

Addressing Representation Collapse in Vector Quantized Models with One Linear Layer

Yongxin Zhu^{1,2}, Bocheng Li^{1,2}, Yifei Xin³, Zhihua Xia⁴, Linli Xu^{1,2*}

¹University of Science and Technology of China

²State Key Laboratory of Cognitive Intelligence, ³Peking University, ⁴Jinan University

zyx2016@mail.ustc.edu.cn, bcli@mail.ustc.edu.cn

xinyifei@stu.pku.edu.cn, xiazhihua@jnu.edu.cn, linlixu@ustc.edu.cn

Abstract

*Vector Quantization (VQ) is essential for discretizing continuous representations in unsupervised learning but suffers from representation collapse, causing low codebook utilization and limiting scalability. Existing solutions often rely on complex optimizations or reduce latent dimensionality, which compromises model capacity and fails to fully solve the problem. We identify the root cause as disjoint codebook optimization, where only a few code vectors are updated via gradient descent. To fix this, we propose **SimpleVQ**, which reparameterizes code vectors through a learnable linear transformation layer over a latent basis, optimizing the entire linear space rather than nearest individual code vectors. Although the multiplication of two linear matrices is equivalent to applying a single linear layer, this simple approach effectively prevents collapse. Extensive experiments on image and audio tasks demonstrate that SimVQ improves codebook usage, is easy to implement, and generalizes well across modalities and architectures.*

1. Introduction

In recent years, vector quantization (VQ) [19, 29] has emerged as a foundational technique in unsupervised representation learning [2, 5] and latent generative models [4, 21, 31, 33, 34, 44]. By converting continuous representations into discrete codes, VQ models can effectively identify the inherent structure of data and enable various discrete modeling methods on continuous data, from high-quality image generation [9] to audio synthesis [6]. The recent success of Large Language Models (LLMs) [1] has highlighted the effectiveness of next-token prediction as a powerful and versatile training objective. Consequently, VQ models are taken as the direct method to transform data from various modalities [25, 28, 37] or scientific domains [10] to discrete

sequences for next token prediction training. However, attempts to integrate VQ models as multimodal tokenizers to leverage the scaling laws of LLMs face significant challenges because of the difficulty of expanding the codebook. For example, the Chameleon model [28] constrains its codebook size to $8k$, which is significantly trailing behind the vocabulary size of LLMs (e.g., LLaMA3’s vocabulary size is $128k$ [8]).

There is a broad agreement that increasing vocabulary size can consistently improve the performance of LLMs [27]. However, recent studies [42] indicate that traditional VQ models often fail to utilize the additional parameters introduced by codebook expansion, leaving most codes inactive during training. The contradiction between codebook expansion and low codebook utilization in VQ models is known as the representation collapse problem [22], where increasing the codebook size fails to improve the performance. To address these discrepancies, we conduct a theoretical analysis of the optimization procedure of VQ models and identify that the disjoint optimization of the codebook is the root cause of representation collapse. As illustrated in Fig. 1(a), the core mechanism of VQ models involves a nearest-neighbor replacement strategy, where the encoder’s output features are replaced by the nearest vector in the codebook to serve as input to the decoder. The indices of the nearest vector are taken as the discrete representation of the data. This nearest-selection operator results in only a subset of codes being updated through gradient descent, while the remaining codes remain unchanged.

Some methods mitigate representation collapse by hand-designing complex optimization strategies, such as stochastic quantization [26], distribution penalty [30, 32], and codebook reset [40] through sophisticated training strategies. Recently, some approaches [16, 33, 35] propose to reduce the dimension of the latent space to a very small scale (e.g., 8 v.s. 128) to alleviate the curse of dimensionality, thereby improving the overlap between the encoder’s fea-

*Corresponding author.




Figure 1. **Comparison of Vanilla VQ and SimVQ.** (a): (left) Disjoint optimization in Vanilla VQ. Only the nearest codes are updated, resulting in a high percentage of “dead” codes that are not updated. (b): (right) Joint optimization in SimVQ. The entire codebook is updated with a latent basis, ensuring all codes remain active.

tures and the codebook. However, while these methods enhance codebook utilization, they do so at the cost of model capacity, leading to worse performance compared to vanilla VQ models when the codebook size is small or representation collapse is not severe. Another approach, VQGAN-LC [42], initializes the codebook with features extracted from the pre-trained CLIP model [17] to create a well-structured latent space that better matches the distribution of the encoder output. Nevertheless, the latent space defined by an external pre-trained model limits the model’s ability to generalize to diverse datasets and reaches a performance plateau as the codebook size increases. These limitations highlight the need for a more effective method to improve codebook utilization without compromising model capacity or relying on external models.

We critically assess prevalent methodologies and reveal that optimizing the latent space rather than individual code vectors is key to preventing representation collapse. Building on this insight, we introduce a simple yet effective method, termed SimVQ, to directly update the latent space spanned by the codebook by linear transforming the code vectors via a learnable latent basis. Specifically, the vectors in the codebook are reparameterized as a linear combination of the basis in the learnable linear layer \mathbf{W} :

$$\mathbf{C} \in \mathbb{R}^{K \times d} \Rightarrow \mathbf{CW} \text{ with } \mathbf{W} \in \mathbb{R}^{d \times d}, \quad (1)$$

where K denotes the codebook size and d represents the dimension of latent space. This reparameterization with linear transformation disentangles the optimization of the codebook into two components: the coefficient matrix \mathbf{C} and the basis of linear space \mathbf{W} respectively. As illustrated in Fig. 1(b), by optimizing the basis matrix \mathbf{W} , the latent space spanned by \mathbf{CW} is rotated and stretched to match encoder’s output feature. The entire codebook is updated jointly to prevent the representation collapse problem. The simplicity of the proposed method makes it highly portable and easily adaptable for improving VQ-based models across a wide range of domains, requiring only *one linear layer*.

