf{Q}_t \in \mathbb{R}^{(K+1) \times (K+1)}$ as,

$$\mathbf{Q}_t = \begin{bmatrix} \alpha_t + \beta_t & \beta_t & \beta_t & \cdots & 0 \\ \beta_t & \alpha_t + \beta_t & \beta_t & \cdots & 0 \\ \beta_t & \beta_t & \alpha_t + \beta_t & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \gamma_t & \gamma_t & \gamma_t & \cdots & 1 \end{bmatrix}. \quad (7)$$

Algorithm 1 Training of the VQ-Diffusion, given transition matrix $\{Q_t\}$, initial network parameters θ , loss weight λ , learning rate η .

- 1: **repeat**
 - 2: $(I, s) \leftarrow$ sample training image-text pair
 - 3: $x_0 \leftarrow$ VQVAE-Encoder(I), $\mathbf{y} \leftarrow$ BPE(s)
 - 4: $t \sim$ Uniform($\{1, \dots, T\}$)
 - 5: $x_t \leftarrow$ sample from $q(x_t|x_0)$ \triangleright Eqn. 4 and 8
 - 6: $\mathcal{L} \leftarrow \begin{cases} \mathcal{L}_0, & \text{if } t = 1 \\ \mathcal{L}_{t-1} + \lambda \mathcal{L}_{x_0}, & \text{otherwise} \end{cases}$ \triangleright Eqn. 9 and 12
 - 7: $\theta \leftarrow \theta - \eta \nabla_{\theta} \mathcal{L}$ \triangleright Update network parameters
 - 8: **until** converged
-

The benefit of this mask-and-replace transition is that: 1) the corrupted tokens are distinguishable to the network, which eases the reverse process. 2) Comparing to the mask only approach in [1], we theoretically prove that it is necessary to include a small amount of uniform noises besides the token masking, otherwise we get a trivial posterior when $x_t \neq x_0$. 3) The random token replacement forces the network to understand the context rather than only focusing on the [MASK] tokens. 4) The cumulative transition matrix \bar{Q}_t and the probability $q(x_t|x_0)$ in Equation 4 can be computed in closed form with:

$$\bar{Q}_t \mathbf{v}(x_0) = \bar{\alpha}_t \mathbf{v}(x_0) + (\bar{\gamma}_t - \bar{\beta}_t) \mathbf{v}(K+1) + \bar{\beta}_t \quad (8)$$

Where $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$, $\bar{\gamma}_t = 1 - \prod_{i=1}^t (1 - \gamma_i)$, and $\bar{\beta}_t = (1 - \bar{\alpha}_t - \bar{\gamma}_t)/K$ can be calculated and stored in advance. Thus, the computation cost of $q(x_t|x_0)$ is reduced from $O(tK^2)$ to $O(K)$. The proof is given in the appendix.

4.2. Learning the reverse process

To reverse the diffusion process, we train a denoising network $p_{\theta}(x_{t-1}|x_t, \mathbf{y})$ to estimate the posterior transition distribution $q(x_{t-1}|x_t, x_0)$. The network is trained to minimize the variational lower bound (VLB) [59]:

$$\begin{aligned} \mathcal{L}_{\text{vlb}} &= \mathcal{L}_0 + \mathcal{L}_1 + \dots + \mathcal{L}_{T-1} + \mathcal{L}_T, \\ \mathcal{L}_0 &= -\log p_{\theta}(x_0|x_1, \mathbf{y}), \\ \mathcal{L}_{t-1} &= D_{KL}(q(x_{t-1}|x_t, x_0) \parallel p_{\theta}(x_{t-1}|x_t, \mathbf{y})), \\ \mathcal{L}_T &= D_{KL}(q(x_T|x_0) \parallel p(x_T)). \end{aligned} \quad (9)$$

Where $p(x_T)$ is the prior distribution of timestep T . For the proposed mask-and-replace diffusion, the prior is:

$$p(x_T) = [\bar{\beta}_T, \bar{\beta}_T, \dots, \bar{\beta}_T, \bar{\gamma}_T]^{\top} \quad (10)$$

Note that since the transition matrix Q_t is fixed in the training, the \mathcal{L}_T is a constant number which measures the gap between the training and inference and can be ignored in the training.

