

# Global Context with Discrete Diffusion in Vector Quantised Modelling for Image Generation

Minghui Hu<sup>1</sup>Yujie Wang<sup>2</sup>Tat-Jen Cham<sup>1</sup>Jianfei Yang<sup>1</sup>P.N.Suganthan<sup>1</sup><sup>1</sup>Nanyang Technological University<sup>2</sup>Sensetime Research

{e200008,yang0478}@e.ntu.edu.sg    {astjcham,epnsugan}@ntu.edu.sg

wangyujie@sensetime.com

## Abstract

The integration of Vector Quantised Variational AutoEncoder (VQ-VAE) with autoregressive models as generation part has yielded high-quality results on image generation. However, the autoregressive models will strictly follow the progressive scanning order during the sampling phase. This leads the existing VQ series models to hardly escape the trap of lacking global information. Denoising Diffusion Probabilistic Models (DDPM) in the continuous domain have shown a capability to capture the global context, while generating high-quality images. In the discrete state space, some works have demonstrated the potential to perform text generation and low resolution image generation. We show that with the help of a content-rich discrete visual codebook from VQ-VAE, the discrete diffusion model can also generate high fidelity images with global context, which compensates for the deficiency of the classical autoregressive model along pixel space. Meanwhile, the integration of the discrete VAE with the diffusion model resolves the drawback of conventional autoregressive models being oversized, and the diffusion model which demands excessive time in the sampling process when generating images. It is found that the quality of the generated images is heavily dependent on the discrete visual codebook. Extensive experiments demonstrate that the proposed Vector Quantised Discrete Diffusion Model (VQ-DDM) is able to achieve comparable performance to top-tier methods with low complexity. It also demonstrates outstanding advantages over other vectors quantised with autoregressive models in terms of image inpainting tasks without additional training.

## 1. Introduction

Vector Quantised Variational AutoEncoder (VQ-VAE) [34] is a popular method developed to compress images into discrete representations for the generation. Typically, after the compression and discretization representation by the convolutional network, an autoregressive model is used




Figure 1. FID v.s. Operations and Parameters. The size of the blobs is proportional to the number of network parameters, the X-axis indicates FLOPs on a log scale and the Y-axis is the FID score.

to model and sample in the discrete latent space, including PixelCNN family [5, 22, 35], transformers family [4, 24], etc. However, in addition to the disadvantage of the huge number of model parameters, these autoregressive models can only make predictions based on the observed pixels (left upper part of the target pixel) due to the inductive bias caused by the strict adherence to the progressive scan order [3, 15]. If the conditional information is located at the end of the autoregressive sequence, it is difficult for the model to obtain relevant information.

A recent alternative generative model is the Denoising Diffusion Model, which can effectively mitigate the lack of global information [10, 29], also achieving comparable or state-of-the-art performance in text [1, 12], image [6, 28, 33] and speech generation [19] tasks. Diffusion models are parameterised Markov chains trained to translate simple distributions to more sophisticated target data distributions in a finite set of steps. Typically the Markov chain begins with an isotropic Gaussian distribution in continuous state space, with the transitions of the chain for reversing a diffusion process that gradually adds Gaussian noise to source images. In the inverse process, as the current step is based on the global information of the previous step in the chain, this

endows the diffusion model with the ability to capture the global information.

However, the diffusion model has a non-negligible disadvantage in that the time and computational effort involved in generating the images are enormous. The main reason is that the reverse process typically contains thousands of steps. Although we do not need to iterate through all the steps when training, all these steps are still required when generating a sample, which is much slower compared to GANs and even autoregressive models. Some recent works [21, 30] have attempted addressing these issues by decreasing the sampling steps, but the computation cost is still high as each step of the reverse process generates a full-resolution image.

In this work, we propose the **Vector Quantised Discrete Diffusion Model (VQ-DDM)**, a versatile framework for image generation consisting of a discrete variational autoencoder and a discrete diffusion model. VQ-DDM consists of two stages: (1) learning an abundant and efficient discrete representation of images, (2) fitting the prior distribution of such latent visual codes via discrete diffusion model.

VQ-DDM substantially reduces the computational resources and required time to generate high-resolution images by using a discrete scheme. Then the common problem of the lack of global content and overly large number of parameters of the autoregressive model is solved by fitting a latent variable prior using the discrete diffusion model. Finally, since a bias of codebook will limit generation quality, while model size is also dependent on the number of categories, we propose a re-build and fine-tune (ReFiT) strategy to construct a codebook with higher utilization, which will also reduce the number of parameters in our model.

In summary, our key contributions include the following:

- VQ-DDM fits the prior over discrete latent codes with a discrete diffusion model. The use of diffusion model allows the generative models consider the global information instead of only focusing on partially seen context to avoid sequential bias.
- We propose a ReFiT approach to improve the utilisation of latent representations in the visual codebook, which can increase the code usage of VQ-GAN from 31.85% to 97.07%, while the FID between reconstruction image and original training image is reduced from 10.18 to 5.64 on CelebA-HQ  $256 \times 256$ .
- VQ-DDM is highly efficient for the both number of parameters and generation speed. As shown in Figure 1, using only 120M parameters, it outperforms VQ-VAE-2 with around 10B parameters and is comparable with VQ-GAN with 1B parameters in image generation tasks in terms of image quality. It is also  $10 \sim 100$  times faster than other diffusion models for image generation [10, 30].

## 2. Preliminaries

### 2.1. Diffusion Models in continuous state space

Given data  $\mathbf{x}_0$  from a data distribution  $q(\mathbf{x}_0)$ , the diffusion model consists of two processes: the *diffusion process* and the *reverse process* [10, 29].