In summary, our contributions to vector quantized models are as follows:

- We theoretically analyze the representation collapse problem in VQ models and reveal that optimizing the latent space spanned by the codebook, rather than individual code vectors, is crucial to addressing this issue.
- We propose a novel method, SimVQ, which reparameterizes the codebook vectors in VQ models via a linear transformation with a learnable latent basis. This simple yet effective approach is highly adaptable and easy to implement, making it broadly applicable across various machine learning contexts.
- We conduct an extensive evaluation of SimVQ across diverse modalities, including image and audio with different model architectures. The results show that SimVQ not only effectively addresses the representation collapse problem by achieving near-complete codebook utilization regardless of the codebook size, but also establishes new state-of-the-art performance. Furthermore, when scaling up the codebook size, SimVQ consistently delivers improved results.

2. Related Work

VQ-VAE [29] is the pioneering work to encode data into discrete representations. Building on these developments, VQGAN [9] combines VQ-VAE with adversarial networks to improve the perceptual quality of generated samples and establish a fundamental quantization protocol in latent generative models [21, 28, 34]. Many traditional VQ works focus on training a better discrete representation rather than codebook utilization to improve reconstruction performance. For example, RVQ [15] and MoVQ [41] enhance the reconstruction details with multichannel quantization. However, these methods suffer from a critical issue of representation collapse, as they struggle to scale the codebook size beyond 10k entries, limiting their scalability. In response to this challenge, several approaches have been pro-

posed recently. DALLE [18] employs the gumbel-softmax trick [12] and stochastic sampling strategies to activate most codes during training. However, during inference, only a small subset of codes is utilized for quantization [38]. Some methods design complex optimization strategies to improve codebook utilization, such as stochastic quantization [26], distribution penalty [30, 32] and codebook reset [40]. Huh et al. [11] proposes rescaling the vectors in the codebook during training to match the distributions in the latent space. VQGAN-FC [33] introduces a method to map latent vectors into a lower-dimensional space followed by l_2 normalization to alleviate representation collapse. FSQ [16] extends this idea by projecting representations into a reduced-dimensional space, where they are quantized into a small set of fixed values. LFQ [35], a variant of FSQ, uses binary values for quantized representations, thereby simplifying the encoding process. While these methods improve the codebook utilization, they do so at the cost of model capacity by significantly reducing the dimensionality of latent space (often to as low as 8), leading to worse performance compared to vanilla VQ models when the codebook size is small and representation collapse is not severe. Additionally, VQGAN-LC [42] proposes to initialize the codebook using features extracted from the pre-trained CLIP model to avoid representation collapse. However, the reliance on the pre-trained model limits the VQ model’s ability to generalize to diverse datasets and results in a performance plateau as the codebook size increases. In contrast, our method, SimVQ, effectively addresses the representation collapse problem with a simple linear layer, without sacrificing model capacity or relying on external pre-trained models.

3. Representation Collapse in VQ Models

3.1. Preliminaries

A vector quantized model is typically a reconstructive encoder-decoder architecture that includes a vector quantization layer to convert continuous representations into discrete codes. For simplicity, we represent an image with a single random variable x . Formally, the encoder f_θ maps the input image into a latent space, producing a continuous representation $z_e = f_\theta(x) \in \mathbb{R}^d$. This representation is then quantized using a learnable codebook $\mathbf{C} = [q_1, \dots, q_K] \in \mathbb{R}^{K \times d}$, where q_i is a codebook vector. We define $\delta_k \in \{0, 1\}^{1 \times K}$ as a characteristic (one-hot) vector where only the k -th element is 1, such that $q_k = \delta_k \mathbf{C} \in \mathbb{R}^{1 \times d}$. The quantization layer selects the nearest codebook vector q_k by minimizing the Euclidean distance between z_e and the codebook entries [29]:

$$k = \arg \min_j \|z_e - q_j\|_2^2 = \arg \min_j \|z_e - \delta_j \mathbf{C}\|_2^2. \quad (2)$$

The selected vector q_k is then passed to the decoder g_ϕ to reconstruct the input image.

To enable gradient propagation through the non-differentiable characteristic vector δ_k , the straight-through estimator (STE) [3] is applied. During the backward pass, the gradient of $z_q = \delta_k \mathbf{C}$ is copied to z_e as follows,

$$z_q = \text{sg}(\delta_k \mathbf{C} - z_e) + z_e, \quad \Rightarrow \frac{\partial z_q}{\partial z_e} = 1 \quad (3)$$

where sg is the stop gradient operator, ensuring the gradient for $\delta_k \mathbf{C}$ is discarded during the backward pass.

The learning objective is the combination of a reconstruction loss and commitment loss that ensures that the encoder commits to an embedding and the encoder’s output does not drift:

$$\mathcal{L} = \log p(x|z_q) + \|\text{sg}(\delta_k \mathbf{C}) - z_e\|_2^2 + \beta \|\delta_k \mathbf{C} - \text{sg}(z_e)\|_2^2, \quad (4)$$

where $\log p(x|z_q)$ is typically the mean squared error (MSE) loss $\|x - g_\phi(z_q)\|_2^2$ for image and audio data.

3.2. Disjoint Optimization of Codebook

In VQ models, only the nearest code is selected and updated via gradient descent. Ideally, all codebook entries should be updated and utilized for decoding. However, experimental evidence shows that only a small fraction of the codebook gets updated and utilized, leading to what is known as the representation collapse problem [22]. To investigate the root cause of this issue, we provide a theoretical analysis of the optimization dynamics in VQ models.