Algorithm 2 Inference of the VQ-Diffusion, given fast inference time stride Δ_t , input text s .

- 1: $t \leftarrow T$, $\mathbf{y} \leftarrow$ BPE(s)
 - 2: $x_t \leftarrow$ sample from $p(x_T)$ \triangleright Eqn. 10
 - 3: **while** $t > 0$ **do**
 - 4: $x_t \leftarrow$ sample from $p_{\theta}(x_{t-\Delta_t}|x_t, \mathbf{y})$ \triangleright Eqn. 13
 - 5: $t \leftarrow t - \Delta_t$
 - 6: **end while**
 - 7: **return** VQVAE-Decoder(x_t)
-

Reparameterization trick on discrete stage. The network parameterization affects the synthesis quality significantly. Instead of directly predicting the posterior $q(x_{t-1}|x_t, x_0)$, recent works [1, 23, 26] find that approximating some surrogate variables, *e.g.*, the noiseless target data $q(x_0)$ gives better quality. In the discrete setting, we let the network predict the noiseless token distribution $p_{\theta}(\tilde{x}_0|x_t, \mathbf{y})$ at each reverse step. We can thus compute the reverse transition distribution according to:

$$p_{\theta}(x_{t-1}|x_t, \mathbf{y}) = \sum_{\tilde{x}_0=1}^K q(x_{t-1}|x_t, \tilde{x}_0) p_{\theta}(\tilde{x}_0|x_t, \mathbf{y}). \quad (11)$$

Based on the reparameterization trick, we can introduce an auxiliary denoising objective, which encourages the network to predict noiseless token x_0 :

$$\mathcal{L}_{x_0} = -\log p_{\theta}(x_0|x_t, \mathbf{y}) \quad (12)$$

We find that combining this loss with \mathcal{L}_{vlb} improves the image quality.

Model architecture. We propose an encoder-decoder transformer to estimate the distribution $p_{\theta}(\tilde{x}_0|x_t, \mathbf{y})$. As shown in Figure 1, the framework contains two parts: a text encoder and a diffusion image decoder. Our text encoder takes the text tokens \mathbf{y} and yields a conditional feature sequence. The diffusion image decoder takes the image token x_t and timestep t and outputs the noiseless token distribution $p_{\theta}(\tilde{x}_0|x_t, \mathbf{y})$. The decoder contains several transformer blocks and a softmax layer. Each transformer block contains a full attention, a cross attention to combine text information and a feed forward network block. The current timestep t is injected into the network with Adaptive Layer Normalization [2](AdaLN) operator, *i.e.*, $\text{AdaLN}(h, t) = a_t \text{LayerNorm}(h) + b_t$, where h is the intermediate activations, a_t and b_t are obtained from a linear projection of the timestep embedding.

Fast inference strategy In the inference stage, by leveraging the reparameterization trick, we can skip some steps in diffusion model to achieve a faster inference.

Specifically, assuming the time stride is Δ_t , instead of sampling images in the chain of $x_T, x_{T-1}, x_{T-2} \dots x_0$, we sample images in the chain of $x_T, x_{T-\Delta_t}, x_{T-2\Delta_t} \dots x_0$