The *diffusion process* progressively destroys the data  $\mathbf{x}_0$  into  $\mathbf{x}_T$  over  $T$  steps, via a fixed Markov chain that gradually introduces Gaussian noise to the data according to a variance schedule  $\beta_{1:T} \in (0, 1]^T$  as follows:

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

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}). \quad (2)$$

With an adequate number of steps  $T$  and a suitable variance schedule  $\beta$ ,  $p(\mathbf{x}_T)$  becomes an isotropic Gaussian distribution.

The *reverse process* is defined as a Markov chain parameterised by  $\theta$ , which is used to restore the data from the noise:

$$p_\theta(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t), \quad (3)$$

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t), \Sigma_\theta(\mathbf{x}_t, t)). \quad (4)$$

The objective of training is to find the best  $\theta$  to fit the data distribution  $q(\mathbf{x}_0)$  by optimizing the variational lower bound (VLB) [18]

$$\begin{aligned} & \mathbb{E}_{q(\mathbf{x}_0)}[\log p_\theta(\mathbf{x}_0)] \\ &= \mathbb{E}_{q(\mathbf{x}_0)} \log \mathbb{E}_{q(\mathbf{x}_{1:T}|\mathbf{x}_0)} \left[ \frac{p_\theta(\mathbf{x}_{0:T})}{q(\mathbf{x}_{1:T}|\mathbf{x}_0)} \right] \\ &\geq \mathbb{E}_{q(\mathbf{x}_{0:T})} \left[ \log \frac{p_\theta(\mathbf{x}_{0:T})}{q(\mathbf{x}_{1:T}|\mathbf{x}_0)} \right] =: L_{\text{vlb}}. \end{aligned} \quad (5)$$

Ho *et al.* [10] revealed that the variational lower bound in Eq. 5 can be calculated with closed form expressions instead of Monte Carlo estimates as the *diffusion process* posteriors and marginals are Gaussian, which allows sampling  $\mathbf{x}_t$  at an arbitrary step  $t$  with  $\alpha_t = 1 - \beta_t$ ,  $\bar{\alpha}_t = \prod_{s=0}^t \alpha_s$  and  $\tilde{\beta}_t = \frac{1 - \bar{\alpha}_t - 1}{1 - \bar{\alpha}_t}$ :

$$q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t|\sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1 - \bar{\alpha}_t)\mathbf{I}), \quad (6)$$

$$\begin{aligned} L_{\text{vlb}} &= \mathbb{E}_{q(\mathbf{x}_0)}[D_{\text{KL}}(q(\mathbf{x}_T|\mathbf{x}_0)||p(\mathbf{x}_T)) - \log p_\theta(\mathbf{x}_0|\mathbf{x}_1)] \\ &\quad + \sum_{t=2}^T D_{\text{KL}}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)||p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)). \end{aligned} \quad (7)$$

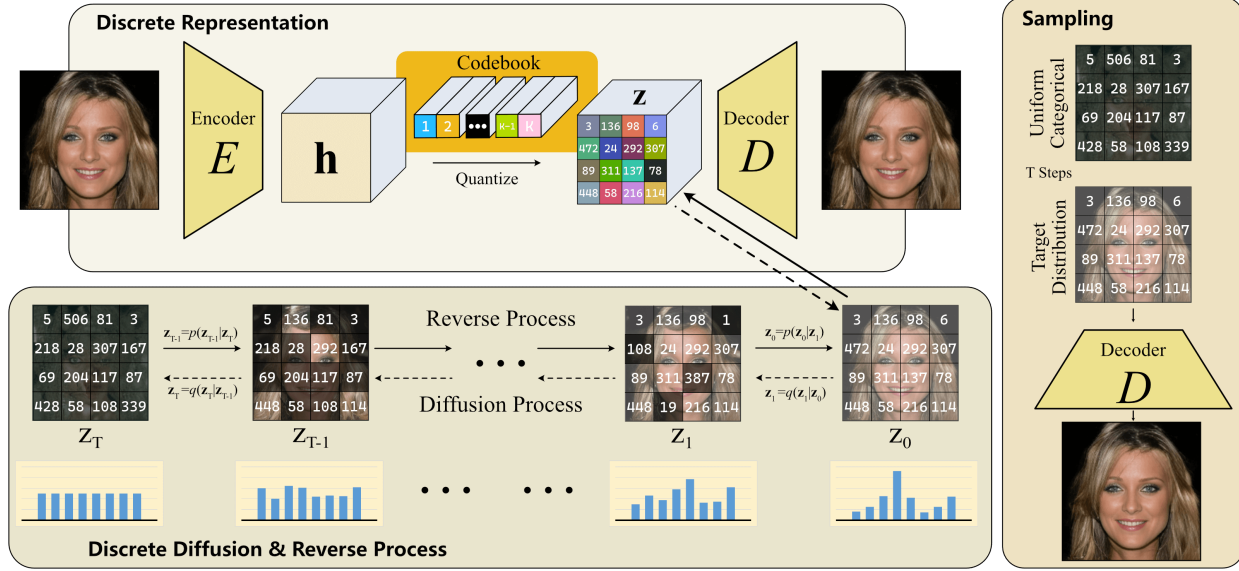


Figure 2. The proposed VQ-DDM pipeline contains 2 stages: (1) Compress the image into discrete variables via discrete VAE. (2) Fit a prior distribution over discrete coding by a diffusion model. Black squares in the diffusion diagram illustrate states when the underlying distributions are uninformative, but which become progressively more specific during the reverse process. The bar chart at the bottom of the image represents the probability of a particular discrete variable being sampled.

