

---

# SQ-VAE: Variational Bayes on Discrete Representation with Self-annealed Stochastic Quantization

---

Yuhta Takida<sup>1</sup> Takashi Shibuya<sup>1</sup> WeiHsiang Liao<sup>1</sup> Chieh-Hsin Lai<sup>1</sup> Junki Ohmura<sup>1</sup>  
Toshimitsu Uesaka<sup>1</sup> Naoki Murata<sup>1</sup> Shusuke Takahashi<sup>1</sup> Toshiyuki Kumakura<sup>2</sup> Yuki Mitsufuji<sup>1</sup>

## Abstract

One noted issue of vector-quantized variational autoencoder (VQ-VAE) is that the learned discrete representation uses only a fraction of the full capacity of the codebook, also known as codebook collapse. We hypothesize that the training scheme of VQ-VAE, which involves some carefully designed heuristics, underlies this issue. In this paper, we propose a new training scheme that extends the standard VAE via novel stochastic dequantization and quantization, called stochastically quantized variational autoencoder (SQ-VAE). In SQ-VAE, we observe a trend that the quantization is stochastic at the initial stage of the training but gradually converges toward a deterministic quantization, which we call self-annealing. Our experiments show that SQ-VAE improves codebook utilization without using common heuristics. Furthermore, we empirically show that SQ-VAE is superior to VAE and VQ-VAE in vision- and speech-related tasks.

## 1. Introduction

The use of a variational autoencoder (VAE) (Kingma & Welling, 2014; Higgins et al., 2017; Zhao et al., 2019) is one of the popular approaches to generative modeling. VAE consists of a pair of an encoder and a decoder, which are jointly trained by maximizing the evidence lower bound (ELBO) of the observed data (Jordan et al., 1999). The encoder maps the input data to a variable in a latent space, whereas the decoder converts the latent variable back into a sample of data space. A new sample is generated by decoding a latent variable that was sampled from a prior distribution.

Apart from VAE, a variant called Vector Quantized VAE (VQ-VAE) (van den Oord et al., 2017), shows its superiority

---

<sup>1</sup>Sony Group Corporation, Japan <sup>2</sup>Sony Corporation of America, USA. Correspondence to: Yuhta Takida <yuta.takida@sony.com>.

in several sample generation tasks (Dhariwal et al., 2020; Ramesh et al., 2021; Esser et al., 2021). In VQ-VAE, the encoded latent variables are quantized to their nearest neighbors in a learnable codebook, and the data samples are decoded from the quantized latent variables. Samples are generated by first sampling the discrete latent variables from an approximated prior, then the sampled latent variables are decoded into synthetic samples. The approximated prior can be, for example, a PixelCNN (van den Oord et al., 2016) that is trained on the latent space of training samples. Although VQ-VAE shares some similarities with VAE, its training does not follow the standard variational Bayes framework (Ghosh et al., 2020). Instead, it relies on carefully designed heuristics such as the use of a stop-gradient operator and the straight-through estimation of gradients. Even so, VQ-VAE often suffers from *codebook collapse*, which means that most of the codebook elements are not being used at all. This results in the deterioration of reconstruction accuracy (Kaiser et al., 2018). To address this problem, techniques such as the exponential moving average (EMA) (Polyak & Juditsky, 1992) update scheme, codebook reset (Dhariwal et al., 2020; Williams et al., 2020), and hyperparameter tuning are often employed (Roy et al., 2018).

We suspect that deterministic quantization is the cause of codebook collapse. Although the original approach with straight-through estimation is intuitive and elegant, some codebook elements can never be selected in cases of bad initialization. Therefore, we propose a framework that combines stochastic quantization and VAE, called stochastically quantized VAE (SQ-VAE)<sup>1</sup>. It can address the low codebook utilization issue of VQ-VAE and can be explained within the scope of the usual variational Bayes framework. Moreover, its training requires no exhaustive hyperparameter tuning and does not rely on heuristic techniques such as stop-gradient, codebook reset, or EMA update.

SQ-VAE introduces a pair of stochastic dequantization and quantization processes in the latent space. These processes are characterized by probability distributions with trainable parameters. This setup allows us to train the model within the usual variational Bayes framework without the need of conventional heuristics. Training the model requires only

---

<sup>1</sup>Our code is available at <https://github.com/sony/sqvae>.

one hyperparameter, which can be treated in straightforward ways as in (Jang et al., 2017). Optimizing the ELBO gradually reduces the stochasticity of the quantization process during the training, which we call *self-annealing*. In general, SQ-VAE does not impose any assumption on the data distribution; hence, we can model the stochastic quantization and dequantization processes via Gaussian distributions for example. However, we found that when the data distribution is categorical and cross entropy (CE) loss is used, this setup often yields an unsatisfactory performance. As a remedy, we propose the use of the von Mises–Fisher (vMF) distribution (see Appendix A) in place of the Gaussian distribution and show its effectiveness in Section 3.4.

We summarize our contributions below.

1. We propose SQ-VAE, which is variational autoencoder equipped with stochastic quantization and trainable posterior categorical distribution. SQ-VAE can be explained within an ordinary variational Bayes framework and may serve as a drop-in replacement of conventional VQ-VAE.
2. In SQ-VAE, the annealing of the stochasticity of the quantization process leads to a greater codebook utilization. We provide a theoretical insight into this self-annealing and validate it with an empirical study.
3. We design two instances of SQ-VAE: Gaussian SQ-VAE for general cases and vMF SQ-VAE specialized for categorical data distribution.
4. We evaluate SQ-VAE in vision- and speech-related generation tasks. The evaluation shows that SQ-VAE achieves better reconstruction than VQ-VAE. Furthermore, the performance of SQ-VAE can be improved by simply increasing the codebook size, which is not the case for VQ-VAE.

Throughout this paper, we use  $\{b_j\}_{j=1}^J$  to denote a set of elements  $b_j$ ;  $[N]$  denotes the set of positive integers less than or equal to  $N$ ; the capital letters  $P$  and  $Q$  denote the probability mass functions, whereas the lower case letters  $p$  and  $q$  denote the probability density functions.

## 2. Background

**VAE** Consider an observation  $\mathbf{x} \in \mathbb{R}^D$  and a target data distribution  $p_{\text{data}}(\mathbf{x})$ , which models finite samples. The standard VAE consists of a stochastic encoder–decoder pair: a decoder  $p_{\theta}(\mathbf{x}|\mathbf{z})$  and an approximated posterior  $q_{\phi}(\mathbf{z}|\mathbf{x})$ , where  $\theta$  and  $\phi$  are trainable parameters. The latent variables  $\mathbf{z} \in \mathbb{R}^{d_z}$  are assumed to follow a prior distribution  $p(\mathbf{z})$ . Data are generated by first sampling  $\mathbf{z}$  from the prior  $p(\mathbf{z})$  then obtaining  $\mathbf{x}$  by feeding  $\mathbf{z}$  into the stochastic decoder,  $p_{\theta}(\mathbf{x}|\mathbf{z})$ . The negative ELBO per sample  $\mathbf{x}$  is expressed as  $\mathcal{L}_{\text{VAE}} =$

$$\mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})} [-\log p_{\theta}(\mathbf{x}|\mathbf{z})] + D_{\text{KL}}(q_{\phi}(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z})). \quad (1)$$

To compute the ELBO of likelihood of samples  $\mathbf{x}$  analytically, the approximated posterior is usually modeled with conditional Gaussian as  $q_{\phi}(\mathbf{z}|\mathbf{x}) = \mathcal{N}(g_{\phi}(\mathbf{x}), \text{diag}(\sigma_{\phi}(\mathbf{x})))$  with two mappings  $g_{\phi} : \mathbb{R}^D \rightarrow \mathbb{R}^{d_z}$  and  $\sigma_{\phi} : \mathbb{R}^D \rightarrow \mathbb{R}^{d_z}$ .

If the target data distribution is continuous, the stochastic decoder can be modeled by a Gaussian distribution with a mapping  $f_{\theta} : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^D$  as

$$p_{\theta}(\mathbf{x}|\mathbf{z}) = \mathcal{N}(f_{\theta}(\mathbf{z}), \sigma^2 \mathbf{I}), \quad (2)$$

which reduces the first term in (1) into the mean squared error (MSE). In contrast, if the data distribution is discrete and has  $C_{\text{all}}$  categories, the stochastic decoder for the  $d$ th element of  $\mathbf{x}$ ,  $x_d$ , can be modeled as a categorical distribution with  $f_{\theta,d}^c : \mathbb{R}^{d_z} \rightarrow \mathbb{R}$  ( $c \in [C_{\text{all}}]$ ) as

$$P_{\theta}(x_d = c|\mathbf{z}) = \text{softmax}_c(\{f_{\theta,d}^c(\mathbf{z})\}_{c'=1}^{C_{\text{all}}}), \quad (3)$$

where the softmax is operated among  $c'$  (see Appendix A). In this case, the first term in (1) becomes the CE loss.

**VQ-VAE** In contrast to VAE, VQ-VAE consists of a deterministic encoder–decoder path and a trainable *codebook*. The codebook is a set  $\mathbf{B}$ , which contains  $K$   $d_b$ -dimensional vectors  $\{\mathbf{b}_k\}_{k=1}^K$ . A  $d_z$ -dimensional discrete latent space related to the codebook can be interpreted as the  $d_z$ -ary Cartesian power of  $\mathbf{B}$ ,  $\mathbf{B}^{d_z} \subset \mathbb{R}^{d_b \times d_z}$ . We denote a latent variable in  $\mathbf{B}^{d_z}$  and its  $i$ th column vector as  $\mathbf{Z}_q \in \mathbf{B}^{d_z}$  and  $\mathbf{z}_{q,i} \in \mathbf{B}$ , respectively. The deterministic encoding process from  $\mathbf{x}$  to  $\mathbf{Z}_q$  includes a mapping  $\hat{\mathbf{Z}}_q = g_{\phi}(\mathbf{x})$  with  $g_{\phi} : \mathbb{R}^D \rightarrow \mathbb{R}^{d_b \times d_z}$  and the quantization process of  $\hat{\mathbf{Z}}_q$  onto  $\mathbf{B}^{d_z}$ . The quantization process is modeled as a deterministic categorical posterior distribution, in which  $\hat{\mathbf{z}}_{q,i}$  is always mapped to its nearest neighbor  $\mathbf{z}_{q,i}$ , i.e.,  $\mathbf{z}_{q,i} = \arg \min_{\mathbf{b}_k} \|\hat{\mathbf{z}}_{q,i} - \mathbf{b}_k\|_2$ . The objective function of VQ-VAE is

$$\begin{aligned} \mathcal{L}_{\text{VQ}} = & -\log p_{\theta}(\mathbf{x}|\mathbf{Z}_q) + \|\text{sg}[g_{\phi}(\mathbf{x})] - \mathbf{Z}_q\|_F^2 \\ & + \beta \|g_{\phi}(\mathbf{x}) - \text{sg}[\mathbf{Z}_q]\|_F^2, \end{aligned} \quad (4)$$

where  $\text{sg}[\cdot]$  denotes the stop-gradient operator and  $\beta$  is set between 0.1 and 2.0 (van den Oord et al., 2017). To improve performance and convergence rate, EMA update is often applied only to the second term, which corresponds to the update of the codebook.