Due to the use of the straight-through estimator (STE) for gradient propagation, the codebook \mathbf{C} can only be updated through the gradient of the commitment loss, which is defined as:

$$\mathcal{L}_{\text{commit}}(\mathbf{C}) = \|z_e - \delta_k \mathbf{C}\|_2^2. \quad (5)$$

The codebook \mathbf{C} is updated according to the following equation, where η is the learning rate:

$$\mathbf{C}^{(t+1)} = \mathbf{C}^{(t)} + \eta \mathbb{E}_{z_e} \left[\frac{\partial \mathcal{L}_{\text{commit}}(\mathbf{C}^{(t)})}{\partial \mathbf{C}^{(t)}} \right] \quad (6)$$

$$= \mathbf{C}^{(t)} - \eta \mathbb{E}_{z_e} \left[\delta_k^T \delta_k \mathbf{C}^{(t)} \right] + \eta \mathbb{E}_{z_e} \left[\delta_k^T z_e \right] \quad (7)$$

where $\delta_k^T \delta_k$ is the Kronecker delta matrix, defined as:

$$(\delta_k^T \delta_k)_{ij} = \begin{cases} 1 & \text{if } i = j = k, \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

All vectors in \mathbf{C} will be updated and utilized if and only if the expectation $\mathbb{E}_{z_e} [\delta_k^T \delta_k]$ converges to the identity matrix. Unlike variational autoencoders (VAEs) [14], which enforce a Gaussian distribution on the latent space via a

KL-divergence penalty, VQ models optimize z_e towards the selected codebook vectors $\mathbb{E}_{z_e} [\delta_k^T \delta_k \mathbf{C}]$. At the same time, the selected codebook vectors are optimized towards the distribution of z_e , resulting in the same selected subset of vectors moving closer to z_e , somewhat akin to a cocoon effect. However, this disjoint optimization of the codebook leads to part of the codebook, specifically $(\mathbf{I} - \mathbb{E}_{z_e} [\delta_k^T \delta_k]) \mathbf{C}$, remaining un-updated and underutilized once the optimization process begins. This phenomenon occurs because the optimization focuses only on a subset of codebook vectors, leaving other vectors stagnant.

This analysis reveals the fundamental cause of representation collapse in VQ models: the disjoint optimization process that updates only a subset of codebook vectors. This insight forms the basis for our proposed solution, SimVQ, which aims to address this issue by optimizing the entire latent space spanned by the codebook, rather than individual code vectors.

4. Addressing Collapse with Latent Linear Transformation

4.1. Reparameterize Codes with Latent Basis

Let $\mathbf{W} = \{\mathbf{w}_1, \dots, \mathbf{w}_n\}$ be a basis of a linear space. Any vector \mathbf{v} in the space can be uniquely expressed as a linear combination of the basis vectors with coefficients $c_1, \dots, c_n \in \mathbb{R}$:

$$\mathbf{v} = c_1 \mathbf{w}_1 + \dots + c_n \mathbf{w}_n = \mathbf{cW}. \quad (9)$$

Given the equivalence between \mathbf{v} and \mathbf{cW} in the linear space, we can reparameterize each vector in the codebook of VQ models with a new basis matrix $\mathbf{W} \in \mathbb{R}^{d \times d}$. Specifically, the codebook $\mathbf{C} = \{\mathbf{c}_1, \dots, \mathbf{c}_K\}$ can be reparameterized as:

$$\{\hat{\mathbf{c}}_1 \mathbf{W}, \dots, \hat{\mathbf{c}}_N \mathbf{W}\} = \hat{\mathbf{C}} \mathbf{W} \in \mathbb{R}^{K \times d}. \quad (10)$$

This reparameterization introduces two components: the basis matrix \mathbf{W} and the coefficient matrix $\hat{\mathbf{C}}$. In the following, we will discuss the optimization of both the basis matrix \mathbf{W} and the coefficient matrix $\hat{\mathbf{C}}$. For simplicity, we slightly abuse \mathbf{C} and $\hat{\mathbf{C}}$ below.

4.2. Asymmetric Optimization Dynamics

While it is commonly accepted that multiplying two linear matrices is equivalent to a single linear layer, we argue that the disjoint optimization problem of the codebook in VQ models can be addressed by linear transformation. In vanilla VQ models, only the codebook \mathbf{C} is responsible for minimizing commitment loss, leading to the disjoint optimization problem where only the selected codes will be updated.




Figure 2. (a): (left) The optimization trajectory of the objective $\|\mathbf{x} - \mathbf{q}\|_2^2$, which is the same as vanilla VQ. Only a small fraction of points are updated while others remain inactive. (b): (right) The optimization trajectory of the objective $\|\mathbf{x} - \mathbf{q}\mathbf{w}\|_2^2$ with \mathbf{q} frozen, which is the same as SimVQ. All the points are updated towards targets \mathbf{x} .




Figure 3. (a): (left) The optimization trajectory of the optimization objective: $\|\mathbf{x} - \mathbf{q}\mathbf{w}\|_2^2$ with both \mathbf{q} and \mathbf{w} unfrozen. (b): (right) The Frobenius norm of the projection matrix \mathbf{w} and loss curves. The loss quickly converges to 0 with \mathbf{w} almost unchanged.

In contrast, when the codebook is reparameterized as \mathbf{CW} , both the basis \mathbf{W} and the coefficient matrix \mathbf{C} contribute to minimizing the commitment loss. The gradients $\frac{\partial \mathcal{L}}{\partial \mathbf{W}}$ and $\frac{\partial \mathcal{L}}{\partial \mathbf{C}}$ can simultaneously reduce the loss. As a result, the optimization of the reparameterized codebook can be divided into three scenarios:

- Updating \mathbf{C} with \mathbf{W} frozen: Only the **selected** codes adapt to the latent distribution of z_e , as depicted on Fig. 1(a). The vanilla VQ is a special case with $\mathbf{W} = \mathbf{I}$.
- Updating \mathbf{W} with \mathbf{C} frozen: The **entire** codebook \mathbf{CW} adjusts to the latent distribution of z_e . The basis matrix \mathbf{W} rotates and stretches the space as shown in Fig. 1(b).
- Updating both \mathbf{C} and \mathbf{W} : The selected subset of codes moves towards z_e while the space spanned by \mathbf{W} undergoes simultaneous rotation and stretching.