Thus the reverse process can be parameterised by neural networks  $\epsilon_\theta$  and  $v_\theta$ , which can be defined as:

$$\mu_\theta(\mathbf{x}_t, t) = \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}} \epsilon_\theta(\mathbf{x}_t, t) \right), \quad (8)$$

$$\Sigma_\theta(\mathbf{x}_t, t) = \exp(v_\theta(\mathbf{x}_t, t) \log \beta_t + (1 - v_\theta(\mathbf{x}_t, t)) \log \tilde{\beta}_t). \quad (9)$$

Using a modified variant of the VLB loss as a simple loss function will offer better results in the case of fixed  $\Sigma_\theta$  [10]:

$$L_{\text{simple}} = \mathbb{E}_{t, \mathbf{x}_0, \epsilon} [||\epsilon - \epsilon_\theta(\mathbf{x}_t, t)||^2], \quad (10)$$

which is a reweighted version resembling denoising score matching over multiple noise scales indexed by  $t$  [31].

Nichol *et al.* [21] used an additional  $L_{\text{vlb}}$  to the simple loss for guiding a learned  $\Sigma_\theta(\mathbf{x}_t, t)$ , while keeping the  $\mu_\theta(\mathbf{x}_t, t)$  still the dominant component of the total loss:

$$L_{\text{hybrid}} = L_{\text{simple}} + \lambda L_{\text{vlb}}. \quad (11)$$

## 2.2. Discrete Representation of Images

van den Oord *et al.* [34] presented a discrete variational autoencoder with a categorical distribution as the latent prior, which is able to map the images into a sequence of discrete latent variables by an encoder and reconstruct the image according to those variables with a decoder.

Formally, given a codebook  $\mathbb{Z} \in \mathbb{R}^{K \times d}$ , where  $K$  represents the capacity of latent variables in the codebook and

$d$  is the dimension of each latent variable, after compressing the high dimension input data  $\mathbf{x} \in \mathbb{R}^{c \times H \times W}$  into latent vectors  $\mathbf{h} \in \mathbb{R}^{h \times w \times d}$  by an encoder  $E$ ,  $\mathbf{z}$  is the quantised  $\mathbf{h}$ , which substitutes the vectors  $h_{i,j} \in \mathbf{h}$  by the nearest neighbor  $z_k \in \mathbb{Z}$ . The decoder  $D$  is trained to reconstruct the data from the quantised encoding  $\mathbf{z}_q$ :

$$\mathbf{z} = \text{Quantise}(\mathbf{h}) := \arg \min_k ||h_{i,j} - z_k||, \quad (12)$$

$$\hat{\mathbf{x}} = D(\mathbf{z}) = D(\text{Quantise}(E(\mathbf{x}))). \quad (13)$$

As  $\text{Quantise}(\cdot)$  has a non-differentiable operation  $\arg \min$ , the straight-through gradient estimator is used for back-propagating the reconstruction error from decoder to encoder. The whole model can be trained in an end-to-end manner by minimizing the following function:

$$L = ||\mathbf{x} - \hat{\mathbf{x}}||^2 + ||sg[E(\mathbf{x})] - \mathbf{z}|| + \beta ||sg[\mathbf{z}] - E(\mathbf{x})||, \quad (14)$$

where  $sg[\cdot]$  denotes stop gradient and broadly the three terms are reconstruction loss, codebook loss and commitment loss, respectively.

VQ-GAN [8] extends VQ-VAE [34] in multiple ways. It substitutes the L1 or L2 loss of the original VQ-VAE with a perceptual loss [40], and adds an additional discriminator to distinguish between real and generated patches [41].

The codebook update of the discrete variational autoencoder is intrinsically a dictionary learning process. Its objective uses L2 loss to narrow the gap between the codes  $\mathbb{Z}_t \in \mathbb{R}^{K_t \times d}$  and the encoder output  $\mathbf{h} \in \mathbb{R}^{h \times w \times d}$  [34],

where  $K_t$  is constant during all diffusion steps. In other words, the codebook training is like  $k$ -means clustering, where cluster centers are the discrete latent codes. However, since the volume of the codebook space is dimensionless and  $\mathbf{h}$  is updated each iteration, the discrete codes  $\mathbb{Z}$  typically do not follow the encoder training quickly enough. Only a few codes get updated during training, with most unused after initialization.

### 3. Methods

Our goal is to leverage the powerful generative capability of the diffusion model to perform high fidelity image generation tasks with a low number of parameters. Our proposed method, VQ-DDM, is capable of generating high fidelity images with a relatively small number of parameters and FLOPs, as summarised in Figure 2. Our solution starts by compressing the image into discrete variables via the discrete VAE and then constructs a powerful model to fit the joint distribution over the discrete codes by a diffusion model. During diffusion training, the darker coloured parts in Figure 2 represent noise introduced by uniform resampling. When the last moment is reached, the latent codes have been completely corrupted into noise. In the sampling phase, the latent codes are drawn from an uniform categorical distribution at first, and then resampled by performing reverse process  $T$  steps to get the target latent codes. Eventually, target latent codes are pushed into the decoder to generate the image.