Note that the objective functions of VAE and VQ-VAE can both be interpreted as the sum of the reconstruction error and the latent regularization penalty.

## 3. Stochastically Quantized VAE

In this section, we propose SQ-VAE and its two instances, which are Gaussian SQ-VAE and vMF SQ-VAE. This framework bridges the training schemes of VAE and VQ-VAE. It relieves VQ-VAE from the heuristic techniques and reduces the difficulty of hyperparameter tuning. Moreover, it incorporates the self-annealing of a trainable categorical posterior distribution, which gradually approaches the deterministic quantization of VQ-VAE during the training. Furthermore, we provide theoretical and empirical support about the benefit of the self-annealing mechanism. The pseudo-codes

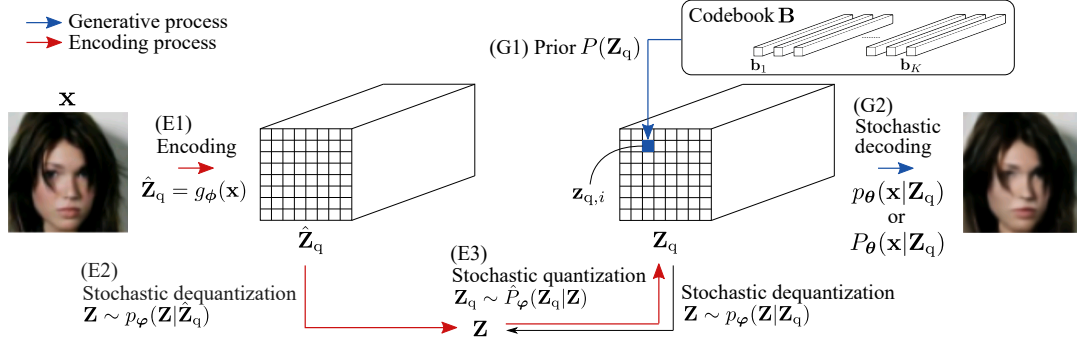


Figure 1. Encoding and generative processes of SQ-VAE. The encoding path from  $\mathbf{x}$  to  $\mathbf{Z}_q$  consists of (E1) deterministic encoding, (E2) stochastic dequantization, and (E3) quantization processes. For generation, in (G1) we first sample  $\mathbf{Z}_q \in \mathbf{B}^{d_z}$  from the prior  $p(\mathbf{Z}_q)$ . Then, in (G2) we feed  $\mathbf{Z}_q$  into the stochastic decoder to generate data samples.

of the two instances can be found in Appendix C.

### 3.1. Overview of SQ-VAE

The outline of SQ-VAE is shown in Figure 1. Identical to VQ-VAE, SQ-VAE also has a trainable codebook  $\mathbf{B} := \{\mathbf{b}_k\}_{k=1}^K$ . As a generative model, the goal of SQ-VAE is to learn a generative process  $\mathbf{x} \sim p_\theta(\mathbf{x}|\mathbf{Z}_q)$  with  $\mathbf{Z}_q \sim P(\mathbf{Z}_q)$  to generate samples that belong to the data distribution  $p_{\text{data}}(\mathbf{x})$ , where  $P(\mathbf{Z}_q)$  denotes the prior distribution of the discrete latent space  $\mathbf{B}^{d_z}$ . The prior  $P(\mathbf{Z}_q)$  is assumed to be an i.i.d. uniform distribution in the main training stage as in VQ-VAE, i.e.,  $P(\mathbf{z}_{q,i} = \mathbf{b}_k) = 1/K$  for  $k \in [K]$ . A second training will take place to learn  $P(\mathbf{Z}_q)$  after the main training stage. Since the exact evaluation of  $p_\theta(\mathbf{Z}_q|\mathbf{x})$  is intractable, the approximated posterior  $q_\phi(\mathbf{Z}_q|\mathbf{x})$  is used instead.

In this setup, although we can establish the generative process following that in VQ-VAE, the construction of the encoding process from  $\mathbf{x}$  to  $\mathbf{Z}_q$  is not straightforward owing to the discrete property of  $\mathbf{Z}_q$ . Therefore, we introduce two auxiliary variables to ease the explanation:  $\mathbf{Z}$  and  $\hat{\mathbf{Z}}_q$ .  $\mathbf{Z}$  is the continuous variable converted from  $\mathbf{Z}_q$  via the dequantization process  $p_\varphi(\mathbf{Z}|\mathbf{Z}_q)$ , where  $\varphi$  indicates its parameters. Furthermore, we may derive the inverse process of  $p_\varphi(\mathbf{Z}|\mathbf{Z}_q)$ , i.e., the stochastic quantization process  $\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})$ , from Bayes' theorem  $\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z}) \propto p_\varphi(\mathbf{Z}|\mathbf{Z}_q)P(\mathbf{Z}_q)$ . On the other hand,  $\hat{\mathbf{Z}}_q$  is defined as  $\hat{\mathbf{Z}}_q = g_\phi(\mathbf{x})$ , which is the output of the deterministic encoder  $g_\phi: \mathbb{R}^D \rightarrow \mathbb{R}^{d_b \times d_z}$  given a sample  $\mathbf{x}$ . Ideally,  $\hat{\mathbf{Z}}_q$  should be close to  $\mathbf{Z}_q$ . Similarly, the dequantization process of  $\hat{\mathbf{Z}}_q$  can be written as  $\mathbf{Z}|\hat{\mathbf{Z}}_q \sim p_\varphi(\mathbf{Z}|\hat{\mathbf{Z}}_q)$ . As in Figure 1, stacking the processes  $p_\varphi(\mathbf{Z}|\hat{\mathbf{Z}}_q)$  and  $\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})$  connects  $\hat{\mathbf{Z}}_q$  and  $\mathbf{Z}_q$ , and thus establishes the stochastic encoding process from  $\mathbf{x}$  to  $\mathbf{Z}_q$  as  $Q_\omega(\mathbf{Z}_q|\mathbf{x}) := \mathbb{E}_{q_\omega(\mathbf{Z}|\mathbf{x})}[\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})]$ , where  $\omega := \{\phi, \varphi\}$  and  $q_\omega(\mathbf{Z}|\mathbf{x}) := p_\varphi(\mathbf{Z}|g_\phi(\mathbf{x}))$ .

At this point, we can derive the ELBO for SQ-VAE as

$$\log p_\theta(\mathbf{x}) \geq -\mathcal{L}_{\text{SQ}}(\mathbf{x}; \theta, \omega, \mathbf{B}) :=$$

$$\begin{aligned} & \mathbb{E}_{q_\omega(\mathbf{Z}|\mathbf{x})\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})} \left[ \log \frac{p_\theta(\mathbf{x}|\mathbf{Z}_q)p_\varphi(\mathbf{Z}|\mathbf{Z}_q)P(\mathbf{Z}_q)}{q_\omega(\mathbf{Z}|\mathbf{x})\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})} \right] \\ & = \mathbb{E}_{q_\omega(\mathbf{Z}|\mathbf{x})\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})} \left[ \log \frac{p_\theta(\mathbf{x}|\mathbf{Z}_q)p_\varphi(\mathbf{Z}|\mathbf{Z}_q)}{q_\omega(\mathbf{Z}|\mathbf{x})} \right] \\ & \quad + \mathbb{E}_{q_\omega(\mathbf{Z}|\mathbf{x})} H(\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})) + \text{const.}, \end{aligned} \quad (5)$$

where  $H(P)$  denotes the entropy of  $P$ . In (5), since  $P(\mathbf{Z}_q)$  is assumed to follow a uniform distribution, it results into a constant term and is thus omitted. We hereafter omit the parameters of  $\mathcal{L}_{\text{SQ}}$  for simplicity. In the end, the main training is carried out by minimizing  $\mathbb{E}_{p_{\text{data}}(\mathbf{x})} \mathcal{L}_{\text{SQ}}(\mathbf{x})$ . The encoder, the decoder, and the codebook are all optimized simultaneously during the process. In this way, the codebook optimization no longer requires heuristic techniques such as the stop-gradient, EMA, and codebook reset (Dhariwal et al., 2020; Williams et al., 2020). The expectation in the first term of (5) involves the categorical distribution  $\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})$ , which can be approximated by the Gumbel–softmax relaxation (Jang et al., 2017; Maddison et al., 2017) to use the reparameterization trick in the backward pass of conventional VAE.

### 3.2. Gaussian SQ-VAE

We design Gaussian SQ-VAE by assuming that the dequantization process follows a Gaussian distribution. On the basis of the assumption, the dequantization process is modeled as

$$p_\varphi(\mathbf{z}_i|\mathbf{Z}_q) = \mathcal{N}(\mathbf{z}_{q,i}, \Sigma_\varphi), \quad (6)$$

in which  $\Sigma_\varphi$  is trainable. From Bayes' theorem, we may recover  $\mathbf{Z}_q$  with the inverse of (6), i.e., the stochastic quantization process, as  $\hat{P}_\varphi(\mathbf{z}_{q,i} = \mathbf{b}_k|\mathbf{Z}) =$

$$\text{softmax}_k \left( \left\{ -\frac{(\mathbf{b}_j - \mathbf{z}_i)^\top \Sigma_\varphi^{-1} (\mathbf{b}_j - \mathbf{z}_i)}{2} \right\}_{j=1}^K \right), \quad (7)$$

Table 1. Different parameterizations of the variance  $\Sigma_\varphi$  in Gaussian SQ-VAE.