To highlight the difference in optimization between \mathbf{C} and \mathbf{CW} , we conduct a toy experiment in a two-dimensional setting and visualize the optimization process in Fig. 2 and Fig. 3.

4.2.1. Toy Examples

We randomly sample two target points \mathbf{x} from Gaussian distribution as follows:

$$\mathbf{x}_1 \sim \mathcal{N}\left(\begin{pmatrix} 2 \\ 2 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\right), \quad \mathbf{x}_2 \sim \mathcal{N}\left(\begin{pmatrix} -2 \\ -2 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\right). \quad (11)$$

Then we initialize 10 learnable vectors \mathbf{q} from a Gaussian

Algorithm 1 Training Procedure for SimVQ

Input: Encoder f_θ , Decoder g_ϕ , Codebook $\mathbf{C} \in \mathbb{R}^{K \times d}$, Linear projector matrix \mathbf{W}_ψ , commitment weight β .
Output: Model parameters θ, ϕ, ψ and Codebook \mathbf{C} .
 Initialize Codebook with Gaussian distribution and **freeze** the parameter of Codebook \mathbf{C} ;
repeat
 Draw $x \sim p_{data}(\mathbf{x})$;
 $z_e = f_\theta(x)$;
/ Replace q_j in vanilla VQ with proposed $q_j \mathbf{W}_\psi$.*
 Nearest code search: $k = \arg \min_j \|z_e - q_j \mathbf{W}_\psi\|_2^2$, where $q_j \in \mathbf{C}$;
 Straight Through Estimation: $z_q = \text{sg}(q_k \mathbf{W}_\psi - z_e) + z_e$;
 $\hat{x} = g_\phi(z_q)$;
 Minimize $\mathcal{L}(\theta, \phi, \psi) = \text{MSE}(x, \hat{x}) + \beta \|z_e - q_k \mathbf{W}_\psi\|_2^2 + \|\text{sg}(z_e) - q_k \mathbf{W}_\psi\|_2^2$;
until converged

distribution:

$$\{q_i\}_{i=1}^{10} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\right), \quad (12)$$

During training with gradient descent, we introduce perturbation noise $\mathcal{N}(0, 0.01)$ to the targets. In Fig. 2(a), the optimization objective is similar to vanilla VQ: $\|\mathbf{x} - \mathbf{q}\|_2^2$. Only the nearest points q_4 and q_{10} are updated. In contrast, in Fig. 2(b), the optimization objective $\|\mathbf{x} - \mathbf{q}\mathbf{w}\|_2^2$ is similar to SimVQ with the points reparameterized by a learnable latent basis \mathbf{w} and \mathbf{q} frozen, resulting in the entire codebook $\{\mathbf{q}\}_{i=1}^{10}$ being *jointly* updated.

Remark 1. The simultaneous optimization of the latent basis \mathbf{w} and the coefficient matrix \mathbf{q} **may** lead to the collapse.

We provide an example in Fig. 3(a) where the optimization objective is $\|\mathbf{x} - \mathbf{q}\mathbf{w}\|_2^2$ with \mathbf{q} unfrozen this time. In the training process, only the nearest point q_1 and point q_{10} move towards the target point, while other points remain almost unchanged. We also visualize the loss curve in Fig. 3(b). The optimization objective with both \mathbf{q} and \mathbf{w} unfrozen converges quickly, where the norm of basis \mathbf{w} is much smaller than the objective with \mathbf{q} frozen. This indicates that the disjoint optimization of the codebook persists: \mathbf{q} can directly commit to the loss and dominate the optimization process, with \mathbf{w} being ignored, leading to the collapse quickly.

4.3. Joint Optimization of the Codebook

We propose SimVQ by simply using a learnable basis $\mathbf{W} \in \mathbb{R}^{d \times d}$ to reparameterize the codebook such that the codebook is transformed into $\mathbf{C}\mathbf{W}$. The pseudo-code for this

approach is provided in Algorithm 1. During training, we optimize only the latent basis matrix \mathbf{W} , while keeping the coefficient matrix \mathbf{C} frozen. The commitment loss for SimVQ is defined as:

$$\mathcal{L}_{\text{commit}}(z_e, q_k) = \|z_e - \delta_k \mathbf{C}\mathbf{W}\|_2^2. \quad (13)$$

The vanilla VQ model is a special case of SimVQ, where the linear basis matrix \mathbf{W} is fixed as the identity matrix \mathbf{I} . The update for \mathbf{W} with learning rate η is:

$$\mathbf{W}^{(t+1)} = \mathbf{W}^{(t)} - \eta \frac{\partial \mathcal{L}_{\text{commit}}(z_e, q_k)}{\partial \mathbf{W}^{(t)}} \quad (14)$$

$$= (\mathbf{I} - \eta \mathbb{E}_{z_e} [\mathbf{C}^T \delta_k^T \delta_k \mathbf{C}]) \mathbf{W}^{(t)} + \eta \mathbb{E}_{z_e} [\mathbf{C}^T \delta_k^T z_e]. \quad (15)$$

The term $\mathbb{E} [\mathbf{C}^T \delta_k^T \delta_k \mathbf{C}]$ represents the expectation of the quadratic form, and simplifies to $\mathbb{E}[\mathbf{q}_k^T \mathbf{q}_k]$. Since the codes are randomly sampled from a Gaussian distribution, we have:

$$\mathbb{E} [\mathbf{q}_k^T \mathbf{q}_k] = \mathbf{I}, \text{ where } \mathbf{q} \sim \mathcal{N}(0, 1), \quad (16)$$

which ensures that all elements of \mathbf{W} are updated. As training progresses, the latent basis \mathbf{W} converges to:

$$\lim_{t \rightarrow \infty} \mathbf{W}^{(t)} = \mathbb{E}_{z_e} [\mathbf{q}_k^T z_e] \quad (17)$$

Thus, in the limit:

$$\lim_{t \rightarrow \infty} \mathbf{q}_k \mathbf{W}^{(t)} = \mathbb{E} [\mathbf{q}_k \mathbf{q}_k^T \mathbf{e}] = \mathbb{E} [\mathbf{e}] \quad (18)$$

At convergence, the product $\mathbf{q}_k \mathbf{W}$ equals the nearest feature.