#### 3.1. Discrete Diffusion Model

Assume the discretization is done with  $K$  categories, i.e.  $z_t \in \{1, \dots, K\}$ , with the one-hot vector representation given by  $\mathbf{z}_t \in \{0, 1\}^K$ . The corresponding probability distribution is expressed by  $\mathbf{z}_t^{\text{logits}}$  in logits. We formulate the discrete diffusion process as

$$q(\mathbf{z}_t | \mathbf{z}_{t-1}) = \text{Cat}(\mathbf{z}_t; \mathbf{z}_{t-1}^{\text{logits}} \mathbf{Q}_t), \quad (15)$$

where  $\text{Cat}(\mathbf{x} | \mathbf{p})$  is the categorical distribution parameterised by  $\mathbf{p}$ , while  $\mathbf{Q}_t$  is the process transition matrix. In our method,  $\mathbf{Q}_t = (1 - \beta_t) \mathbf{I} + \beta_t / K$ , which means  $\mathbf{z}_t$  has  $1 - \beta_t$  probability to keep the state from last timestep and  $\beta_t$  chance to resample from a uniform categorical distribution. Formally, it can be written as

$$q(\mathbf{z}_t | \mathbf{z}_{t-1}) = \text{Cat}(\mathbf{z}_t; (1 - \beta_t) \mathbf{z}_{t-1}^{\text{logits}} + \beta_t / K). \quad (16)$$

It is straightforward to get  $\mathbf{z}_t$  from  $\mathbf{z}_0$  under the schedule  $\beta_t$  with  $\alpha_t = 1 - \beta_t$ ,  $\bar{\alpha}_t = \prod_{s=0}^t \alpha_s$ :

$$q(\mathbf{z}_t | \mathbf{z}_0) = \text{Cat}(\mathbf{z}_t; \bar{\alpha}_t \mathbf{z}_0 + (1 - \bar{\alpha}_t) / K) \quad (17)$$

$$\text{or } q(\mathbf{z}_t | \mathbf{z}_0) = \text{Cat}(\mathbf{z}_t; \mathbf{z}_0 \bar{\mathbf{Q}}_t); \quad \bar{\mathbf{Q}}_t = \prod_{s=0}^t \mathbf{Q}_s. \quad (18)$$

We use the same cosine noise schedule as [12, 21] because our discrete model is also established on the latent codes with a small  $16 \times 16$  resolution. Mathematically, it can be expressed in the case of  $\bar{\alpha}$  by

$$\bar{\alpha} = \frac{f(t)}{f(0)}, \quad f(t) = \cos \left( \frac{t/T + s}{1 + s} \times \frac{\pi}{2} \right)^2. \quad (19)$$

By applying Bayes' rule, we can compute the posterior  $q(\mathbf{z}_{t-1} | \mathbf{z}_t, \mathbf{z}_0)$  as:

$$\begin{aligned} q(\mathbf{z}_{t-1} | \mathbf{z}_t, \mathbf{z}_0) &= \text{Cat} \left( \mathbf{z}_t; \frac{\mathbf{z}_t^{\text{logits}} \mathbf{Q}_t^\top \odot \mathbf{z}_0 \bar{\mathbf{Q}}_{t-1}}{\mathbf{z}_0 \bar{\mathbf{Q}}_t \mathbf{z}_t^{\text{logits}^\top}} \right) \\ &= \text{Cat}(\mathbf{z}_t; \boldsymbol{\theta}(\mathbf{z}_t, \mathbf{z}_0) / \sum_{k=1}^K \theta_k(z_{t,k}, z_{0,k})), \end{aligned} \quad (20)$$

$$\begin{aligned} \boldsymbol{\theta}(\mathbf{z}_t, \mathbf{z}_0) &= [\alpha_t \mathbf{z}_t^{\text{logits}} + (1 - \alpha_t) / K] \\ &\odot [\bar{\alpha}_{t-1} \mathbf{z}_0 + (1 - \bar{\alpha}_{t-1}) / K]. \end{aligned} \quad (21)$$

It is worth noting that  $\boldsymbol{\theta}(\mathbf{z}_t, \mathbf{z}_0) / \sum_{k=1}^K \theta_k(z_{t,k}, z_{0,k})$  is the normalised version of  $\boldsymbol{\theta}(\mathbf{z}_t, \mathbf{z}_0)$ , and we use  $\mathbf{N}[\boldsymbol{\theta}(\mathbf{z}_t, \mathbf{z}_0)]$  to denote  $\boldsymbol{\theta}(\mathbf{z}_t, \mathbf{z}_0) / \sum_{k=1}^K \theta_k(z_{t,k}, z_{0,k})$  below.

Hoogeboom *et al.* [12] predicted  $\hat{\mathbf{z}}_0$  from  $\mathbf{z}_t$  with a neural network  $\mu(\mathbf{z}_t, t)$ , instead of directly predicting  $p_\theta(\mathbf{z}_{t-1} | \mathbf{z}_t)$ . Thus the reverse process can be parameterised by the probability vector from  $q(\mathbf{z}_{t-1} | \mathbf{z}_t, \hat{\mathbf{z}}_0)$ . Generally, the reverse process  $p_\theta(\mathbf{z}_{t-1} | \mathbf{z}_t)$  can be expressed by

$$\begin{aligned} p_\theta(\mathbf{z}_0 | \mathbf{z}_1) &= \text{Cat}(\mathbf{z}_0 | \hat{\mathbf{z}}_0), \\ p_\theta(\mathbf{z}_{t-1} | \mathbf{z}_t) &= \text{Cat}(\mathbf{z}_t | \mathbf{N}[\boldsymbol{\theta}(\mathbf{z}_t, \hat{\mathbf{z}}_0)]). \end{aligned} \quad (22)$$

Inspired by [13, 20], we use a neural network  $\mu(\mathbf{Z}_t, t)$  to learn and predict the a noise  $n_t$  and obtain the logits of  $\hat{\mathbf{z}}_0$ .