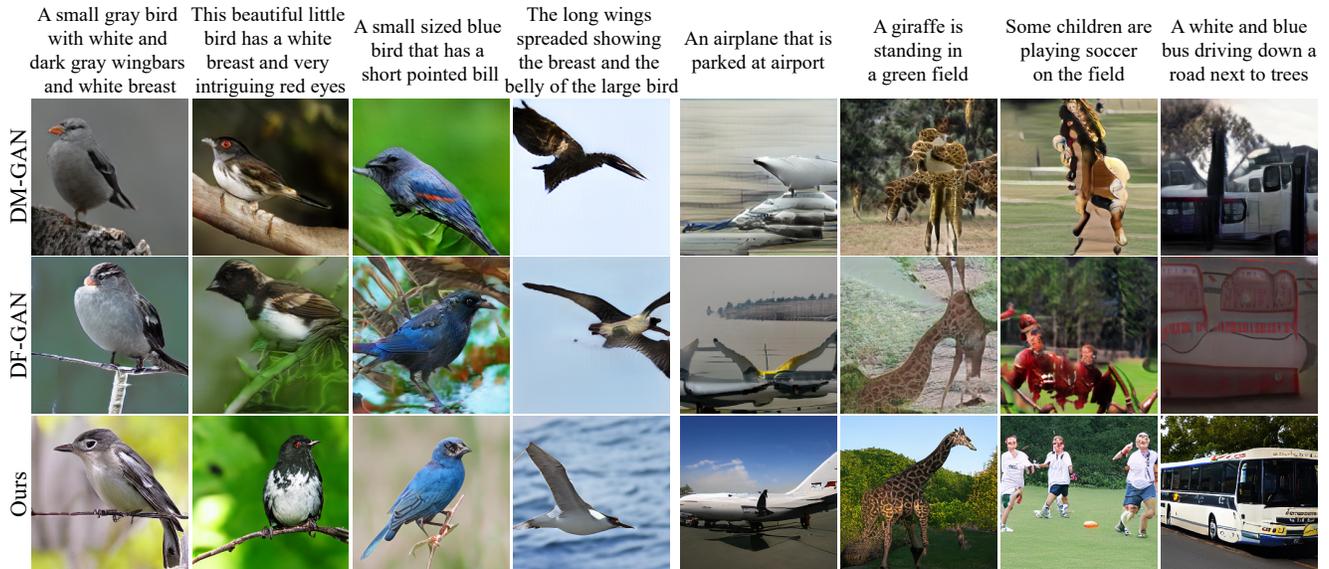


Figure 2. Comparison with GAN-based method on CUB-200 and MSCOCO datasets.

with the reverse transition distribution:

$$p_{\theta}(\mathbf{x}_{t-\Delta_t}|\mathbf{x}_t, \mathbf{y}) = \sum_{\tilde{\mathbf{x}}_0=1}^K q(\mathbf{x}_{t-\Delta_t}|\mathbf{x}_t, \tilde{\mathbf{x}}_0)p_{\theta}(\tilde{\mathbf{x}}_0|\mathbf{x}_t, \mathbf{y}). \quad (13)$$

We found it makes the sampling more efficient which only causes little harm to quality. The whole training and inference algorithm is shown in Algorithm 1 and 2.

5. Experiments

In this section, we first introduce the overall experiment setups and then present extensive results to demonstrate the superiority of our approach in text-to-image synthesis. Finally, we point out that our method is a general image synthesis framework that achieves great performance on other generation tasks, including unconditional and class conditional image synthesis.

Datasets. To demonstrate the capability of our proposed method for text-to-image synthesis, we conduct experiments on CUB-200 [66], Oxford-102 [40], and MSCOCO [36] datasets. The CUB-200 dataset contains 8855 training images and 2933 test images belonging to 200 bird species. Oxford-102 dataset contains 8189 images of flowers of 102 categories. Each image in CUB-200 and Oxford-102 dataset contains 10 text descriptions. MSCOCO dataset contains 82k images for training and 40k images for testing. Each image in this dataset has five text descriptions.

To further demonstrate the scalability of our method, we also train our model on large scale datasets, including Conceptual Captions [6, 57] and LAION-400M [55]. The Conceptual Caption dataset, including both CC3M [57] and CC12M [6] datasets, contains 15M images. To balance the

text and image distribution, we filter a 7M subset according to the word frequency. The LAION-400M dataset contains 400M image-text pairs. We train our model on three subsets from LAION, i.e., cartoon, icon, and human, each of them contains 0.9M, 1.3M, 42M images, respectively. For each subset, we filter the data according to the text.