4.4. Efficiency Analysis

SimVQ demonstrates greater efficiency than vanilla VQ due to its asymmetric training strategy, wherein the codebook \mathbf{C} remains static and only the linear projection \mathbf{W} is optimized. This approach results in a significant reduction in memory usage during the gradient backpropagation process. In vanilla VQ, the memory cost for codebook optimization is $O(Kd)$, where K is the number of vectors in the codebook, and d is the dimension of each vector. In our experiments, $K = 65,536$ is much larger than $d = 128$. As the vocabulary size increases, the memory required for backpropagation grows proportionally, significantly impacting resource consumption. In contrast, SimVQ's memory cost for backpropagation is only $O(d^2)$ because the codebook \mathbf{C} is fixed, and only the linear layer \mathbf{W} is updated. This results in a constant memory requirement in backpropagation, independent of the vocabulary size. The $d \times d$ scaling becomes particularly advantageous as K increases in practical applications. This structural design minimizes the computational overhead and improves training efficiency, especially when dealing with large vocabularies.

Method	Latent dim	Codebook size	Util↑	rFID↓	LPIPS↓	PSNR↑	SSIM↑
VQGAN [9]	128	65,536	1.4%	3.74	0.17	22.20	70.6
VQGAN-EMA [19]	128	65,536	4.5%	3.23	0.15	22.89	72.3
VQGAN-FC [33]	128	65,536	1.4%	5.33	0.18	21.45	68.8
VQGAN-FC [33]	8	65,536	100.0%	2.63	0.13	23.79	77.5
FSQ [†] [16]	6	64,000	100.0%	2.80	0.13	23.63	75.8
LFQ [35]	16	65,536	100.0%	2.88	0.13	23.60	77.2
VQGAN-LC-CLIP ⁺ [42]	768	65,536	100.0%	2.40	0.13	23.98	77.3
SimVQ (ours)	128	65,536	100.0%	2.24	0.12	24.15	78.4
SimVQ (ours)	128	262,144	100.0%	1.99	0.11	24.68	80.3

Table 1. Reconstruction performance on ImageNet-1k with a resolution of 128×128 . All models are trained using images downsampled into 16×16 tokens. [†] Results are reproduced using the codebook size of $[8, 8, 8, 5, 5, 5]$ to approximately match 65,536. + Following VQGAN-LC, we extract CLIP features with the codebook frozen. The codebook utilization is calculated as the fraction of the codes that are activated at least once when encoding the validation set.

5. Experiments

To assess the efficacy and versatility of the proposed SimVQ, we conduct experiments across both image and audio modalities. Subsequently, we analyze the learned linear layer to investigate the latent basis. The experimental configurations are listed in Appendix 8.1.

5.1. Vision Modality

5.1.1. Baselines

Our baseline selection focuses on methods that enhance codebook utilization through architectural designs. We include VQGAN-FC [33], FSQ [16], LFQ [35] and VQGAN-LC-CLIP [42] as our primary baselines, as they represent the current state-of-the-art in improving reconstruction performance through codebook architecture innovations. Traditional VQ variants such as RVQ [15] and MoVQ [41] address fundamentally different technical challenges - they focus on training better discrete representations, rather than enhancing codebook utilization to improve reconstruction performance like our work. Another important research direction explores optimization-based solutions, where methods like stochastic quantization [26], distribution penalty [30, 32], and codebook reset [40] tackle codebook utilization through sophisticated training strategies. Given our focus on architectural innovations in codebook design, we evaluate SimVQ against methods that share this technical foundation for a meaningful assessment of our contributions.

5.1.2. Implementation Details

To rigorously evaluate the proposed SimVQ, we reproduce all the VQ models listed in Tab. 1 using the same architecture of VQGAN [9] with the quantization layer different only. Among the baselines, for VQGAN-FC [33], we follow the original setting to reduce the dimension of the

latent space to 8 followed by l_2 normalization to improve codebook utilization. For FSQ [16], we adopt a codebook size of $[8, 8, 8, 5, 5, 5]$ as recommended, to approximately match the default codebook size. For VQGAN-LC [42], we follow them and leverage an external pre-trained CLIP model to extract features of the training dataset in advance for a well-defined latent space. All models are trained on the ImageNet [7] dataset for 50 epochs with a batch size of 256. Input images are processed at a resolution of 128×128 pixels and downsampled by a factor of 8, yielding a feature map of $16 \times 16 \times 128$, where 128 is the dimension of the latent space. We set the default codebook size to a large number of $2^{16} = 65536$ rather than the traditional number 8192 to highlight the representation collapse problem. Performance is evaluated using rFID, LPIPS, PSNR, and SSIM metrics on the ImageNet validation set.

5.1.3. Main Results

Tab. 1 presents the reconstruction performance of various VQ models on image data. We make three key observations: 1) Traditional VQGAN models utilize only a very small subset of the codebook, with a utilization rate of just 1.4%. Although VQGAN-EMA is proposed to improve codebook utilization, especially when the codebook size scales up to 65k, it still suffers from severe representation collapse. 2) Recently proposed methods, such as LFQ, FSQ, and VQGAN-FC, effectively improve codebook utilization to 100%. However, these methods require reducing the latent space to a very low dimension. For example, applying VQGAN-FC to the standard latent dimension of 128 results in severe representation collapse and degraded reconstruction performance. Additionally, these models face limitations in model capacity due to the low-dimensional latent space. While they achieve full codebook utilization, their reconstruction quality on rFID score lags significantly behind SimVQ. 3) VQGAN-LC-CLIP leverages an exter-