	Variance $\Sigma_\varphi$	Unnormalized log-probability	Regularization objective $\mathcal{R}_\varphi^{\mathcal{N}}(\mathbf{Z}, \mathbf{Z}_q)$
(I)	$\sigma_\varphi^2 \mathbf{I}$	$\ \mathbf{b}_k - \mathbf{z}_i(\mathbf{x})\ _2^2 / 2\sigma_\varphi^2$	$\ \mathbf{Z} - \mathbf{Z}_q\ _F^2 / 2\sigma_\varphi^2$
(II)	$\sigma_\varphi^2(\mathbf{x}) \mathbf{I}$	$\ \mathbf{b}_k - \mathbf{z}_i(\mathbf{x})\ _2^2 / 2\sigma_\varphi^2(\mathbf{x})$	$\ \mathbf{Z} - \mathbf{Z}_q\ _F^2 / 2\sigma_\varphi^2(\mathbf{x})$
(III)	$\sigma_{\varphi,i}^2(\mathbf{x}) \mathbf{I}$	$\ \mathbf{b}_k - \mathbf{z}_i(\mathbf{x})\ _2^2 / 2\sigma_{\varphi,i}^2(\mathbf{x})$	$\sum_{i=1}^{d_z} \ \mathbf{z}_i(\mathbf{x}) - \mathbf{z}_{q,i}\ _2^2 / 2\sigma_{\varphi,i}^2(\mathbf{x})$
(IV)	$\text{diag}(\sigma_{\varphi,i}^2(\mathbf{x}))$	$\sum_{j=1}^{d_b} (b_{k,j} - z_{i,j}(\mathbf{x}))^2 / 2\sigma_{\varphi,i,j}^2(\mathbf{x})$	$\sum_{i=1}^{d_z} \sum_{j=1}^{d_b} (z_{i,j}(\mathbf{x}) - z_{q,i,j})^2 / 2\sigma_{\varphi,i,j}^2(\mathbf{x})$

where the unnormalized log-probabilities for  $\mathbf{b}_k$  in (7) correspond to Mahalanobis' distance from  $\mathbf{z}_i$  with the variance  $\Sigma_\varphi$ . We further consider several parameterizations of  $\Sigma_\varphi$  and summarize them in Table 1 with the corresponding unnormalized negative log-probabilities<sup>2</sup>. We examine their effectiveness in Section 5. The decoding and encoding setups of Gaussian SQ-VAE are described as follows.

**Decoding** The usual Gaussian setup is adopted in the decoding such that  $p_\theta(\mathbf{x}|\mathbf{Z}_q) = \mathcal{N}(f_\theta(\mathbf{Z}_q), \sigma^2 \mathbf{I})$ , where  $\sigma^2 \in \mathbb{R}_+$  and  $\theta$  are trainable parameters.

**Encoding** The encoding follows the process depicted in Figure 1, and the dequantization process applied to  $\hat{\mathbf{Z}}_q$  is  $p_\varphi(\mathbf{z}_i|\hat{\mathbf{Z}}_q) = \mathcal{N}(\hat{\mathbf{z}}_{q,i}, \Sigma_\varphi)$ .

**Objective Function** The substitution of the encoding and decoding processes above into (5) gives  $\mathcal{L}_{\mathcal{N}\text{-SQ}} =$

$$\mathbb{E}_{q_\omega(\mathbf{z}|\mathbf{x})\hat{P}_\varphi(\mathbf{z}_q|\mathbf{z})} \left[ \frac{1}{2\sigma^2} \|\mathbf{x} - f_\theta(\mathbf{Z})\|_2^2 + \mathcal{R}_\varphi^{\mathcal{N}}(\mathbf{Z}, \mathbf{Z}_q) \right] - \mathbb{E}_{q_\omega(\mathbf{z}|\mathbf{x})} H(\hat{P}_\varphi(\mathbf{z}_q|\mathbf{z})) + \frac{D}{2} \log \sigma^2 + \text{const.}, \quad (8)$$

where  $\mathcal{R}_\varphi^{\mathcal{N}}(\mathbf{Z}, \mathbf{Z}_q)$  denotes the regularization objective in Table 1, depending on the parameterization of  $\Sigma_\varphi$ . The derivation detail can be found in Appendix B.1.

### 3.3. Self-annealed Quantization

Before proposing the next SQ-VAE instance, we would like to demonstrate the effectiveness of trainable parameters in (de)quantization processes. In this subsection, we adopt the parameterization  $\Sigma_\varphi = \sigma_\varphi^2 \mathbf{I}$  (type I in Table 1) for simplicity.

According to (7),  $\Sigma_\varphi$  controls the degree of stochasticity of the quantization during the training. We first consider two extreme cases,  $\sigma^2 \rightarrow \infty$  and  $\sigma^2 \rightarrow 0$ , with the following proposition whose proof is given in Appendix D.2.

**Proposition 1.** *Assume that  $p_{\text{data}}(\mathbf{x})$  has finite support, whereas  $g_\phi$  and  $\{\mathbf{b}_k\}_{k=1}^K$  are bounded. Let  $\omega^* = \{\phi^*, \varphi^*\}$  be a minimizer of  $\mathbb{E}_{p_{\text{data}}(\mathbf{x})} D_{\text{KL}}(Q_\omega(\mathbf{Z}_q|\mathbf{x}) \parallel P_\theta(\mathbf{Z}_q|\mathbf{x}))$  with fixed  $\theta$ ,  $\sigma^2$  and  $\{\mathbf{b}_k\}_{k=1}^K$ . If  $\sigma^2 \rightarrow 0$ , then  $\sigma_{\varphi^*}^2 \rightarrow 0$ .*

When  $\sigma^2 \rightarrow \infty$ , the first term in (8) diminishes. It is mini-

<sup>2</sup>Although more complicated parameterizations exist, they often lead to an unstable optimization in our experiments.

mized when  $\sigma_\varphi^2 \rightarrow \infty$ , where  $P_\varphi(\mathbf{z}_{q,i} = \mathbf{b}_k|\mathbf{Z})$  approaches a uniform distribution. On the other hand, according to Proposition 1, when  $\sigma^2 \rightarrow 0$ , it leads to  $\sigma_\varphi^2 \rightarrow 0$ . This means that  $P_\varphi(\mathbf{z}_{q,i} = \mathbf{b}_k|\mathbf{Z})$  converges to the Kronecker delta function  $\delta_{k,\hat{k}}$ , where  $\hat{k} = \arg \min_k \|\mathbf{z}_i - \mathbf{b}_k\|_2$ . This deterministic quantization is exactly the posterior categorical distribution of VQ-VAE. According to the two cases above, if  $\sigma^2$  decreases gradually during the training, the quantization process will also gradually decrease its stochasticity and approach the deterministic quantization. We refer to this as *self-annealing*.

**Dynamics of the variance parameter** To verify whether *self-annealing* happens during the training, we conduct an experiment on MNIST (LeCun et al., 1998). We train Gaussian SQ-VAE with  $\Sigma_\varphi = \sigma_\varphi^2 \mathbf{I}$ . As targets for comparison, we also train models with  $\sigma_\varphi^2$  fixed to a designated  $\sigma_q^2$ . The details of the experimental setup can be found in Appendix D.3. The results are summarized in Figure 2.

In Figure 2(a), as the epoch grows,  $\sigma_\varphi^2$  decreases along with  $\sigma^2$ , which agrees with Proposition 1 and our expectation. As shown in Figure 2(b), with trainable  $\sigma_\varphi^2$ , the average entropy decreases as the training progresses. These two results suggest that self-annealing occurs in practical situations. On the other hand, Figure 2(b) indicates that with a fixed  $\sigma_q^2$ , the average entropy stays relatively constant. Moreover, as shown in Figure 2(c), the MSE is greatly affected by the  $\sigma_q^2$  selected. Although there is an optimum for the fixed  $\sigma_q^2$ , the trainable  $\sigma_\varphi^2$ , which is indicated by the blue line, achieves the lowest MSE among all cases. Therefore, we showed that the stochastic quantization and self-annealing together yields a codebook that effectively covers a larger support in the latent space, especially in the beginning of the training stage. This leads to the improvement of reconstruction accuracy, which will be demonstrated further in Section 5.

### 3.4. vMF SQ-VAE for Categorical Distributions

An intuitive way to adapting SQ-VAE for categorical data distribution is to model the decoder output as a categorical distribution as (3). Consider a typical classification scenario that the last layer of a decoder is a linear layer followed by a softmax. The decoder can be represented as the combination of the linear layer  $\mathbf{w}_{\text{last},c} \in \mathbb{R}^F$  and the rest  $\tilde{f}_{\theta^-,d}^{\text{rest}} : \mathbf{B}^{d_z} \rightarrow \mathbb{R}^F$ . It becomes  $f_{\theta^-,d}^c(\mathbf{Z}_q) = \mathbf{w}_{\text{last},c}^\top \tilde{f}_{\theta^-,d}^{\text{rest}}(\mathbf{Z}_q)$ , where  $\theta^-$  denotes the trainable parameters excluding  $\mathbf{w}_{\text{last},c}$ . We may represent

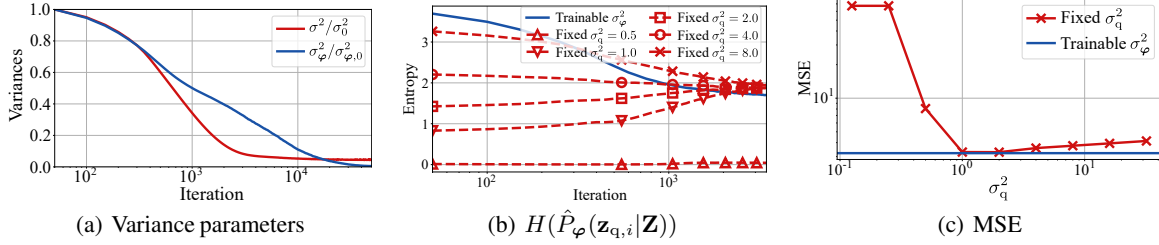


Figure 2. Empirical study on the dynamics related to  $\sigma_\varphi^2$  in Section 3.3. (a) The variance parameter  $\sigma_\varphi^2$  (blue) decreased with  $\sigma^2$  (red), where  $\sigma_0^2$  and  $\sigma_{\varphi,0}^2$  are their initial values. (b) Average entropy of the quantization process w.r.t. the iteration, which is obtained by Monte Carlo estimation. (c) MSE for trainable  $\sigma_\varphi^2$  and various values of  $\sigma_q^2$  on the test set.