Traning Details. Our VQ-VAE’s encoder and decoder follow the setting of VQGAN [16] which leverages the GAN loss to get a more realistic image. We directly adopt the publicly available VQGAN model trained on Open-Images [30] dataset for all text-to-image synthesis experiments. It converts 256×256 images into 32×32 tokens. The codebook size $K = 2886$ after removing useless codes. We adopt a publicly available tokenizer of the CLIP model [45] as text encoder, yielding a conditional sequence of length 77. We fix both image and text encoders in our training.

For fair comparison with previous text-to-image methods under similar parameters, we build two different diffusion image decoder settings: 1) **VQ-Diffusion-S** (Small), it contains 18 transformer blocks with dimension of 192. The model contains 34M parameters. 2) **VQ-Diffusion-B** (Base), it contains 19 transformer blocks with dimension of 1024. The model contains 370M parameters.

In order to show the scalability of our method, we also train our base model on a larger database Conceptual Captions, and then fine-tune it on each database. This model is denoted as **VQ-Diffusion-F**.

For the default setting, we set timesteps $T = 100$ and loss weight $\lambda = 0.0005$. For the transition matrix, we linearly increase $\bar{\gamma}_t$ and $\bar{\beta}_t$ from 0 to 0.9 and 0.1, respectively. We optimize our network using AdamW [37] with $\beta_1 = 0.9$ and $\beta_2 = 0.96$. The learning rate is set to 0.00045 after 5000 iterations of warmup. More training details are provided in the appendix.

5.1. Comparison with state-of-the-art methods

We qualitatively compare the proposed method with several state-of-the-art methods, including some GAN-based methods [60, 61, 63, 67, 70, 71, 73], DALL-E [48] and CogView [13], on MSCOCO, CUB-200 and Oxford-102 datasets. We calculate the FID [22] between 30k generated images and 30k real images, and show the results in Table 1.

We can see that our small model, VQ-Diffusion-S, which has the similar parameter number with previous GAN-based models, has strong performance on CUB-200 and Oxford-102 datasets. Our base model, VQ-Diffusion-B, further improves the performance. And our VQ-Diffusion-F model achieves the best results and surpasses all previous methods by a large margin, even surpassing DALL-E [48] and CogView [13], which have ten times more parameters than ours, on MSCOCO dataset.

Some visualized comparison results with DM-GAN [73] and DF-GAN [63] are shown in Figure 2. Obviously, our synthesized images have better realistic fine-grained details and are more consistent with the input text.

5.2. In the wild text-to-image synthesis

To demonstrate the capability of generating in-the-wild images, we train our model on three subsets from LAION-400M dataset, e.g., cartoon, icon and human. We provide our results here in Figure 3. Though our base model is much smaller than previous works like DALL-E and CogView, we also achieved a strong performance.

Compared with the AR method which generates images from top-left to down-right, our method generates images in a global manner. It makes our method can be applied to many vision tasks, e.g., irregular mask inpainting. For this task, we do not need to re-train a new model. We simply set the tokens in the irregular region as [MASK] token, and send them to our model. This strategy supports both unconditional mask inpainting and text conditional mask inpainting. We show these results in the appendix.

5.3. Ablations

Number of timesteps. We investigate the timesteps in training and inference. As shown in Table 2, we perform the experiment on the CUB-200 dataset. We find when the training steps increase from 10 to 100, the result improves, when it further increase to 200, it seems saturated. So we set the default timesteps number to 100 in our experiments. To demonstrate the fast inference strategy, we evaluate the generated images from 10, 25, 50, 100 inference steps on five models with different training steps. We find it still maintains a good performance when dropping 3/4 inference steps, which may save about 3/4 inference times.