Method	Bandwidth	Util↑	UTMOS↑	PESQ↑	STOI↑	V/UV F1↑
GT	-	-	4.06/3.48	-	-	-
EnCodec [6]	3.0kbps	-	2.31/2.09	2.05/2.05	0.90/0.88	0.92/0.89
Vocos [24]	3.0kbps	-	3.53/3.06	2.40/2.19	0.92/0.90	0.94/0.91
SpeechTokenizer [39]	3.0kbps	-	3.56/3.02	1.93/1.74	0.88/0.84	0.93/0.89
WavTokenizer [13]	0.9kbps	100/100%	3.74/3.43*	2.01/2.26*	0.89/0.89*	0.92/0.92*
SimVQ (ours)	0.9kbps	100.0/100.0%	4.00/3.51	2.33/2.08	0.91/0.88	0.94/0.91
WavTokenizer [13]	0.975kbps	68/-%	4.02/-	2.39/-	0.92/-	0.94/-
WavTokenizer [13]	1.05kbps	27/-%	4.00/-	2.36/-	0.81/-	0.94/-
SimVQ (ours)	0.975kbps	99.4/99.4%	4.03/3.52	2.42/2.15	0.92/0.88	0.94/0.92
SimVQ (ours)	1.2kbps	99.4/99.0%	4.03/3.52	2.54/2.26	0.93/0.90	0.94/0.92
SimVQ (ours)	1.35kbps	95.6/94.7%	4.03/3.53	2.61/2.31	0.93/0.90	0.95/0.93

Table 2. Reconstruction performance on LibriTTS test-clean/test-other dataset. * WavTokenizer is trained with a window size of 3 seconds. The bandwidth of 0.9kbps, 0.975kbps, 1.2kbps, 1.35kbps means the codebook size of 4096, 8192, 65536, 262144 respectively.

Method	Codebook Size	Util↑	rFID↓	LPIPS↓	PSNR↑	SSIM↑
VQGAN-LC-CLIP [†] [42]	50,000	99.9%	2.75	0.13	23.8	58.4
VQGAN-LC-CLIP [†] [42]	100,000	99.9%	<u>2.62</u>	0.12	23.8	58.9
VQGAN-LC-CLIP [†] [42]	200,000	99.8%	<u>2.66</u>	0.12	23.9	59.2
SimVQ	1,024	100.0%	3.67	0.16	22.34	70.8
SimVQ	8,192	100.0%	2.98	0.14	23.23	74.7
SimVQ	65,536	100.0%	2.24	0.12	24.15	78.4
SimVQ	262,144	100.0%	1.99	0.11	24.68	80.3

Table 3. Ablation study on the effect of various codebook sizes on ImageNet at a resolution of 128×128 . [†] We directly copy the reported results of VQGAN-LC from the original paper on ImageNet 256×256 resolution.

Initialization	Trainable	Util↑	rFID↓	LPIPS↓	PSNR↑	SSIM↑
Gaussian	Yes	100.0%	2.31	0.12	24.04	77.2
Uniform	No	100.0%	2.31	0.12	24.15	78.4
Gaussian	No	100.0%	2.24	0.12	24.15	78.4

Table 4. Ablation study of codebook optimization strategy.

nal pre-trained CLIP model to provide a well-defined latent space. However, VQGAN-LC relies on CLIP features pre-trained on much larger datasets than ImageNet, which introduces generalization issues and a lower performance ceiling (degradation issue in Tab. 3). In contrast, SimVQ can be applied to a wide range of data types and achieves superior performance (rFID 2.40 vs. 2.24) without the limitations imposed by a pre-trained feature extraction model.

5.1.4. Ablation Study

On the Codebook Size In Tab. 3, we explore the impact of different codebook sizes, ranging from $1k$ to $262k$, which is typically the level of LLM’s vocabulary size. SimVQ consistently improves performance as the codebook size increases. For instance, the rFID score decreases to 1.99, and SSIM surpasses 80.0. In contrast, while VQGAN-LC-CLIP can keep high codebook utilization as increasing codebook size, it encounters performance degradation, with the rFID

score worsening from 2.62 to 2.66 when the codebook size is increased from 100,000 to 200,000.

On the Codebook Optimization Strategy We investigate codebook initialization and the training of the linear layer in Tab. 4. Our findings are as follows: 1) The codebook is robust to different initialization strategies, yielding similar results with both Gaussian and uniform initialization. 2) When the codebook is updated during training, SimVQ continues to address the representation collapse issue, though with a slight degradation in performance.

5.2. Audio Modality

5.2.1. Baselines and Implementation Details

We use the LibriTTS dataset [36] for audio-based VQ model training. The baselines such as EnCodec [6], Vocos [24], and SpeechTokenizer [39] are based on residual vector quantization method. Our SimVQ model adopts the same architecture as WavTokenizer [13] with the only modification being the replacement of their EMA codebook with our one linear layer reparameterization method. We train SimVQ on LibriTTS-580h for 50 epochs with a batch size of 64. Note that WavTokenizer is trained with a 3-second window size for optimal performance, we train SimVQ using a 1-second window to accelerate training. For objective evaluation of the reconstructed audio, we follow Vocos [24] and employ metrics such as UTMOS [23], PESQ [20], STOI, and the F1 score for voiced/unvoiced classification (V/UV F1). UTMOS is particularly valuable as it produces scores highly correlated with human evaluations.

5.2.2. Main Results

Tab. 2 presents the reconstruction performance of various VQ models on audio data. Baseline models using residual vector quantization perform significantly worse than




Figure 4. (a):(left) The rank of latent basis matrix \mathbf{W} over training epochs. (b):(right) The Frobenius norm of latent basis matrix \mathbf{W} over training epochs.

SimVQ, even when utilizing much larger bandwidths. Despite using the same architecture as WavTokenizer, our model, which replaces the quantization layer with SimVQ, achieves superior performance with a 1-second window size and maintains nearly 100% codebook utilization when scaling up to a size of 262,144. The consistent performance of the SimVQ model across both image and audio data demonstrates that SimVQ is a general method for addressing the representation collapse problem in VQ models and can be effectively applied across multiple modalities.