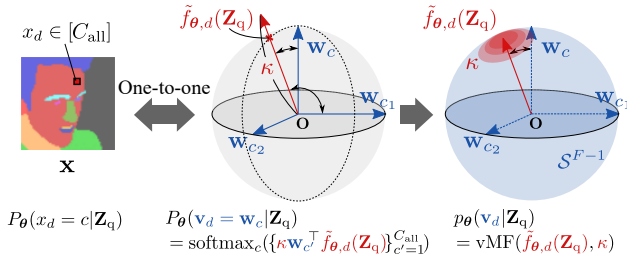


Figure 3. vMF decoder.

the ELBO of this model in terms of the decomposition as

$$\begin{aligned} \mathcal{L}_{\text{CE-SQ}}^{\text{naive}} = & \mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{x})\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})} \left[ - \sum_{d=1}^D \log(P_{\theta}(x_d = c|\mathbf{Z}_q)) \right. \\ & \left. + \mathcal{R}_{\varphi}^{\mathcal{N}}(\mathbf{Z}, \mathbf{Z}_q) \right] - \mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{x})} H(\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})) + \text{const.} \end{aligned} \quad (9a)$$

with

$$P_{\theta}(x_d = c|\mathbf{Z}_q) = \text{softmax}_c \left( \{ \mathbf{w}_{\text{last},c}^{\top} \tilde{f}_{\theta^-,d}(\mathbf{Z}_q) \}_{c=1}^{C_{\text{all}}} \right). \quad (9b)$$

However, we found that the performance of this *Naive categorical (NC) SQ-VAE* is often unsatisfactory, as shown in Section 5.2. A possible cause can be found by observing the difference between (8) and (9a). In (9a), owing to the replacement of Gaussian with categorical distribution, trainable parameters such as  $\sigma^2$  no longer exist in the objective function. This means that the model cannot be benefited from the self-annealing effect.

To gain the advantage from self-annealing, we introduce the vMF distribution to refine the model as in Figure 3, and we call it vMF SQ-VAE. Consider a hypersphere  $S^{F-1}$  that is embedded in an  $F$ -dimensional space. Let  $\mathbf{w}_c$  denote the projection vector<sup>3</sup> of the  $c$ th data category on the surface of

<sup>3</sup>In the current implementation,  $\{\mathbf{w}_c\}_{c=1}^{C_{\text{all}}}$  are predefined and distributed on  $S^{F-1}$ . We will explore its optimization in future work.

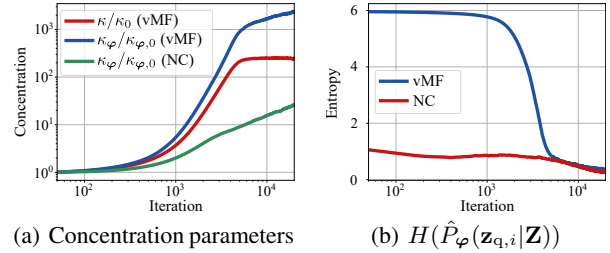


Figure 4. Comparison between vMF and NC decoders: (a) The concentration parameter of vMF decoder  $\kappa_{\varphi}$  increases with  $\kappa$ , whereas the growth of  $\kappa_{\varphi}$  of the NC decoder is relatively small. Here,  $\kappa_0$  and  $\kappa_{\varphi,0}$  indicate initial values. (b) Average entropy of probabilities of quantization processes.

$S^{F-1}$ . Next, we represent the projection of data  $x_d$  on the hypersphere as  $\mathbf{v}_d \in \{\mathbf{w}_c\}_{c=1}^{C_{\text{all}}}$ . If  $x_d$  belongs to a category  $c$ , that is,  $x_d = c$ , then  $\mathbf{v}_d = \mathbf{w}_c$  and vice versa.

**Decoding** The first step is to decode  $\mathbf{Z}_q$  into  $\mathbf{V} := \{\mathbf{v}_d\}_{d=1}^D$  with the decoder  $\tilde{f}_{\theta,d} : \mathbf{B}^{dz} \rightarrow S^{F-1}$ . Then, determine the probability of  $\mathbf{v}_d = \mathbf{w}_c$  with a trainable scalar  $\kappa \in \mathbb{R}_+$  by using

$$P_{\theta}(\mathbf{v}_d = \mathbf{w}_c|\mathbf{Z}_q) = \text{softmax}_c \left( \left\{ \kappa \mathbf{w}_c^{\top} \tilde{f}_{\theta,d}(\mathbf{Z}_q) \right\}_{c=1}^{C_{\text{all}}} \right), \quad (10)$$

which resembles the categorical decoder in (9b) except for the normalization onto  $S^{F-1}$  and the scaling factor  $\kappa$ . Therefore, we may represent the categorical probabilities for the decoded  $\mathbf{Z}_q$  as

$$p_{\theta}(\mathbf{v}_d|\mathbf{Z}_q) \propto \exp \left( \kappa \mathbf{v}_d^{\top} \tilde{f}_{\theta,d}(\mathbf{Z}_q) \right). \quad (11)$$

By normalizing (11) w.r.t.  $\mathbf{v}_d$  over  $S^{F-1}$ , we obtain  $p_{\theta}(\mathbf{v}_d|\mathbf{Z}_q) = \text{vMF}(\tilde{f}_{\theta,d}(\mathbf{Z}_q), \kappa)$ , where  $\tilde{f}_{\theta,d}(\mathbf{Z}_q)$  and  $\kappa$  correspond to the mean direction and the concentration parameter of the vMF distribution, respectively.

**Encoding** Accordingly, we model the stochastic dequantization process of the encoder with the vMF distribution:

$$p_\varphi(\mathbf{z}_i|\mathbf{Z}_q) = \text{vMF}(\mathbf{z}_{q,i}, \kappa_\varphi), \quad (12)$$

where  $\kappa_\varphi$  is the trainable concentration parameter<sup>4</sup>. Similarly to Gaussian SQ-VAE in Section 3.2, the discrete  $\mathbf{Z}_q$  is recovered using Bayes’ theorem as

$$\hat{P}_\varphi(\mathbf{z}_{q,i} = \mathbf{b}_k|\mathbf{Z}) = \text{softmax}_k(\{\kappa_\varphi \mathbf{b}_j^\top \mathbf{z}_i\}_{j=1}^K), \quad (13)$$

where the unnormalized log-probabilities of  $\mathbf{b}_k$  in (13) correspond to the  $\kappa_\varphi$ -scaled cosine similarity between  $\mathbf{b}_k$  and  $\mathbf{z}_i$ .

**Objective Function** Substituting the encoding and decoding processes into (5) leads to  $\mathcal{L}_{\text{vMF-SQ}} =$

$$\begin{aligned} \mathbb{E}_{q_\omega(\mathbf{z}|\mathbf{x})\hat{P}_\varphi(\mathbf{z}_q|\mathbf{z})} & \left[ -\kappa \sum_{d=1}^D \mathbf{v}_d^\top \tilde{f}_{\theta,d}(\mathbf{Z}_q) + \mathcal{R}_\varphi^{\text{vMF}}(\mathbf{Z}, \mathbf{Z}_q) \right] \\ & - \mathbb{E}_{q_\omega(\mathbf{z}|\mathbf{v})} H(\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})) - \log C_F(\kappa) + \text{const.}, \end{aligned} \quad (14)$$

where  $\mathcal{R}_\varphi^{\text{vMF}}(\mathbf{x}, \mathbf{Z}_q)$  is a regularization objective defined by  $\mathcal{R}_\varphi^{\text{vMF}}(\mathbf{Z}, \mathbf{Z}_q) = \sum_{i=1}^{d_z} \kappa_{\varphi,i} (1 - \mathbf{z}_{q,i}^\top \mathbf{z}_i)$  (see Appendix B.2 for details). Here,  $C_F(\kappa)$  denotes the normalizing constant of the vMF distribution (see Appendix A).

**Comparing vMF SQ-VAE with Naïve Categorical SQ-VAE** In (14), the first two terms are scaled with  $\kappa$  and  $\kappa_\varphi$ . Furthermore, vMF SQ-VAE has a property that, if  $\kappa \rightarrow \infty$ , then  $\kappa_{\varphi^*} \rightarrow \infty$ . Its proof can be done similarly to Proposition 1 via setting  $\kappa = 1/\sigma^2$  and  $\kappa_{\varphi^*} = 1/\sigma_{\varphi^*}^2$ . As a result, vMF SQ-VAE can also achieve self-annealing as described in Section 3.3 if  $\kappa \rightarrow \infty$ . In the experiment on CelebAHQ-Mask (Lee et al., 2020) in Section 5.2,  $\kappa_\varphi$  increases together with  $\kappa$  as training progresses, as shown in Figure 4. On the other hand, self-annealing is impossible for NC SQ-VAE owing to the lack of scaling parameters.

## 4. Related Work

**Latent Vector Quantization** This is a common approach in VQ-VAE, end-to-end image compression (Toderici et al., 2016; Theis et al., 2017) and many other studies. van den Oord et al. (2017) trained VQ-VAE with deterministic quantization. They approximated the gradient of the quantization process with the straight-through estimator (Bengio et al., 2013) and utilized the stop-gradient operator. Some VQ-VAE-based models also adopted deterministic quantization (Razavi et al., 2019; Dhariwal et al., 2020).

The common form of stochastic quantization  $Q(\mathbf{z}_{q,i} = \mathbf{b}_k) \propto \exp(-\|\mathbf{b}_k - \mathbf{z}_i\|_2^2)$  has been applied in sev-

<sup>4</sup> $\kappa_\varphi$  can be dependent on either  $\mathbf{x}$  or  $i$ . However, we choose  $\kappa_\varphi$  to be the independent variable.

eral VQ-VAE-based models (Roy et al., 2018; Williams et al., 2020). Roy et al. (2018) connected the EMA update of the codebook with an expectation maximization (EM) algorithm and softened the EM algorithm with a stochastic posterior. Williams et al. (2020) adopted a hierarchical VQ-VAE with a stochastic posterior to compress images at extremely low bit rates. The stochastic quantization schemes in previous literature did not involve trainable parameters in their categorical posterior, which correspond to the schemes with fixed  $\sigma_q^2$  in Section 3.3. We emphasize the fact that our categorical posterior including trainable parameters achieves a higher reconstruction accuracy with the help of self-annealing.