Mask-and-replace diffusion strategy. We explore how the mask-and-replace strategy benefits our performance on the

	MSCOCO	CUB-200	Oxford-102
StackGAN [70]	74.05	51.89	55.28
StackGAN++ [71]	81.59	15.30	48.68
EFF-T2I [60]	-	11.17	16.47
SEGAN [61]	32.28	18.17	-
AttnGAN [67]	35.49	23.98	-
DM-GAN [73]	32.64	16.09	-
DF-GAN [63]	21.42	14.81	-
DAE-GAN [51]	28.12	15.19	-
DALLE [48]	27.50	56.10	-
Cogview [13]	27.10	-	-
VQ-Diffusion-S	30.17	12.97	14.95
VQ-Diffusion-B	19.75	11.94	14.88
VQ-Diffusion-F	13.86	10.32	14.10

Table 1. FID comparison of different text-to-image synthesis method on MSCOCO, CUB-200, and Oxford-102 datasets.



Figure 3. In the wild text-to-image synthesis results.

Oxford-102 dataset. We set different final mask rate ($\bar{\gamma}_T$) to investigate the effect. Both mask only strategy ($\bar{\gamma}_T = 1$) and replace only strategy ($\bar{\gamma}_T = 0$) are special cases of our mask-and-replace strategy. From Figure 4, we find it get the best performance when $M = 0.9$. When $M > 0.9$, it may suffer from the error accumulation problem, when $M < 0.9$, the network may be difficult to find which region needs to pay more attention.

Truncation. We also demonstrate that the truncation sampling strategy is extremely important for our discrete diffusion based method. It may avoid the network sampling from low probability tokens. Specifically, we only keep top r tokens of $p_\theta(\hat{x}_0|x_t, y)$ in the inference stage. We evaluate the results with different truncation rates r on CUB-200

inference steps	training steps					
		10	25	50	100	200
10		32.35	27.62	23.47	19.84	20.96
25		-	18.53	15.25	14.03	16.13
50		-	-	13.82	12.45	13.67
100		-	-	-	11.94	12.27
200		-	-	-	-	11.80

Table 2. Ablation study on training steps and inference steps. Each column shares the same training steps while each row shares the same inference steps.

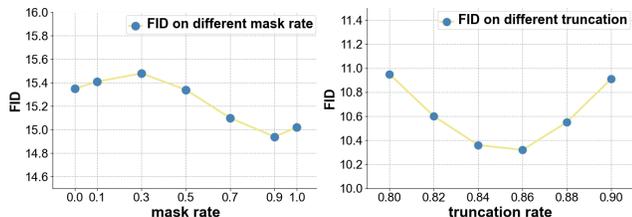


Figure 4. Ablation study on the mask rate and the truncation rate.

Model	steps	FID	throughput(imgs/s)
VQ-AR-S		18.12	0.08
VQ-Diffusion-S	25	15.46	1.25
VQ-Diffusion-S	50	13.62	0.67
VQ-Diffusion-S	100	12.97	0.37
VQ-AR-B		17.76	0.03
VQ-Diffusion-B	25	14.03	0.47
VQ-Diffusion-B	50	12.45	0.24
VQ-Diffusion-B	100	11.94	0.13

Table 3. Comparison between VQ-Diffusion and VQ-AR models. By changing the inference steps, the VQ-Diffusion model is 15 times faster than the VQ-AR model while maintaining better performance.

dataset. As shown in Figure 4, we find that it achieves the best performance when the truncation rate equals 0.86.

VQ-Diffusion vs VQ-AR. For a fair comparison, we replace our diffusion image decoder with an autoregressive decoder with the same network structure and keep other settings the same, including both image and text encoders. The autoregressive model is denoted as VQ-AR-S and VQ-AR-B, corresponding to VQ-Diffusion-S and VQ-Diffusion-B. The experiment is performed on the CUB-200 dataset. As shown in Table 3, on both -S and -B settings the VQ-Diffusion model surpasses the VQ-AR model by a large margin. Meanwhile, we evaluate the throughput of both methods on a V100 GPU with a batch size of 32. The VQ-Diffusion with the fast inference strategy is 15 times faster than the VQ-AR model with a better FID score.