It is worth noting that the neural network  $\mu(\cdot)$  is based on the  $\mathbf{Z}_t \in \mathbb{N}^{h \times w}$ , where all the discrete representation  $\mathbf{z}_t$  of the image are combined. The final noise prior  $\mathbf{Z}_T$  is uninformative, and it is possible to separably sample from each axis during inference. However, the reverse process is jointly informed and evolves towards a highly coupled  $\mathbf{Z}_0$ . We do not define a specific joint prior for  $\mathbf{z}_t$ , but encode the joint relationship into the learned reverse process. This is implicitly done in the continuous domain diffusion. As  $\mathbf{z}_{t-1}$  is based on the whole previous representation  $\mathbf{z}_t$ , the reverse process can sample the whole discrete code map directly while capturing the global information.

The loss function used is the VLB from Eq. 7, where the summed KL divergence for  $T > 2$  is given by

$$\begin{aligned} \text{KL}(q(\mathbf{z}_{t-1} | \mathbf{z}_t, \mathbf{z}_0) || p_\theta(\mathbf{z}_{t-1} | \mathbf{z}_t)) &= \\ \sum_k \mathbf{N}[\boldsymbol{\theta}(\mathbf{z}_t, \mathbf{z}_0)] \times \log \frac{\mathbf{N}[\boldsymbol{\theta}(\mathbf{z}_t, \mathbf{z}_0)]}{\mathbf{N}[\boldsymbol{\theta}(\mathbf{z}_t, \hat{\mathbf{z}}_0)]}. \end{aligned} \quad (23)$$



### 3.2. Re-build and Fine-tune Strategy

Our discrete diffusion model is based on the latent representation of the discrete VAE codebook  $\mathbb{Z}$ . However, the codebooks with rich content are normally large, with some even reaching  $K = 16384$ . This makes it highly unwieldy for our discrete diffusion model, as the transition matrices of discrete diffusion models have a quadratic level of growth to the number of classes  $K$ , *e.g.*  $O(K^2T)$  [1].

To reduce the categories used for our diffusion model, we proposed a Re-build and Fine-tune (ReFit) strategy to decrease the size  $K$  of codebook  $\mathbb{Z}$  and boost the reconstruction performance based on a well-trained discrete VAEs trained by the straight-through method.

From Eq. 14, we can find the second term and the third term are related to the codebook, but only the second term is involved in the update of the codebook.  $\|sg[E(\mathbf{x})] - \mathbf{z}\|$  reveals that only a few selected codes, the same number as the features from  $E(\mathbf{x})$ , are engaged in the update per iteration. Most of the codes are not updated or used after initialization, and the update of the codebook can lapse into a local optimum.

We introduce a re-build and fine-tune strategy to avoid the waste of codebook capacity. With the trained encoder, we reconstruct the codebook so that all codes in the codebook have the opportunity to be selected. This will greatly increase the usage of the codebook. Suppose we desire to obtain a discrete VAE having a codebook with  $\mathbb{Z}_t$  based on a trained discrete VAE with an encoder  $E_s$  and a decoder  $D_s$ . We first encode each image  $\mathbf{x} \in \mathbb{R}^{c \times H \times W}$  to latent features  $\mathbf{h}$ , or loosely speaking, each image gives us  $h \times w$  features with  $d$  dimension. Next we sample  $P$  features uniformly from the entire set of features found in training images, where  $P$  is the sampling number and far larger than the desired codebook capacity  $K_t$ . This ensures that the re-build codebook is composed of valid latent codes. Since the process of codebook training is basically the process of finding cluster centres, we directly employ k-means with AFK-MC<sup>2</sup> [2] on the sampled  $P$  features and utilise the centres to re-build the codebook  $\mathbb{Z}_t$ . We then replace the original codebook with the re-build  $\mathbb{Z}_t$  and fine-tune it on top of the well-trained discrete VAE.

## 4. Experiments and Analysis

### 4.1. Datasets and Implementation Details

We show the effectiveness of the proposed VQ-DDM on *CelebA-HQ* [14] and *LSUN-Church* [39] datasets and verify the proposed Re-build and Fine-tune strategy on *CelebA-HQ* and *ImageNet* datasets. The details of the dataset are given in the Appendix.

The discrete VAE follows the same training strategy as VQ-GAN [8]. All training images are processed to  $256 \times 256$ , and the compress ratio is set to 16, which means

the latent vector  $\mathbf{z} \in \mathbb{R}^{1 \times 16 \times 16}$ . When conducting Rebuild and Fine-tune, the sampling number  $P$  is set to  $20k$  for *LSUN* and *CelebA*. For the more content-rich case, we tried a larger  $P$  value  $50k$  for *ImageNet*. In practical experiments, we sample  $P$  images with replacement uniformly from the whole training data and obtained corresponding latent features. For each feature map, we make another uniform sampling over the feature map size  $16 \times 16$  to get the desired features. In the fine-tuning phase, we freeze the encoder and set the learning rate of the decoder to  $1e-6$  and the learning rate of the discriminator to  $2e-6$  with 8 instances per batch.

With regard to the diffusion model, the network for estimating  $n_t$  has the same structure as [10], which is a U-Net [26] with self-attention [36]. The detailed settings of hyperparameters are provided in the Appendix. We set timestep  $T = 4000$  in our experiments and the noise schedule is the same as [21].

### 4.2. Codebook Quality

A large codebook dramatically increases the cost of DDM. To reduce the cost to an acceptable scale, we proposed a resample and fine-tune strategy to compress the size of the codebook, while maintaining quality. To demonstrate the effectiveness of the proposed strategy, we compare the codebook usage and FID of reconstructed images of our method to VQ-GAN [8], VQ-VAE-2 [25] and DALL-E [24].