Agustsson et al. (2017) proposed a controllable quantization scheme for image compression. In their work, annealing is achieved with a predefined hyperparameter scheduling. However, they mentioned that the annealing rate must be controlled carefully, otherwise, either the annealing cannot progress or the gradient will vanish in the early stage. In contrast, our proposed method does not rely on such tuning.

Wu & Flierl (2020) proposed a bottleneck regularizer of the latent space in VAE to extract meaningful representations for downstream tasks. Although the regularizer is *inspired from* the stochastic quantization of latent variables, how the method relates continuous latent features with a discrete codebook is beyond the scope of this work.

**Learning Categorical Distributions with VAE** Kingma & Welling (2014) first proposed a VAE with the Bernoulli decoder for binary data. The Bernoulli decoder can be easily generalized into a multi-class categorical decoder as in (3). This categorical decoder has been adopted to the VQ-VAE by Chorowski et al. (2019).

Polykovskiy & Vetrov (2020) proposed a deterministic decoding scheme for discrete data. The categorical probability of its decoder has been limited to be one-hot by applying an argmax operator. The non-differentiability of argmax is solved by the smooth relaxation of the argmax operator with a temperature parameter. In practice, the hyperparameter is manually annealed to 0 in the training phase. By comparing vMF SQ-VAE with this work, we found that the scaling parameter  $\kappa$  is self-annealed in the training without a predefined scheduling scheme.

## 5. Experiments

We apply SQ-VAE in several vision- and speech-related tasks to demonstrate its improvement over the conventional VQ-VAE and VAE. All the experiments are repeated with three different random seeds, unless otherwise stated. Moreover, all the models including VAE are adapted to ensure that they all use the same amount of bits to represent an encoded input. In our experiments, we approximate the categorical distributions  $\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})$  included in the expectation operator of (1) with Gumbel–softmax distributions. The gap between the two distributions is gradually reduced by annealing

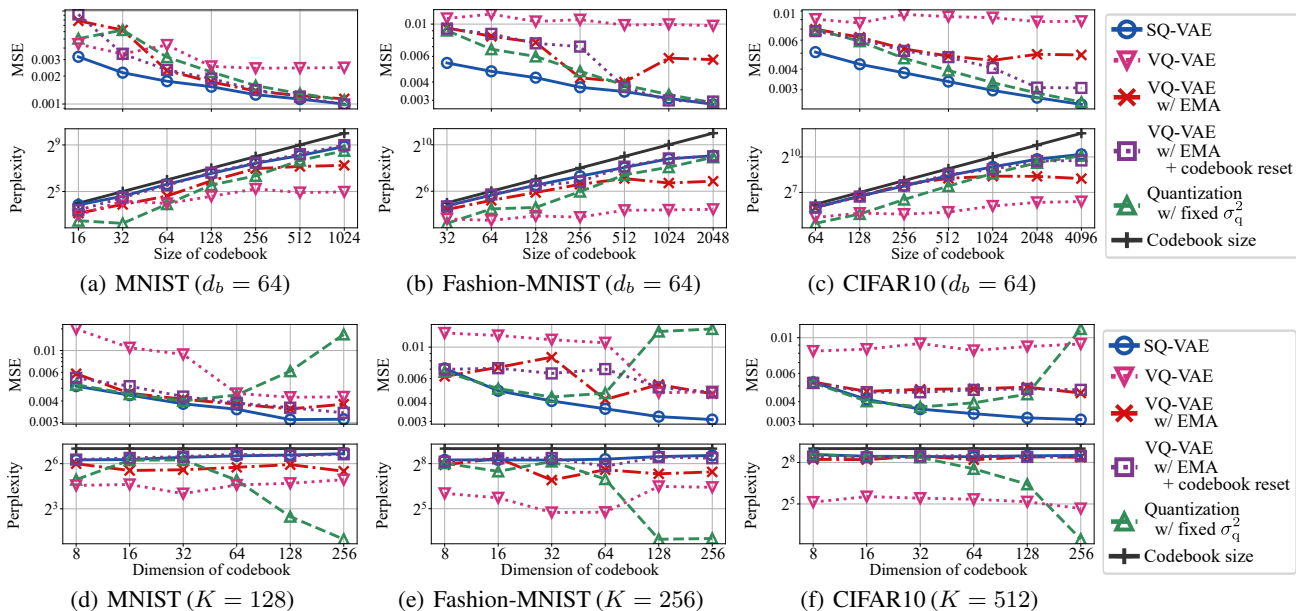


Figure 5. Empirical studies on the impact of codebook capacity examined on MNIST Fashion-MNIST and CIFAR10. (a)–(c) The size  $K$  is swept with the dimension  $d_b$  fixed to 64. (d)–(f) Various  $d_b$  values are tested with the size  $K$  fixed as 128, 256, and 512, respectively. The black lines with “+” marks indicate the upper bounds of the perplexities, i.e.,  $K$ . All the y-axes are in log-scale.

the temperature parameter of Gumbel–softmax using the predefined schedule in Jang et al. (2017). This temperature parameter does not affect the categorical probabilities of the original distribution. The effect of  $\sigma_\varphi^2$  (or  $\kappa_\varphi$ ) is definitely different from that of this temperature parameter (see Section 3.3). We refer to Appendix E for the detailed setup of evaluation and supplemental experiments.

### 5.1. Continuous Data Distribution

**Vision** First, we compare the reconstruction accuracy and codebook utilization of SQ-VAE (I) with VQ-VAEs (van den Oord et al., 2017) under different codebook capacity settings, i.e., the codebook dimension  $d_b$  and the codebook size  $K$ . The reconstruction accuracy is measured by MSE, and the codebook utilization is measured by the perplexity of latent variables. This experiment is performed on MNIST, Fashion-MNIST (Xiao et al., 2017) and CIFAR10 (Krizhevsky et al., 2009). As a target for comparison, we also train models with fixed stochastic quantization  $P_\varphi(\mathbf{z}_{q,i} = \mathbf{b}_k) \propto \exp(-\|\mathbf{b}_k - \mathbf{z}_i\|_2^2 / 2\sigma_q^2)$  (Roy et al., 2018; Williams et al., 2020). We choose  $\sigma_q^2 = 1.0$ , which achieves the best MSE with  $K = 64$  in our preliminary results (see Section 3.3). In Figure 5, SQ-VAE achieved the lowest MSE with the highest perplexity in most of settings. Moreover, its performance is proportional to the size and dimension of the codebook. On the other hand, the performance of VQ-VAE is less correlated with the codebook settings and therefore suggests the need for careful tuning. Moreover, even when all the heuristic techniques are applied to VQ-VAE, SQ-VAE still outperforms VQ-VAE in

terms of MSE, especially when the codebook size is small.

We also test SQ-VAE with various parameterizations on CelebA  $64 \times 64$  (Liu et al., 2015) with the codebook capacity set to  $(n_b, K) = (64, 512)$ .<sup>5</sup> A separately trained PixelCNN (van den Oord et al., 2016) is used for each model. To keep a fair comparison, these PixelCNNs are trained such that they have the same test log-likelihood. MSE and Fréchet Inception (FID) (Heusel et al., 2017) are used as the metrics for quality assessment. To examine the feasibility of the latent space, in addition to the reconstruction and the generation, we apply latent manipulations to latent variables and evaluate the FID of the images reconstructed with modified latent variables. The first manipulation is replacing latent vector elements by the  $k$ th nearest codebook elements, recorded as Neighbor- $k$ . The second manipulation is that we apply linear interpolation of pairs of encoded points then project the interpolated vectors to their nearest codes. Their mixing ratios are randomly distributed within  $[0, 1]$ . If the codebook elements are evenly distributed within the support of the dataset in the latent space, the resulting FID would be less affected by these manipulations.

From Table 2, SQ-VAE achieves the best performance against all the baselines. It also shows that the training of SQ-VAE is stable even for complicated parameterization such as Gaussian SQ-VAE (III). On the other hand, the simplest Gaussian SQ-VAE (I) still yields performance improvement over the fixed stochastic quantization scheme. In addition, the result

<sup>5</sup>We test SQ-VAE (I–III) models in this experiment because the training of SQ-VAE (IV) model is unstable on the CelebA dataset.

Table 2. Evaluation on CelebA. The MSE ( $\times 10^3$ ) and reconstructed FID (rFID) are evaluated using the test set. The codebook capacity for the discrete latent space is set to  $(n_b, k) = (64, 512)$ . The Roman numerals for Gaussian SQ-VAEs correspond to those in Table 1. We also show the FID of samples generated with a prior learned with PixelCNN.

Model	Reconstruction		Generation (FID)	Latent manipulation (FID)			
	MSE	rFID		Neighbor-3	Neighbor-5	Neighbor-10	Interp.
VAE	4.79 $\pm$ 0.01	40.3 $\pm$ 0.3	–	–	–	–	–
VQ-VAE (EMA)	1.33 $\pm$ 0.41	18.5 $\pm$ 5.1	42.0 $\pm$ 11.5	31.9 $\pm$ 14.8	42.8 $\pm$ 20.7	70.7 $\pm$ 35.4	28.2 $\pm$ 6.4
VQ-VAE (EMA+code reset)	1.62 $\pm$ 0.36	22.0 $\pm$ 5.9	51.8 $\pm$ 10.8	39.7 $\pm$ 12.0	52.7 $\pm$ 14.7	83.2 $\pm$ 20.4	32.6 $\pm$ 7.1
Quantization w/ fixed $\sigma_q^2$	1.09 $\pm$ 0.01	15.9 $\pm$ 0.1	38.2 $\pm$ 0.9	20.0 $\pm$ 0.4	26.4 $\pm$ 0.8	41.5 $\pm$ 2.1	18.6 $\pm$ 0.3
Gaussian SQ-VAE (I)	<b>0.96</b> $\pm$ 0.01	14.8 $\pm$ 0.3	28.2 $\pm$ 0.9	17.8 $\pm$ 0.1	21.9 $\pm$ 0.1	<b>33.1</b> $\pm$ 0.3	<b>17.6</b> $\pm$ 0.6
Gaussian SQ-VAE (II)	0.98 $\pm$ 0.01	14.3 $\pm$ 0.2	<b>27.7</b> $\pm$ 1.1	17.8 $\pm$ 0.2	22.2 $\pm$ 0.4	34.0 $\pm$ 0.9	<b>17.6</b> $\pm$ 0.1
Gaussian SQ-VAE (III)	<b>0.96</b> $\pm$ 0.00	<b>13.9</b> $\pm$ 0.1	28.1 $\pm$ 0.3	<b>17.3</b> $\pm$ 0.2	<b>21.6</b> $\pm$ 0.3	33.5 $\pm$ 0.6	18.5 $\pm$ 0.4

Table 3. Evaluation on VCTK and ZeroSpeech 2019. The MSE (dB<sup>2</sup>) of sample reconstruction is evaluated using the test set. We do not apply SQ-VAE (II) in this evaluation because of the variable length property of speech data and the different manipulations of speech signals between training and inference (see Appendix E.2).