5.4. Unified generation model

Our method is general, which can also be applied to other image synthesis tasks, *e.g.*, unconditional image syn-

Model	ImageNet	FFHQ
StyleGAN2 [29]	-	3.8
BigGAN [3]	7.53	12.4
BigGAN-deep [3]	6.84	-
IDDPM [39]	12.3	-
ADM-G [12]	10.94	-
VQGAN [16]	15.78	9.6
ImageBART [15]	21.19	9.57
Ours	11.89	6.33
ADM-G (1.0guid) [12]	4.59	-
VQGAN (acc0.05) [16]	5.88	-
ImageBART (acc0.05) [15]	7.44	-
Ours (acc0.05)	5.32	-

Table 4. FID score comparison for class-conditional synthesis on ImageNet, and unconditional synthesis on FFHQ dataset. 'guid' denotes using classifier guidance [12], 'acc' denotes adopting acceptance rate [16].

thesis and image synthesis conditioned on labels. To generate images from a given class label, we first remove the text encoder network and cross attention part in transformer blocks, and inject the class label through the AdaLN operator. Our network contains 24 transformer blocks with dimension 512. We train our model on the ImageNet dataset. For VQ-VAE, we adopt the publicly available model from VQ-GAN [16] trained on ImageNet dataset, which down-samples images from 256×256 to 16×16 . For unconditional image synthesis, we trained our model on the FFHQ256 dataset, which contains 70k high quality face images. The image encoder also downsamples images to 16×16 tokens.

We assess the performance of our model in terms of FID and compare with a variety of previously established models [3, 12, 15, 16, 39]. For a fair comparison, we calculate FID between 50k generated images and all real images. Following [16] we can further increase the quality by only accepting images with a top 5% classification score, denoted as acc0.05. We show the quantitative results in Table 4. While some task-specialized GAN models report better FID scores, our approach provides a unified model that works well across a wide range of tasks.

6. Conclusion

In this paper, we present a novel text-to-image architecture named VQ-Diffusion. The core design is to model the VQ-VAE latent space in a non-autoregressive manner. The proposed mask-and-replace diffusion strategy avoids the accumulation of errors of the AR model. Our model has the capacity to generate more complex scenes, which surpasses previous GAN-based text-to-image methods. Our method is also general and produces strong results on unconditional and conditional image generation.

References

- [1] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. *arXiv preprint arXiv:2107.03006*, 2021. [2](#), [3](#), [4](#), [5](#)
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. [5](#)
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. [8](#)
- [4] 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. [1](#), [2](#)
- [5] Miriam Cha, Youngjune L Gwon, and HT Kung. Adversarial learning of semantic relevance in text to image synthesis. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 3272–3279, 2019. [2](#)
- [6] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3558–3568, 2021. [6](#)
- [7] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pre-training from pixels. In *International Conference on Machine Learning*, pages 1691–1703. PMLR, 2020. [2](#), [3](#)
- [8] Jun Cheng, Fuxiang Wu, Yanling Tian, Lei Wang, and Dapeng Tao. Rifegan: Rich feature generation for text-to-image synthesis from prior knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10911–10920, 2020. [2](#)
- [9] Ayushman Dash, John Cristian Borges Gamboa, Sheraz Ahmed, Marcus Liwicki, and Muhammad Zeshan Afzal. Tac-gan-text conditioned auxiliary classifier generative adversarial network. *arXiv preprint arXiv:1703.06412*, 2017. [2](#)
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [2](#)
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. [1](#), [2](#), [3](#), [4](#)
- [12] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis. *arXiv preprint arXiv:2105.05233*, 2021. [1](#), [3](#), [8](#)
- [13] Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, et al. Cogview: Mastering text-to-image generation via transformers. *arXiv preprint arXiv:2105.13290*, 2021. [2](#), [3](#), [7](#)
- [14] Alaaeldin El-Nouby, Shikhar Sharma, Hannes Schulz, Devon Hjelm, Layla El Asri, Samira Ebrahimi Kahou, Yoshua Bengio, and Graham W Taylor. Tell, draw, and repeat: Generating and modifying images based on continual linguistic instruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10304–10312, 2019. [2](#)
- [15] Patrick Esser, Robin Rombach, Andreas Blattmann, and Björn Ommer. Imagebart: Bidirectional context with multinomial diffusion for autoregressive image synthesis. *arXiv preprint arXiv:2108.08827*, 2021. [1](#), [2](#), [3](#), [8](#)
- [16] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12873–12883, 2021. [2](#), [3](#), [6](#), [8](#)
- [17] Lianli Gao, Daiyuan Chen, Jingkuan Song, Xing Xu, Dongxiang Zhang, and Heng Tao Shen. Perceptual pyramid adversarial networks for text-to-image synthesis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8312–8319, 2019. [2](#)
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014. [2](#)
- [19] Shuyang Gu, Jianmin Bao, Dong Chen, and Fang Wen. Giga: Generated image quality assessment. In *European Conference on Computer Vision*, pages 369–385. Springer, 2020. [2](#)
- [20] Shuyang Gu, Jianmin Bao, Dong Chen, and Fang Wen. Priorgan: Real data prior for generative adversarial nets. *arXiv preprint arXiv:2006.16990*, 2020. [2](#)
- [21] Shuyang Gu, Jianmin Bao, Hao Yang, Dong Chen, Fang Wen, and Lu Yuan. Mask-guided portrait editing with conditional gans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3436–3445, 2019. [2](#)
- [22] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. [7](#)
- [23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *arXiv preprint arXiv:2006.11239*, 2020. [1](#), [3](#), [5](#)
- [24] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *arXiv preprint arXiv:2106.15282*, 2021. [3](#)
- [25] Seunghoon Hong, Dingdong Yang, Jongwook Choi, and Honglak Lee. Inferring semantic layout for hierarchical text-to-image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7986–7994, 2018. [2](#)
- [26] Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Towards non-autoregressive language models. *arXiv preprint arXiv:2102.05379*, 2021. [3](#), [4](#), [5](#)