5.3. Analysis

In Fig. 4(a), we plot the rank of the latent basis matrix over training epochs. Notably, SimVQ demonstrates the ability to adaptively adjust the rank of the latent space. Specifically, when the codebook size increases from 65k to 262k, the rank of the latent basis matrix decreases more rapidly and converges to a lower value. This observation suggests that a larger codebook can effectively alleviate the pressure on the latent space dimensionality, allowing the model to learn to represent data more efficiently. Additionally, despite the rank decreasing to a lower-rank space, SimVQ maintains 100% codebook utilization, highlighting its superiority over VQGAN-FC, which struggles when increasing the latent dimension from 8 to 128. We also calculate the Frobenius norm of the latent basis matrix, as shown in Fig. 4. The norm of a codebook size of 262k is slightly larger than for 65k, indicating that a larger codebook can span a broader area in the linear space. For a comprehensive evaluation, we also provide the reconstruction loss curve on the ImageNet validation dataset in Appendix 8.2. The results consistently show that SimVQ achieves improved performance, further validating the effectiveness of our approach.

5.4. Qualitative Evaluation

We visualize the distribution of encoder features and codebook embeddings in Fig. 5 and the frequency in Appendix 8.3. Compared to vanilla VQ models that most codebook vectors are not unused, SimVQ can update the whole codebook and distribute the features evenly in the whole space.




Figure 5. Visualization of the divergence between the encoder features and codebook embeddings on a random subset of ImageNet validation dataset. The left figure is vanilla VQ model and the right one is SimVQ.

We qualitatively compare the reconstruction quality of both images and audio in Appendix 8.4. SimVQ can preserve more details with an enlarged codebook size, such as “eyes” and “text”, which are challenging for vanilla VQ models.

6. Discussion

VQ Performance and Generative Models Many generative models [21, 25, 28, 34] utilize VQ models as a “tokenizer” to obtain discrete tokens. While it is intuitive that VQ reconstruction quality should impact generation performance, recent studies [35, 43] have revealed that the relationship is more nuanced: improved VQ reconstruction metrics do not necessarily translate to better generative outcomes. This complex relationship suggests that evaluating VQ models primarily through downstream generation tasks may not provide the most insightful assessment of their fundamental properties. Therefore, we focus on addressing the critical issue of representation collapse in VQ models, aiming to advance our understanding of their core representation learning mechanisms.

7. Conclusion

In this paper, we explore the representation collapse problem in VQ models. We conduct a theoretical analysis of the optimization process in VQ models and propose a simple yet effective method, SimVQ, to address this issue. Our method addresses the representation collapse by jointly optimizing the latent space through a linear transformation with one linear layer. Experimental results demonstrate that SimVQ outperforms previous approaches on both image and audio datasets, highlighting its broad applicability across diverse machine learning tasks.

Acknowledgement

This research was supported by the National Natural Science Foundation of China (Grant No.62276245).

References

[1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 1

[2] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. In *Advances in Neural Information Processing Systems*, pages 12449–12460. Curran Associates, Inc., 2020. 1

[3] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 3

[4] Zalán Borsos, Raphaël Marinier, Damien Vincent, Eugene Kharitonov, Olivier Pietquin, Matt Sharifi, Dominik Roblek, Olivier Teboul, David Grangier, Marco Tagliasacchi, and Neil Zeghidour. Audiolm: A language modeling approach to audio generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:2523–2533, 2023. 1

[5] Jake Bruce, Michael D Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, Yusuf Aytar, Sarah Maria Elisabeth Bechtle, Feryal Behbahani, Stephanie C.Y. Chan, Nicolas Heess, Lucy Gonzalez, Simon Osindero, Sherjil Ozair, Scott Reed, Jingwei Zhang, Konrad Zolna, Jeff Clune, Nando De Freitas, Satinder Singh, and Tim Rocktäschel. Genie: Generative interactive environments. In *Proceedings of the 41st International Conference on Machine Learning*, pages 4603–4623. PMLR, 2024. 1

[6] Alexandre Défossez, Jade Copet, Gabriel Synnaeve, and Yossi Adi. High fidelity neural audio compression. *Transactions on Machine Learning Research*, 2023. Featured Certification, Reproducibility Certification. 1, 7

[7] 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, 2009. 6

[8] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. 1

[9] 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 (CVPR)*, pages 12873–12883, 2021. 1, 2, 6

[10] Zhangyang Gao, Cheng Tan, Jue Wang, Yufei Huang, Lirong Wu, and Stan Z Li. Foldtoken: Learning protein language via vector quantization and beyond. *arXiv preprint arXiv:2403.09673*, 2024. 1

[11] Minyoung Huh, Brian Cheung, Pulkit Agrawal, and Phillip Isola. Straightening out the straight-through estimator: Overcoming optimization challenges in vector quantized net- works. In *Proceedings of the 40th International Conference on Machine Learning*, pages 14096–14113. PMLR, 2023. 3

[12] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. 3

[13] Shengpeng Ji, Ziyue Jiang, Xize Cheng, Yifu Chen, Minghui Fang, Jialong Zuo, Qian Yang, Ruiqi Li, Ziang Zhang, Xiaoda Yang, et al. Wavtokenizer: an efficient acoustic discrete codec tokenizer for audio language modeling. *arXiv preprint arXiv:2408.16532*, 2024. 7

[14] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2013. 3

[15] Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Autoregressive image generation using residual quantization. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11513–11522, 2022. 2, 6

[16] Fabian Mentzer, David Minnen, Eirikur Agustsson, and Michael Tschannen. Finite scalar quantization: VQ-VAE made simple. In *The Twelfth International Conference on Learning Representations*, 2024. 1, 3, 6

[17] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2

[18] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *Proceedings of the 38th International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021. 3

[19] 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*. Curran Associates, Inc., 2019. 1, 6