Model	MSE (dB <sup>2</sup> )	
	VCTK	ZeroSpeech 2019
VQ-VAE w/ EMA	29.59 $\pm$ 0.25	34.33 $\pm$ 1.57
Gaussian SQ-VAE (I)	25.52 $\pm$ 0.08	33.17 $\pm$ 1.11
Gaussian SQ-VAE (III)	25.94 $\pm$ 0.22	34.35 $\pm$ 1.07
Gaussian SQ-VAE (IV)	<b>24.68</b> $\pm$ 0.21	<b>32.32</b> $\pm$ 0.88

of applying Gaussian SQ-VAE to the CelebA HQ 256 $\times$ 256 dataset (Karras et al., 2018) is shown in Appendix F.

**Speech** In this experiment, we test SQ-VAE and VQ-VAE by the reconstruction of the normalized log-Mel spectrogram using two speech datasets: VCTK version 0.80 (Veaux et al., 2017) and ZeroSpeech 2019 English (Dunbar et al., 2019). We adopt the VQ-VAE model of van Niekerk et al. (2020) as the baseline and replace its RNN-based vocoder with a projection layer. The codebook dimension  $d_b$  and the codebook size  $K$  are set to 64 and 512, respectively. For VQ-VAE, the first term of (4),  $-\log p_{\theta}(\mathbf{x}|\mathbf{Z}_q)$ , is set as  $\|\mathbf{x} - f_{\theta}(\mathbf{Z}_q)\|_2^2/2\sigma^2$ , following Eloff et al. (2019). The hyperparameter  $\sigma^2$  is determined by a grid search with  $\sigma^2 = \{10^{-2}, 10^{-1}, 10^0, 10^1\}$ . In this experiment, we run each experiment with five different random seeds and report the average and standard deviation of MSE values. As shown in Table 3, SQ-VAE achieves better average MSE than VQ-VAE even though VQ-VAE has been tuned with the hyperparameter  $\sigma^2$ .

Although SQ-VAE improves log-Mel spectrogram reconstruction, we observed no significant improvement when we applied SQ-VAE to a downstream task called acoustic unit discovery. Further exploration would be conducted to investigate the cause. The details about the experiment and the samples of reconstructed log-Mel spectrograms are described in Appendix E.2.

Table 4. Evaluation on CelebA-Mask. The pixel error (%), mIoU, and perplexity are evaluated using the test set.

Model	Pixel error	mIoU	Perplexity
VAE	8.79 $\pm$ 0.01	55.8 $\pm$ 0.3	–
VQ-VAE w/ EMA	6.95 $\pm$ 0.14	59.7 $\pm$ 0.7	46.2 $\pm$ 2.0
NC SQ-VAE	6.63 $\pm$ 1.38	64.1 $\pm$ 5.4	12.6 $\pm$ 5.2
vMF SQ-VAE	<b>3.51</b> $\pm$ 0.17	<b>74.6</b> $\pm$ 0.0	<b>52.4</b> $\pm$ 0.8

Table 5. Evaluation on MNIST and gray-scaled CelebA. The MSE ( $\times 10^3$ ) is evaluated using the test set.

Model	MNIST	Gray-CelebA
VAE	22.80 $\pm$ 0.32	17.73 $\pm$ 0.23
VQ-VAE w/ EMA	6.24 $\pm$ 0.18	5.19 $\pm$ 0.06
NC SQ-VAE	10.89 $\pm$ 0.47	3.88 $\pm$ 0.02
vMF SQ-VAE	<b>1.63</b> $\pm$ 0.21	<b>2.37</b> $\pm$ 0.01

## 5.2. Categorical Distributions

We test vMF SQ-VAE using CelebAHQ-Mask, which is a categorical image dataset with  $L = 19$  categories. In this experiment, the segmentation maps are rescaled to 64 $\times$ 64 with the nearest neighbor interpolation. We compare vMF SQ-VAE with VAE, VQ-VAE, and NC SQ-VAE (9a). For VAE and VQ-VAE, softmax is applied to the output of the decoders, and CE loss is adopted as the reconstruction objective. The reconstruction quality is measured using the pixel error and the mean of the class-wise intersection over union (mIoU) (Cordts et al., 2016).

The result is shown in Table 4. As expected, the codebook perplexity in NC SQ-VAE significantly deteriorates, whereas it still achieves better pixel error and mIoU than VQ-VAE. On the other hand, vMF SQ-VAE outperforms all the baselines with the highest codebook perplexity.

Next, we apply these methods to gray-scaled (256 categories) image datasets: MNIST and gray-scaled CelebA. The reconstruction accuracy is shown in Table 5. Again, vMF SQ-VAE outperforms the baselines, whereas NC SQ-VAE performs worse than VQ-VAE on MNIST. The result shows that vMF SQ-VAE is applicable even when the dataset has a large number of categories.



## 6. Conclusion

We proposed SQ-VAE, a VAE-based framework with discretized latent space and combined with a pair of stochastic dequantization and quantization processes that includes a trainable categorical posterior. This framework bridges the training schemes of VAE and VQ-VAE, and the trainable categorical posterior yields the self-annealing effect, which improves the reconstruction quality. Both Gaussian SQ-VAE and vMF SQ-VAE outperform the conventional approaches through experiments using vision and speech tasks. Moreover, SQ-VAE requires only a predefined scheduling of the temperature parameter of Gumbel–softmax. This hints a potential application of SQ-VAE on data compression.

Exploring the best parameterization of the trainable categorical posterior for different datasets and model architectures would be a direction of our future work.

## Acknowledgements

We would like to thank Masato Ishii, Lukas Mauch and Bac NguyenCong for many helpful comments during the preparation of this manuscript. Besides, we thank anonymous reviewers for their valuable suggestions and comments.

## References

- Agustsson, E., Mentzer, F., Tschannen, M., Cavigelli, L., Timofte, R., Benini, L., and Gool, L. V. Soft-to-hard vector quantization for end-to-end learning compressible representations. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- Bengio, Y., Léonard, N., and Courville, A. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- Chen, X., Mishra, N., Rohaninejad, M., and Abbeel, P. PixelSNAIL: An improved autoregressive generative model. In *Proc. International Conference on Machine Learning (ICML)*, pp. 864–872, 2018.
- Child, R., Gray, S., Radford, A., and Sutskever, I. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proc. 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734, 2014.
- Chorowski, J., Weiss, R. J., Bengio, S., and van den Oord, A. Unsupervised speech representation learning using WaveNet autoencoders. *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 27(12):2041–2053, 2019.
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., and Schiele, B. The cityscapes dataset for semantic urban scene understanding. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3213–3223, 2016.
- Davidson, T. R., Falorsi, L., De Cao, N., Kipf, T., and Tomczak, J. M. Hyperspherical variational auto-encoders. In *Proc. Conference on Uncertainty in Artificial Intelligence (UAI)*, 2018.
- Dhariwal, P., Jun, H., Payne, C., Kim, J. W., Radford, A., and Sutskever, I. Jukebox: A generative model for music. *arXiv preprint arXiv:2005.00341*, 2020.
- Dunbar, E., Algayres, R., Karadayi, J., Bernard, M., Benjumea, J., Cao, X.-N., Miskic, L., Dugrain, C., Ondel, L., Black, A. W., Besacier, L., Sakti, S., and Dupoux, E. The Zero Resource Speech Challenge 2019: TTS without T. In *Proc. Interspeech 2019*, pp. 1088–1092, 2019.
- Eloff, R., Nortje, A., van Niekerk, B., Govender, A., Nortje, L., Pretorius, A., van Biljon, E., van der Westhuizen, E., van Staden, L., and Kamper, H. Unsupervised acoustic unit discovery for speech synthesis using discrete latent-variable neural networks. In *Proc. Interspeech 2019*, pp. 1103–1107, 2019.
- Esser, P., Rombach, R., and Ommer, B. Taming transformers for high-resolution image synthesis. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 12873–12883, 2021.
- Ghosh, P., Sajjadi, M. S., Vergari, A., Black, M., and Schölkopf, B. From variational to deterministic autoencoders. In *Proc. International Conference on Learning Representation (ICLR)*, 2020.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pp. 6626–6637, 2017.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-VAE: Learning basic visual concepts with a constrained variational framework. In *Proc. International Conference on Learning Representation (ICLR)*, 2017.
- Jang, E., Gu, S., and Poole, B. Categorical reparameterization with gumbel-softmax. In *Proc. International Conference on Learning Representation (ICLR)*, 2017.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. An introduction to variational methods for graphical models. *Machine Learning*, 37(2):183–233, 1999.
- Kaiser, L., Bengio, S., Roy, A., Vaswani, A., Parmar, N., Uszkoreit, J., and Shazeer, N. Fast decoding in