- [27] Yupan Huang, Hongwei Xue, Bei Liu, and Yutong Lu. Unifying multimodal transformer for bi-directional image and text generation. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 1138–1147, 2021. 2
- [28] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. 2
- [29] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. 8
- [30] Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Andreas Veit, et al. Openimages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from <https://github.com/openimages>*, 2(3):18, 2017. 6
- [31] Qicheng Lao, Mohammad Havaei, Ahmad Pesaranghader, Francis Dutil, Lisa Di Jorio, and Thomas Fevens. Dual adversarial inference for text-to-image synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7567–7576, 2019. 2
- [32] Bowen Li, Xiaojuan Qi, Thomas Lukasiewicz, and Philip HS Torr. Controllable text-to-image generation. *arXiv preprint arXiv:1909.07083*, 2019. 2
- [33] Wenbo Li, Pengchuan Zhang, Lei Zhang, Qiuyuan Huang, Xiaodong He, Siwei Lyu, and Jianfeng Gao. Object-driven text-to-image synthesis via adversarial training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12174–12182, 2019. 2
- [34] Jiadong Liang, Wenjie Pei, and Feng Lu. Cpgan: Content-parsing generative adversarial networks for text-to-image synthesis. In *European Conference on Computer Vision*, pages 491–508. Springer, 2020. 2
- [35] Junyang Lin, Rui Men, An Yang, Chang Zhou, Ming Ding, Yichang Zhang, Peng Wang, Ang Wang, Le Jiang, Xianyan Jia, et al. M6: A chinese multimodal pretrainer. *arXiv preprint arXiv:2103.00823*, 2021. 2
- [36] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 2, 6
- [37] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [38] Anh Nguyen, Jeff Clune, Yoshua Bengio, Alexey Dosovitskiy, and Jason Yosinski. Plug & play generative networks: Conditional iterative generation of images in latent space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4467–4477, 2017. 2
- [39] Alex Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. *arXiv preprint arXiv:2102.09672*, 2021. 3, 8
- [40] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729. IEEE, 2008. 2, 6
- [41] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. *arXiv preprint arXiv:1711.00937*, 2017. 2, 3
- [42] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *International Conference on Machine Learning*, pages 4055–4064. PMLR, 2018. 2
- [43] Tingting Qiao, Jing Zhang, Duanqing Xu, and Dacheng Tao. Learn, imagine and create: Text-to-image generation from prior knowledge. *Advances in Neural Information Processing Systems*, 32:887–897, 2019. 2
- [44] Tingting Qiao, Jing Zhang, Duanqing Xu, and Dacheng Tao. Mirrorgan: Learning text-to-image generation by redescription. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1505–1514, 2019. 2
- [45] 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. *arXiv preprint arXiv:2103.00020*, 2021. 6
- [46] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. 2018. 1, 2
- [47] 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. 1, 2
- [48] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. *arXiv preprint arXiv:2102.12092*, 2021. 1, 2, 3, 7
- [49] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *Advances in neural information processing systems*, pages 14866–14876, 2019. 2
- [50] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *International Conference on Machine Learning*, pages 1060–1069. PMLR, 2016. 2
- [51] Shulan Ruan, Yong Zhang, Kun Zhang, Yanbo Fan, Fan Tang, Qi Liu, and Enhong Chen. Dae-gan: Dynamic aspect-aware gan for text-to-image synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13960–13969, 2021. 2, 7
- [52] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *arXiv preprint arXiv:2104.07636*, 2021. 3
- [53] Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P Kingma. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. *arXiv preprint arXiv:1701.05517*, 2017. 2
- [54] Florian Schmidt. Generalization in generation: A closer look at exposure bias. *arXiv preprint arXiv:1910.00292*, 2019. 3