In this experiment, we compressed the images from  $3 \times 256 \times 256$  to  $1 \times 16 \times 16$  with two different codebook capacities  $K = \{512, 1024\}$ . We also proposed an indicator to measure the usage rate of the codebook, which is the number of discrete features that have appeared in the test set or training set divided by the codebook capacity. The quantitative comparison results are shown in Table 1 while the reconstruct images are demonstrated in Figs. 3 & 4. Reducing the codebook capacity from 1024 to 512 only brings  $\sim 0.1$  decline in CelebA and  $\sim 1$  in ImageNet. As seen in Figure 4, the reconstructed images (c,d) after ReFiT strategy are richer in colour and more realistic in expression than the reconstructions from VQ-GAN (b). The codebook usage of our method has improved significantly compared to other methods, nearly 3x high than the second best. Our method also achieves the equivalent reconstruction quality at the same compression rate and with  $32\times$  lower capacity  $K$  of codebook  $\mathbb{Z}$ .

For VQ-GAN with capacity 16384, although it only has 976 effective codes, which is smaller than 1024 in our ReFiT method when  $P = 20k$ , it achieves a lower FID in reconstructed images vs validation images. One possible reason is that the value of  $P$  is not large enough to cover some infrequent combinations of features during the re-build phase. As the results in Table 1, after we increase the sampling number  $P$  from  $20k$  to  $100k$ , we observe that

Model	Latent Size	Capacity	Usage of $\mathbb{Z}$		FID $\downarrow$	
			CelebA	ImageNet	CelebA	ImageNet
VQ-VAE-2	Cascade	512	$\sim 65\%$	-	-	$\sim 10$
DALL-E	32x32	8192	-	-	-	32.01
VQ-GAN	16x16	16384	-	5.96%	-	4.98
VQ-GAN	16x16	1024	31.85%	33.67%	10.18	7.94
<i>ours</i> ( $P = 100k$ )	16x16	1024	-	100%	-	4.98
<i>ours</i> ( $P = 20k$ )	16x16	1024	97.07%	100%	5.59	5.99
<i>ours</i> ( $P = 20k$ )	16x16	512	93.06%	100%	5.64	6.95

<sup>1</sup> All methods are trained straight-through, except DALL-E with Gumbel-Softmax [24].

<sup>2</sup> CelebA-HQ at  $256 \times 256$ . Reported FID is between 30k reconstructed data vs training data.

<sup>3</sup> Reported FID is between 50k reconstructed data vs validation data

Table 1. FID between reconstructed images and original images on CelebA-HQ and ImageNet

increasing the value of  $P$  achieved higher performance.

### 4.3. Generation Quality

We evaluate the performance of VQ-DDM for the unconditional image generation on *CelebA-HQ*  $256 \times 256$ . Specifically, we evaluated the performance of our approach in terms of FID and compared it with various likelihood-based methods including GLOW [16], NVAE [32], VAEBM [38], DC-VAE [23], VQ-GAN [8] and likelihood-free method, e.g., PGGAN [14]. We also conducted an experiment on *LSUN-Church*.

In *CelebA-HQ* experiments, the discrete diffusion model was trained with  $K = 512$  and  $K = 1024$  codebooks respectively. We also report the different FID from  $T = 2$  to  $T = 4000$  with corresponding time consumption in Figure 6. Regarding the generation speed, it took about 1000 hours to generate  $50k$   $256 \times 256$  images using DDPM with 1000 steps on a NVIDIA 2080Ti GPU, 100 hours for DDIM with 100 steps [30], and around 10 hours for our VQ-DDM with 1000 steps.

Table 2 shows the main results on VQ-DDM along with other established models. Although VQ-DDM is also a likelihood-based method, the training phase relies on the negative log-likelihood (NLL) of discrete hidden variables, so we do not compare the NLL between our method and the other methods. The training NLL is around 1.258 and test NLL is 1.286 while the FID is 13.2. Fig. 7a shows the generated samples from VQ-DDM trained on the *CelebA-HQ*.

For *LSUN-Church*, the codebook capacity  $K$  is set to 1024, while the other parameters are set exactly the same. The training NLL is 1.803 and the test NLL is 1.756 while the FID between the generated images and the training set is 16.9. Some samples are shown in Fig. 7b.

After utilizing ReFiT, the generation quality of the model is significantly improved, which implies a decent codebook can have a significant impact on the subsequent generative phase. Within a certain range, the larger the codebook capacity leads to a better performance. However, excessive number of codebook entries will cause the model collapse [12].

## 4.4. Image Inpainting

Autoregressive models have recently demonstrated superior performance in the image inpainting tasks [4, 8]. However, one limitation of this approach is that if the important context is found at the end of the autoregressive series, the models will not be able to correctly complete the images. As mentioned in Sec. 3.1, the diffusion model will directly sample the full latent code map, with sampling steps based on the *full* discrete map of the previous step. Hence it can significantly improve inpainting as it does not depend on context sequencing.

We perform the mask diffusion and reverse process in the discrete latent space. After encoding the masked image  $x_0 \sim q(\mathbf{x}_0)$  to discrete representations  $z_0 \sim q(\mathbf{z}_0)$ , we diffuse  $\mathbf{z}_0$  with  $t$  steps to  $\tilde{\mathbf{z}}_t \sim q(\mathbf{z}_t|\mathbf{z}_0)$ . Thus the last step with mask  $\tilde{\mathbf{z}}_T^m$  can be demonstrated as  $\tilde{\mathbf{z}}_T^m = (1 - m) \times \tilde{\mathbf{z}}_T + m \times \mathbb{C}$ , where  $\mathbb{C} \sim \text{Cat}(K, 1/K)$  is the sample from a uniform categorical distribution and  $m \in \{0, 1\}^K$  is the mask,  $m = 0$  means the context there is masked and  $m = 1$  means that given the information there. In the reverse process,  $\mathbf{z}_{T-1}$  can be sampled from  $p_\theta(\mathbf{z}_{T-1}|\tilde{\mathbf{z}}_T^m)$  at  $t = T$ , otherwise,  $\mathbf{z}_{t-1} \sim p_\theta(\mathbf{z}_{t-1}|\mathbf{z}_t^m)$ , and the masked  $\mathbf{z}_{t-1}^m = (1 - m) \times \mathbf{z}_{t-1} + m \times \tilde{\mathbf{z}}_{t-1}$ .