- sequence models using discrete latent variables. In *Proc. International Conference on Machine Learning (ICML)*, pp. 2390–2399, 2018.
- Karras, T., Aila, T., Laine, S., and Lehtinen, J. Progressive growing of gans for improved quality, stability, and variation. In *Proc. International Conference on Learning Representation (ICLR)*, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *Proc. International Conference on Learning Representation (ICLR)*, 2015.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. In *Proc. International Conference on Learning Representation (ICLR)*, 2014.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Kumar, K., Kumar, R., de Boissiere, T., Gestin, L., Teoh, W. Z., Sotelo, J., de Brébisson, A., Bengio, Y., and Courville, A. C. MelGAN: Generative adversarial networks for conditional waveform synthesis. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE*, 86(11):2278–2324, 1998.
- Lee, C.-H., Liu, Z., Wu, L., and Luo, P. MaskGAN: Towards diverse and interactive facial image manipulation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5549–5558, 2020.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3730–3738, 2015.
- Maddison, C. J., Mnih, A., and Teh, Y. W. The concrete distribution: A continuous relaxation of discrete random variables. In *Proc. International Conference on Learning Representation (ICLR)*, 2017.
- Polyak, B. T. and Juditsky, A. B. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- Polykovskiy, D. and Vetrov, D. Deterministic decoding for discrete data in variational autoencoders. In *Proc. International Conference on Artificial Intelligence and Statistics (AISTAS)*, pp. 3046–3056, 2020.
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., and Sutskever, I. Zero-shot text-to-image generation. In *Proc. International Conference on Machine Learning (ICML)*, pp. 8821–8831, 2021.
- Razavi, A., van den Oord, A., and Vinyals, O. Generating diverse high-fidelity images with VQ-VAE-2. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pp. 14866–14876, 2019.
- Roy, A., Vaswani, A., Neelakantan, A., and Parmar, N. Theory and experiments on vector quantized autoencoders. *arXiv preprint arXiv:1805.11063*, 2018.
- Schatz, T., Peddinti, V., Bach, F., Jansen, A., Hermansky, H., and Dupoux, E. Evaluating speech features with the minimal-pair ABX task: analysis of the classical MFC/PLP pipeline. In *Proc. Interspeech 2013*, pp. 1781–1785, 2013.
- Takida, Y., Liao, W.-H., Uesaka, T., Takahashi, S., and Mitsufuji, Y. Preventing posterior collapse induced by oversmoothing in Gaussian VAE. *arXiv preprint arXiv:2102.08663*, 2021.
- Theis, L., Shi, W., Cunningham, A., and Huszár, F. Lossy image compression with compressive autoencoders. In *Proc. International Conference on Learning Representation (ICLR)*, 2017.
- Tjandra, A., Sakti, S., and Nakamura, S. Transformer VQ-VAE for unsupervised unit discovery and speech synthesis: ZeroSpeech 2020 challenge. In *Proc. Interspeech 2020*, pp. 4851–4855, 2020.
- Toderici, G., O’Malley, S. M., Hwang, S. J., Vincent, D., Minnen, D., Baluja, S., Covell, M., and Sukthankar, R. Variable rate image compression with recurrent neural networks. In *Proc. International Conference on Learning Representation (ICLR)*, 2016.
- Tolstikhin, I., Bousquet, O., Gelly, S., and Schoelkopf, B. Wasserstein auto-encoders. In *Proc. International Conference on Learning Representation (ICLR)*, pp. 5885–5892, 2018.
- van den Oord, A., Kalchbrenner, N., and Kavukcuoglu, K. Pixel recurrent neural networks. In *Proc. International Conference on Machine Learning (ICML)*, pp. 1747–1756, 2016.
- van den Oord, A., Vinyals, O., and Kavukcuoglu, K. Neural discrete representation learning. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pp. 6306–6315, 2017.
- van Niekerk, B., Nortje, L., and Kamper, H. Vector-quantized neural networks for acoustic unit discovery in the ZeroSpeech 2020 challenge. In *Proc. Interspeech 2020*, pp. 4836–4840, 2020.
- Veaux, C., Yamagishi, J., and MacDonald, K. CSTR VCTK corpus: English multi-speaker corpus for CSTR Voice Cloning Toolkit. <https://datashare.ed.ac.uk/handle/10283/2651>, 2017.

- Williams, W., Ringer, S., Ash, T., MacLeod, D., Dougherty, J., and Hughes, J. Hierarchical quantized autoencoder. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pp. 4524–4535, 2020.
- Wu, H. and Flierl, M. Vector quantization-based regularization for autoencoders. In *Proc. AAAI Conference on Artificial Intelligence (AAAI)*, pp. 6380–6387, 2020.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- Yamamoto, R., Song, E., and Kim, J.-M. Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram. In *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, pp. 6199–6203, 2020.
- Zhao, S., Song, J., and Ermon, S. InfoVAE: Balancing learning and inference in variational autoencoders. In *Proc. AAAI Conference on Artificial Intelligence (AAAI)*, pp. 5885–5892, 2019.

---

## Supplementary Material

---

### A. Notations and Definitions

Here, we introduce several notations and definitions we used in the article. Notations used in formulation of SQ-VAE is listed in Table 6.

**Softmax function** We use the following notation to represent the  $k$ -th component of a softmax function that operates among the elements indexed by  $k$ . It can also be regarded as a function that maps  $\mathbb{R}^K$  to  $(0, 1)^K$ .

$$\text{softmax}_k(\{z_j\}_{j=1}^K) = \frac{\exp(z_k)}{\sum_{j=1}^K \exp(z_j)}. \quad (15)$$

**von Mises-Fisher distribution** The von Mises-Fischer distribution is parameterized with mean direction and concentration  $(\mathbf{m}, \kappa)$ , which is represented as

$$\text{vMF}(\mathbf{m}, \kappa) = C_F(\kappa) \exp(\kappa \mathbf{m}^\top \mathbf{v}), \quad (16)$$

where  $C_F(\kappa)$  is the normalizing constant obtained by using

$$C_F(\kappa) = \int_{\mathbf{v} \in \mathcal{S}^{F-1}} \exp(\kappa \mathbf{m}^\top \mathbf{v}) d\mathbf{v}. \quad (17)$$

The normalizing constant can also be expressed with the modified Bessel function of the first kind  $I_\nu(\cdot)$ <sup>6</sup> as

$$\log C_F(\kappa) = \left(\frac{F}{2} - 1\right) \log \kappa - \log I_{F/2-1}(\kappa) - \frac{F}{2} \log(2\pi). \quad (18)$$

### B. Derivation Details

The ELBO of SQ-VAE can be calculated using the Bayes' theorem  $P_\theta(\mathbf{Z}_q|\mathbf{x})p_\theta(\mathbf{x}) = p_\theta(\mathbf{x}|\mathbf{Z}_q)P(\mathbf{Z}_q)$  as

$$\begin{aligned} \log p_\theta(\mathbf{x}) &\geq \log p_\theta(\mathbf{x}) - D_{\text{KL}}(q_\omega(\mathbf{Z}|\mathbf{x})\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z}) \parallel P_\theta(\mathbf{Z}_q|\mathbf{x})p_\varphi(\mathbf{Z}|\mathbf{Z}_q)) \\ &= \mathbb{E}_{q_\omega(\mathbf{Z}|\mathbf{x})\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})} \left[ \log \frac{p_\theta(\mathbf{x})P_\theta(\mathbf{Z}_q|\mathbf{x})p_\varphi(\mathbf{Z}|\mathbf{Z}_q)}{q_\omega(\mathbf{Z}|\mathbf{x})\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})} \right] \\ &= \mathbb{E}_{q_\omega(\mathbf{Z}|\mathbf{x})\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})} \left[ \log \frac{p_\theta(\mathbf{x}|\mathbf{Z}_q)p_\varphi(\mathbf{Z}|\mathbf{Z}_q)}{q_\omega(\mathbf{Z}|\mathbf{x})} - \log \frac{\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})}{P(\mathbf{Z}_q)} \right] \\ &= \mathbb{E}_{q_\omega(\mathbf{Z}|\mathbf{x})\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})} \left[ \log \frac{p_\theta(\mathbf{x}|\mathbf{Z}_q)p_\varphi(\mathbf{Z}|\mathbf{Z}_q)}{q_\omega(\mathbf{Z}|\mathbf{x})} \right] + \mathbb{E}_{q_\omega(\mathbf{Z}|\mathbf{x})} H(\hat{P}_\varphi(\mathbf{Z}_q|\mathbf{Z})) - d_z \log K. \end{aligned} \quad (19)$$

---

<sup>6</sup>In our experiments, we use the following implementation provided by Davidson et al. (2018) to evaluate the modified Bessel function: [https://github.com/nicola-decao/s-vae-pytorch/tree/master/hyperspherical\\_vae](https://github.com/nicola-decao/s-vae-pytorch/tree/master/hyperspherical_vae)

### B.1. Gaussian SQ-VAE

By applying the Gaussian assumption as mentioned in Section 3.2 into (19), we obtain the first two terms as

$$\begin{aligned} \log p_{\theta}(\mathbf{x}|\mathbf{Z}_q) &= \log \mathcal{N}(f_{\theta}(\mathbf{Z}_q), \sigma^2 \mathbf{I}) \\ &= -\frac{1}{2\sigma^2} \|\mathbf{x} - f_{\theta}(\mathbf{Z}_q)\|_2^2 + \frac{D}{2} \log(2\pi\sigma^2), \end{aligned} \quad (20)$$

$$\begin{aligned} \mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{x})\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})} \left[ \log \frac{p_{\varphi}(\mathbf{Z}|\mathbf{Z}_q)}{q_{\omega}(\mathbf{Z}|\mathbf{x})} \right] &= \mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{x})\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})} \left[ \log \frac{\prod_{i=1}^{d_z} p_{\varphi}(\mathbf{z}_i|\mathbf{Z}_q)}{\prod_{i=1}^{d_z} q_{\omega}(\mathbf{z}_i|\mathbf{x})} \right] \\ &= \sum_{i=1}^{d_z} \mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{x})\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})} \left[ \log \frac{p_{\varphi}(\mathbf{z}_i|\mathbf{Z}_q)}{p_{\varphi}(\mathbf{z}_i|g_{\phi}(\mathbf{x}))} \right] \end{aligned} \quad (21)$$

$$\begin{aligned} &= \frac{1}{2} \sum_{i=1}^{d_z} \left( -\mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{x})\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})} [(\mathbf{z}_i - \mathbf{z}_{q,i})^{\top} \Sigma_{\varphi}^{-1}(\mathbf{z}_i - \mathbf{z}_{q,i})] \right. \\ &\quad \left. + \mathbb{E}_{p_{\varphi}(\mathbf{Z}|g_{\phi}(\mathbf{x}))} [(\mathbf{z}_i - g_{\phi,i}(\mathbf{x}))^{\top} \Sigma_{\varphi}^{-1}(\mathbf{z}_i - g_{\phi,i}(\mathbf{x}))] \right) \\ &= -\mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{x})\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})} \sum_{i=1}^{d_z} \left[ \frac{1}{2} (\mathbf{z}_i - \mathbf{z}_{q,i})^{\top} \Sigma_{\varphi}^{-1}(\mathbf{z}_i - \mathbf{z}_{q,i}) \right] + \frac{d_b d_z}{2}. \end{aligned} \quad (22)$$

This leads to the objective function of Gaussian SQ-VAE  $\mathcal{L}_{\mathcal{N}\text{-SQ}}$  as (8).