[20] A.W. Rix, J.G. Beerends, M.P. Hollier, and A.P. Hekstra. Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs. In *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No.01CH37221)*, pages 749–752 vol.2, 2001. 7

[21] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022. 1, 2, 8

[22] Aurko Roy, Ashish Vaswani, Arvind Neelakantan, and Niki Parmar. Theory and experiments on vector quantized autoencoders. *arXiv preprint arXiv:1805.11063*, 2018. 1, 3

[23] Takaaki Saeki, Detai Xin, Wataru Nakata, Tomoki Koriyama, Shinnosuke Takamichi, and Hiroshi Saruwatari. Utmos: Utokyo-sarulab system for voicemos challenge 2022. *ArXiv*, abs/2204.02152, 2022. 7

[24] Hubert Siuzdak. Vocos: Closing the gap between time-domain and fourier-based neural vocoders for high-quality audio synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. 7

[25] Peize Sun, Yi Jiang, Shoufa Chen, Shilong Zhang, Bingyue Peng, Ping Luo, and Zehuan Yuan. Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv preprint arXiv:2406.06525*, 2024. 1, 8

[26] Yuhta Takida, Takashi Shibuya, Weihsiang Liao, Chieh-Hsin Lai, Junki Ohmura, Toshimitsu Uesaka, Naoki Murata, Shusuke Takahashi, Toshiyuki Kumakura, and Yuki Mitsu-fuji. SQ-VAE: Variational Bayes on discrete representation with self-annealed stochastic quantization. In *Proceedings of the 39th International Conference on Machine Learning*, pages 20987–21012. PMLR, 2022. 1, 3, 6

[27] Chaofan Tao, Qian Liu, Longxu Dou, Niklas Muennighoff, Zhongwei Wan, Ping Luo, Min Lin, and Ngai Wong. Scaling laws with vocabulary: Larger models deserve larger vocabularies. *arXiv preprint arXiv:2407.13623*, 2024. 1

[28] Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024. 1, 2, 8

[29] Aaron van den Oord, Oriol Vinyals, and koray kavukcuoglu. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2017. 1, 2, 3

[30] Tung-Long Vuong, Trung Le, He Zhao, Chuanxia Zheng, Mehrtash Harandi, Jianfei Cai, and Dinh Phung. Vector quantized wasserstein auto-encoder. In *Proceedings of the 40th International Conference on Machine Learning*. JMLR.org, 2023. 1, 3, 6

[31] Chengyi Wang, Sanyuan Chen, Yu Wu, Ziqiang Zhang, Long Zhou, Shujie Liu, Zhuo Chen, Yanqing Liu, Huaming Wang, Jinyu Li, et al. Neural codec language models are zero-shot text to speech synthesizers. *arXiv preprint arXiv:2301.02111*, 2023. 1

[32] Pan Xiao, Peijie Qiu, and Aristeidis Sotiras. Scvae: Sparse coding-based variational autoencoder. *ArXiv*, abs/2303.16666, 2023. 1, 3, 6

[33] Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong Xu, Jason Baldridge, and Yonghui Wu. Vector-quantized image modeling with improved VQGAN. In *International Conference on Learning Representations*, 2022. 1, 3, 6

[34] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, Ben Hutchinson, Wei Han, Zarana Parekh, Xin Li, Han Zhang, Jason Baldridge, and Yonghui Wu. Scaling autoregressive models for content-rich text-to-image generation. *Transactions on Machine Learning Research*, 2022. Featured Certification. 1, 2, 8

[35] Lijun Yu, Jose Lezama, Nitesh Bharadwaj Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G Hauptmann, Boqing Gong, Ming-Hsuan Yang, Irfan Essa, David A Ross, and Lu Jiang. Language model beats diffusion - tokenizer is key to visual generation. In *The Twelfth International Conference on Learning Representations*, 2024. 1, 3, 6, 8

[36] Heiga Zen, Viet Dang, Rob Clark, Yu Zhang, Ron J Weiss, Ye Jia, Zhifeng Chen, and Yonghui Wu. Libritts: A corpus derived from librispeech for text-to-speech. *arXiv preprint arXiv:1904.02882*, 2019. 7

[37] Dong Zhang, Shimin Li, Xin Zhang, Jun Zhan, Pengyu Wang, Yaqian Zhou, and Xipeng Qiu. SpeechGPT: Empowering large language models with intrinsic cross-modal conversational abilities. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15757–15773, Singapore, 2023. Association for Computational Linguistics. 1

[38] Jiahui Zhang, Fangneng Zhan, Christian Theobalt, and Shijian Lu. Regularized vector quantization for tokenized image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18467–18476, 2023. 3

[39] Xin Zhang, Dong Zhang, Shimin Li, Yaqian Zhou, and Xipeng Qiu. Speechtokenizer: Unified speech tokenizer for speech language models. In *The Twelfth International Conference on Learning Representations*, 2024. 7

[40] Chuanxia Zheng and Andrea Vedaldi. Online clustered codebook. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22798–22807, 2023. 1, 3, 6

[41] Chuanxia Zheng, Guoxian Song, Tat-Jen Cham, Jianfei Cai, Dinh Q. Phung, and Linjie Luo. High-quality plurilistic image completion via code shared vqgan. *ArXiv*, abs/2204.01931, 2022. 2, 6

[42] Lei Zhu, Fangyun Wei, Yanye Lu, and Dong Chen. Scaling the codebook size of vqgan to 100,000 with a utilization rate of 99%. *ArXiv*, abs/2406.11837, 2024. 1, 2, 3, 6, 7

[43] Yongxin Zhu, Bocheng Li, Hang Zhang, Xin Li, Linli Xu, and Lidong Bing. Stabilize the latent space for image autoregressive modeling: A unified perspective. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 8

[44] Yongxin Zhu, Dan Su, Liqiang He, Linli Xu, and Dong Yu. Generative pre-trained speech language model with efficient hierarchical transformer. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1764–1775, Bangkok, Thailand, 2024. Association for Computational Linguistics. 1