We compare our approach and another that exploits a transformer with a sliding attention window as an autoregressive generative model [8]. The completions are shown in Fig. 8, in the first row, the upper 62.5% (160 out of 256 in latent space) of the input image is masked and the lower 37.5% (96 out of 256) is retained, and in the second row, only a quarter of the image information in the lower right corner is retained as input. We also tried masking in an arbitrary position. In the third row, we masked the perimeter, leaving only a quarter part in the middle. Since the reverse diffusion process captures the global relationships, the image completions of our model performs much better. Our method can make a consistent completions based on arbitrary contexts, whereas the inpainting parts from transformer lack consistency. It is also worth noting that our model requires no additional training in solving the task of image inpainting.

## 5. Related Work

### 5.1. Vector Quantised Variational Autoencoders

VQ-VAE [34] leads a trend of discrete representation of images. The common practice is to model the discrete representations using an autoregressive model, e.g. PixelCNN [5, 35], transformers [8, 24, 24], etc. Some works had attempted to fit the prior distribution of discrete latent variables using a light non-autoregressive approach, like EM approach [27] and Markov chain with self-organizing map [9], but yet they are struggling to fit a large scale of data. Ho *et al.* [10] have also shown that the diffusion mod-

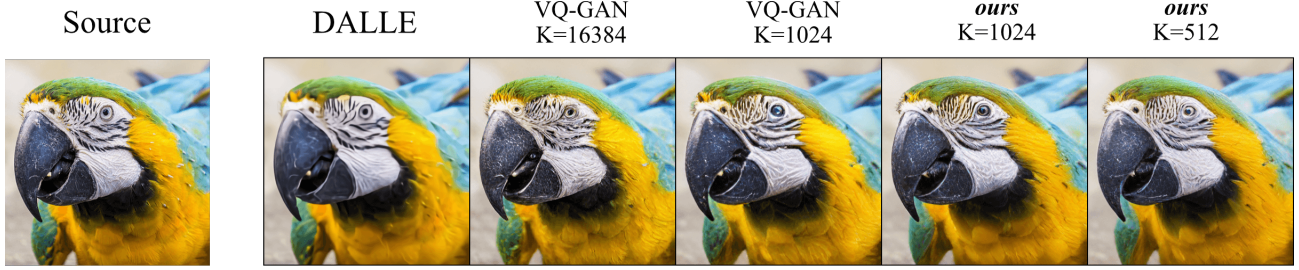


Figure 3. Reconstruction images  $384 \times 384$  from ImageNet based VQ-GAN and ReFiT




Figure 4. Reconstruction images of CelebA HQ  $256 \times 256$  from VQ-GAN and ReFiT.




Figure 5. Steps and corresponding FID during the sampling. The text annotations are hours to sample 50k latent feature maps on 1 NVIDIA 2080Ti GPU




Figure 6. Hours to sampling 50k latent codes by VQ-DDM and generating 50k images with VQ-DDM and DDPM

els can be regarded as an autoregressive model along the time dimension, but in reality, it is non-autoregressive along the pixel dimension. A concurrent work [7] follow a similar pipeline which uses a diffusion model on discrete latent variables, but the work uses parallel modeling of multiple short Markov chains to achieve denoising.

Method	FID ↓	Params	FLOPs
<b>Likelihood-based</b>			
GLOW [16]	60.9	220 M	540 G
NVAE [32]	40.3	1.26 G	185 G
<i>ours</i> ( $K = 1024$ w/o ReFiT)	22.6	117 M	1.06 G
VAEBM [38]	20.4	127 M	8.22 G
<i>ours</i> ( $K = 512$ w/ ReFiT)	18.8	117 M	<b>1.04 G</b>
DC-VAE [23]	15.8	-	-
<i>ours</i> ( $K = 1024$ w/ ReFiT)	13.2	117 M	1.06 G
DDIM( $T=100$ ) [30]	10.9	114 M	124 G
VQ-GAN + Transformer [8]	10.2	802 M	102 G <sup>a</sup>
<b>Likelihood-free</b>			
PG-GAN [14]	8.0	46.1 M	14.1 G

<sup>a</sup> VQ-GAN is an autoregressive model, and the number in the table is the computation needed to generate the full size latent feature map. The FLOPs needed to generate one discrete index out of 256 is 0.399 G.

Table 2. FID on CelebA HQ  $256 \times 256$  dataset. All the FLOPs in the table only consider the generation stage or inference phase for one  $256 \times 256$  images.

## 5.2. Diffusion Models

Sohl-Dickstein *et al.* [29] presented a simple discrete diffusion model, which diffused the target distribution into the independent binomial distribution. Recently, Hooeboom *et al.* [12] have extended the discrete model from binomial to multinomial. Further, Austin *et al.* [1] proposed a generalised discrete diffusion structure, which provides several choices for the diffusion transition process.