- [55] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. [6](#)
- [56] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*, 2015. [3](#)
- [57] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypemymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018. [6](#)
- [58] Shikhar Sharma, Dendi Suhubdy, Vincent Michalski, Samira Ebrahimi Kahou, and Yoshua Bengio. Chatpainter: Improving text to image generation using dialogue. *arXiv preprint arXiv:1802.08216*, 2018. [2](#)
- [59] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. [1](#), [3](#), [5](#)
- [60] Douglas M Souza, Jônatas Wehrmann, and Duncan D Ruiz. Efficient neural architecture for text-to-image synthesis. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2020. [2](#), [7](#)
- [61] Hongchen Tan, Xiuping Liu, Xin Li, Yi Zhang, and Baocai Yin. Semantics-enhanced adversarial nets for text-to-image synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10501–10510, 2019. [2](#), [7](#)
- [62] Hongchen Tan, Xiuping Liu, Meng Liu, Baocai Yin, and Xin Li. Kt-gan: knowledge-transfer generative adversarial network for text-to-image synthesis. *IEEE Transactions on Image Processing*, 30:1275–1290, 2020. [2](#)
- [63] Ming Tao, Hao Tang, Songsong Wu, Nicu Sebe, Xiao-Yuan Jing, Fei Wu, and Bingkun Bao. Df-gan: Deep fusion generative adversarial networks for text-to-image synthesis. *arXiv preprint arXiv:2008.05865*, 2020. [2](#), [7](#)
- [64] Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International Conference on Machine Learning*, pages 1747–1756. PMLR, 2016. [2](#)
- [65] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [1](#)
- [66] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. [2](#), [6](#)
- [67] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1316–1324, 2018. [2](#), [7](#)
- [68] Guojun Yin, Bin Liu, Lu Sheng, Nenghai Yu, Xiaogang Wang, and Jing Shao. Semantics disentangling for text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2327–2336, 2019. [2](#)
- [69] Han Zhang, Jing Yu Koh, Jason Baldridge, Honglak Lee, and Yinfei Yang. Cross-modal contrastive learning for text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 833–842, 2021. [2](#)
- [70] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 5907–5915, 2017. [2](#), [7](#)
- [71] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan++: Realistic image synthesis with stacked generative adversarial networks. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1947–1962, 2018. [2](#), [7](#)
- [72] Zizhao Zhang, Yuanpu Xie, and Lin Yang. Photographic text-to-image synthesis with a hierarchically-nested adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6199–6208, 2018. [2](#)
- [73] Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. Dm-gan: Dynamic memory generative adversarial networks for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5802–5810, 2019. [2](#), [7](#)