In the continuous state space, there are some recent diffusion models that surpassed the state-of-the-art in the image generation area. With the guidance from the classifiers, Dhariwal *et al.* [6] enabled diffusion models called ADM to generate images beyond BigGAN, which was previously one of the most powerful generative models. In CDM [11], the authors performed the cascade pipeline on the diffusion model to generate the image with ultra-high fidelity and reach state-of-the-art on conditional ImageNet generation. In addition, there have been several recent works that have





(a) Samples ( $256 \times 256$ ) from a VQ-DDM model trained on CelebA HQ. FID=13.2



(b) Samples ( $256 \times 256$ ) from a VQ-DDM model trained on LSUN-Church. FID=16.9

Figure 7. Samples from VQ-DDM models.

Raw Image	Masked Input	VQ-DDM	VQ-GAN + Transformer

Figure 8. Completions with the arbitrary masks.

attempted to use diffusion models to modelling the latent variables of VAE [17, 37], while revealed the connection among several diffusion models mentioned above.

## 6. Conclusion

In this paper, we introduce VQ-DDM, a high-fidelity image generation model with a two-stage pipeline. In the first stage, we train a discrete VAE with a well-utilised content-rich codebook. With the help of such an efficient codebook, it is possible to generate high-quality images by a discrete diffusion model with relatively tiny parameters in the second stage. Simultaneously, benefiting from the discrete diffusion model, the sampling process captures the global information and the image inpainting is no longer affected

by the location of the given context and mask. Meanwhile, in comparison with other diffusion models, our approach further reduces the gap in generation speed with respect to GAN. We believe that VQ-DDM can also be utilised for audio, video and multimodal generation.

## Limitations

For a complete diffusion, we need a large number of steps, which will result in a very fluctuating training process and limit the image generation quality. Hence, our model may suffer from underperformance when exposed to the large scale and complex datasets.



## References

- [1] Jacob Austin, Daniel Johnson, Jonathan Ho, Danny Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. *arXiv preprint arXiv:2107.03006*, 2021. 1, 5, 7
- [2] Olivier Bachem, Mario Lucic, Hamed Hassani, and Andreas Krause. Fast and provably good seedings for k-means. *Advances in neural information processing systems*, 29:55–63, 2016. 5
- [3] Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. Scheduled sampling for sequence prediction with recurrent neural networks. *arXiv preprint arXiv:1506.03099*, 2015. 1
- [4] 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. 1, 6
- [5] Xi Chen, Nikhil Mishra, Mostafa Rohaninejad, and Pieter Abbeel. Pixelsnail: An improved autoregressive generative model. In *International Conference on Machine Learning*, pages 864–872. PMLR, 2018. 1, 6
- [6] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis. *arXiv e-prints*, pages arXiv–2105, 2021. 1, 7
- [7] 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. 7
- [8] Patrick Esser, Robin Rombach, and Björn 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. 3, 5, 6, 7
- [9] Vincent Fortuin, Matthias Hüser, Francesco Locatello, Heiko Strathmann, and Gunnar Rätsch. Som-vae: Interpretable discrete representation learning on time series. *arXiv preprint arXiv:1806.02199*, 2018. 6
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *arXiv preprint arXiv:2006.11239*, 2020. 1, 2, 3, 5, 6
- [11] 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. 7
- [12] 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. 1, 4, 6, 7
- [13] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016. 4
- [14] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 5, 6, 7
- [15] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *arXiv preprint arXiv:2101.01169*, 2021. 1
- [16] Diederik P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *arXiv preprint arXiv:1807.03039*, 2018. 6, 7
- [17] Diederik P Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *arXiv preprint arXiv:2107.00630*, 2021. 8
- [18] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2
- [19] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations*, 2020. 1
- [20] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016. 4
- [21] Alex Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. *arXiv preprint arXiv:2102.09672*, 2021. 2, 3, 4, 5
- [22] Aaron van den Oord, Nal Kalchbrenner, Oriol Vinyals, Lasse Espeholt, Alex Graves, and Koray Kavukcuoglu. Conditional image generation with pixelcnn decoders. *arXiv preprint arXiv:1606.05328*, 2016. 1
- [23] Gaurav Parmar, Dacheng Li, Kwong Lee, and Zhuowen Tu. Dual contradistinctive generative autoencoder. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 823–832, 2021. 6, 7
- [24] 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, 5, 6
- [25] 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. 5
- [26] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 5
- [27] Aurko Roy, Ashish Vaswani, Arvind Neelakantan, and Niki Parmar. Theory and experiments on vector quantized autoencoders. *arXiv preprint arXiv:1805.11063*, 2018. 6
- [28] Abhishek Sinha, Jiaming Song, Chenlin Meng, and Stefano Ermon. D2c: Diffusion-denoising models for few-shot conditional generation. *arXiv preprint arXiv:2106.06819*, 2021. 1
- [29] 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, 2, 7
- [30] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2020. 2, 6, 7

- [31] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Proceedings of the 33rd Annual Conference on Neural Information Processing Systems*, 2019. 3
- [32] Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical variational autoencoder. *arXiv preprint arXiv:2007.03898*, 2020. 6, 7
- [33] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. *Advances in Neural Information Processing Systems*, 34, 2021. 1
- [34] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6309–6318, 2017. 1, 3, 6
- [35] Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International Conference on Machine Learning*, pages 1747–1756. PMLR, 2016. 1, 6
- [36] 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. 5
- [37] Antoine Wehenkel and Gilles Louppe. Diffusion priors in variational autoencoders. In *ICML Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models*, 2021. 8
- [38] Zhisheng Xiao, Karsten Kreis, Jan Kautz, and Arash Vahdat. Vaebm: A symbiosis between variational autoencoders and energy-based models. In *International Conference on Learning Representations*, 2020. 6, 7
- [39] 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. 5
- [40] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 3
- [41] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networkss. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2017. 3