### B.2. vMF SQ-VAE

The ELBO for vMF SQ-VAE is obtained by replacing  $\mathbf{x}$  with  $\mathbf{V}$  as

$$\log p_{\theta}(\mathbf{V}) \geq \mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{V})\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})} \left[ \log \frac{p_{\theta}(\mathbf{V}|\mathbf{Z}_q)p_{\varphi}(\mathbf{Z}|\mathbf{Z}_q)}{q_{\omega}(\mathbf{Z}|\mathbf{V})} \right] - \mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{V})} H(\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})) - d_z \log K. \quad (23)$$

By applying the von Mises-Fischer distribution as mentioned in 3.4 into (19), we obtain the first two terms as

$$\begin{aligned} \log p_{\theta}(\mathbf{x}|\mathbf{Z}_q) &= \sum_{d=1}^D \log \text{vMF}(\tilde{f}_{\theta,d}(\mathbf{Z}_q), \kappa) \\ &= \kappa \sum_{d=1}^D \mathbf{v}_d^{\top} \tilde{f}_{\theta,d}(\mathbf{Z}_q) + \log C_F(\kappa), \end{aligned} \quad (24)$$

$$\begin{aligned} \mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{x})\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})} \left[ \log \frac{p_{\varphi}(\mathbf{Z}|\mathbf{Z}_q)}{q_{\omega}(\mathbf{Z}|\mathbf{x})} \right] &= \mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{x})\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})} \left[ \log \frac{\prod_{i=1}^{d_z} p_{\varphi}(\mathbf{z}_i|\mathbf{Z}_q)}{\prod_{i=1}^{d_z} q_{\omega}(\mathbf{z}_i|\mathbf{x})} \right] \\ &= \sum_{i=1}^{d_z} \mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{x})\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})} \left[ \log \frac{p_{\varphi}(\mathbf{z}_i|\mathbf{Z}_q)}{p_{\varphi}(\mathbf{z}_i|g_{\phi,i}(\mathbf{x}))} \right] \\ &= \sum_{i=1}^{d_z} \left( \mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{x})\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})} [\kappa_{\varphi} \mathbf{z}_{q,i}^{\top} \mathbf{z}_i] - \mathbb{E}_{p_{\varphi}(\mathbf{Z}|g_{\phi}(\mathbf{x}))} [\kappa_{\varphi} g_{\phi,i}(\mathbf{x})^{\top} \mathbf{z}_i] \right) \\ &= \sum_{i=1}^{d_z} \left( \mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{x})\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})} [\kappa_{\varphi} \mathbf{z}_{q,i}^{\top} \mathbf{z}_i] - \kappa_{\varphi} \|g_{\phi,i}(\mathbf{x})\|_2^2 \right) \\ &= \mathbb{E}_{q_{\omega}(\mathbf{Z}|\mathbf{x})\hat{P}_{\varphi}(\mathbf{Z}_q|\mathbf{Z})} \left[ \sum_{i=1}^{d_z} \kappa_{\varphi} (\mathbf{z}_{q,i}^{\top} \mathbf{z}_i - 1) \right]. \end{aligned} \quad (25)$$

This leads to the objective function of vMF SQ-VAE  $\mathcal{L}_{\text{vMF-SQ}}$  as (14).

Table 6. Notations used in our formulation of SQ-VAE.

Trainable parameters	
$\mathbf{B}$	A trainable codebook that consists of $K$ codebook elements, $\{\mathbf{b}_k\}_{k=1}^K$ .
$\theta$	A trainable parameter for the decoder. The decoding function is denoted as $f_\theta : \mathbb{R}^{d_b \times d_z} \rightarrow \mathbb{R}^D$ .
$\phi$	A trainable parameter for the encoder. The encoding function is denoted as $g_\phi : \mathbb{R}^D \rightarrow \mathbb{R}^{d_b \times d_z}$ .
$\varphi$	A trainable parameter for probabilistic dequantization and quantization processes in the latent space.
$\omega$	A tuple of $\phi$ and $\varphi$ , which is used in the probabilistic encoding process.
Latent variables	
$\mathbf{Z}_q$	A latent variable that consists of $d_z$ codebook elements, i.e., $\mathbf{Z}_q \in \mathbf{B}^{d_z}$ , whose $i$ th vector is denoted as $\mathbf{z}_i \in \mathbf{B}$ .
$\hat{\mathbf{Z}}_q$	A latent variable encoded directly from the encoder, i.e., $\hat{\mathbf{Z}}_q = g_\phi(\mathbf{x})$ , which approximates $\mathbf{Z}_q$ .
$\mathbf{Z}$	A continuous latent variable obtained by dequantization of $\mathbf{Z}_q$ or $\hat{\mathbf{Z}}_q$ .
Probabilistic processes	
$p_\theta(\mathbf{x} \mathbf{Z}_q)$	A decoder distribution whose mean is given by $f_\theta(\mathbf{Z}_q)$ (e.g., Gaussian and vMF distribution).
$p_\varphi(\mathbf{Z} \mathbf{Z}_q)$	A probability distribution for the dequantization of $\mathbf{Z}_q$ (e.g., Gaussian and vMF distribution).
$q_\omega(\mathbf{Z} \mathbf{x})$	A distribution for the encoding from $\mathbf{x}$ to $\mathbf{Z}$ , which consists of deterministic encoding and stochastic dequantization.
$\hat{P}_\varphi(\mathbf{Z}_q \mathbf{Z})$	A categorical distribution for the quantization of $\mathbf{Z}$ , which is inverse process of $p_\varphi(\mathbf{Z} \mathbf{Z}_q)$ .

## C. Training Procedures of SQ-VAEs

The training procedures of Gaussian SQ-VAE and vMF SQ-VAE are described here in Algorithms 1 and 2, respectively.  $[t]$  indicates the index of training steps.

### Algorithm 1 Gaussian SQ-VAE

**Input:** Dataset  $\mathbf{x}_{data}$   
 Initialize the codebook and parameters:  
 $\mathbf{B}^{[0]}, \theta^{[0]}$  and  $\omega^{[0]} := \{\phi^{[0]}, \varphi^{[0]}\}$

**for**  $t = 1, 2, \dots, T$  **do**  
 $\mathbf{x} \leftarrow$  Random minibatch from  $\mathbf{x}_{data}$

$\hat{\mathbf{Z}}_q \leftarrow g_{\phi^{[t-1]}}(\mathbf{x})$   
 $\mathbf{Z}_q \sim \hat{P}_{\varphi^{[t-1]}}(\mathbf{Z}_q|\hat{\mathbf{Z}}_q)$   
 $\mathbf{g} \leftarrow \nabla_{\theta, \omega, \mathbf{B}} \mathcal{L}_{\mathcal{N}\text{-SQ}}(\theta^{[t-1]}, \omega^{[t-1]}, \mathbf{B}^{[t-1]})$   
 with sampled  $\mathbf{x}$  and  $\mathbf{Z}_q$   
 $\theta^{[t]}, \omega^{[t]}, \mathbf{B}^{[t]} \leftarrow$  Update parameters using  $\mathbf{g}$

**end for**

### Algorithm 2 vMF SQ-VAE

**Input:** Dataset  $\mathbf{x}_{data}$   
 Initialize the codebook and parameters:  
 $\mathbf{B}^{[0]}, \theta^{[0]}$  and  $\omega^{[0]} := \{\phi^{[0]}, \varphi^{[0]}\}$   
 Determine a vector set  $\mathbf{W} := \{\mathbf{w}_l | \mathbf{w}_l \in \mathcal{S}^{F-1}\}_{l=1}^L$

**for**  $t = 1, 2, \dots, T$  **do**  
 $\mathbf{x} \leftarrow$  Random minibatch from  $\mathbf{x}_{data}$   
 $\mathbf{V} \leftarrow$  Project  $\mathbf{x}$  onto  $\mathcal{S}^{F-1}$   
 $\hat{\mathbf{Z}}_q \leftarrow g_{\phi^{[t-1]}}(\mathbf{V})$   
 $\mathbf{Z}_q \sim \hat{P}_{\varphi^{[t-1]}}(\mathbf{Z}_q|\hat{\mathbf{Z}}_q)$   
 $\mathbf{g} \leftarrow \nabla_{\theta, \omega, \mathbf{B}} \mathcal{L}_{\text{vMF-SQ}}(\theta^{[t-1]}, \omega^{[t-1]}, \mathbf{B}^{[t-1]})$   
 with sampled  $\mathbf{x}$  and  $\mathbf{Z}_q$   
 $\theta^{[t]}, \omega^{[t]}, \mathbf{B}^{[t]} \leftarrow$  Update parameters using  $\mathbf{g}$

**end for**

## D. Self-Annealed Quantization

### D.1. Similarity between SQ-VAE and conventional VAE

A property that is similar to Proposition 1 can be observed in Gaussian VAE with the posterior  $q_\phi(\mathbf{z}|\mathbf{x}) = \mathcal{N}(g_\phi(\mathbf{x}), s^2\mathbf{I})$ . In this case, the latent variables are perturbed by adding Gaussian noises with their variance  $s^2\mathbf{I}$  to the encoded points. When  $s^2$  is trained, it approaches zero (Takida et al., 2021) as the training progresses, which means that the stochastic encoding becomes almost deterministic, i.e., no perturbation.

### D.2. Proof of Proposition 1

**Proposition 1.** Assume that  $p_{data}(\mathbf{x})$  has finite support, whereas  $g_\phi$  and  $\{\mathbf{b}_k\}_{k=1}^K$  are bounded. Let  $\omega^* = \{\phi^*, \varphi^*\}$  be a minimizer of  $\mathbb{E}_{p_{data}(\mathbf{x})} D_{\text{KL}}(Q_\omega(\mathbf{Z}_q|\mathbf{x}) || P_\theta(\mathbf{Z}_q|\mathbf{x}))$  with fixed  $\theta, \sigma^2$ , and  $\{\mathbf{b}_k\}_{k=1}^K$ . If  $\sigma^2 \rightarrow 0$ , then  $\sigma_{\varphi^*}^2 \rightarrow 0$ .

*Proof.* We denote the number of samples included in the support of  $p_{data}(\mathbf{x})$  as  $N_{data}$ . Let  $\mathbf{Z}_q^{\mathbf{k}} \in \mathbf{B}^{d_z}$  with a set  $\mathbf{k} \in [K]^{d_z}$

be a discrete tensor with its  $i$ th elements  $\mathbf{b}_{k_i}$ , where  $k_i$  indicates the  $i$ th coordinate of  $\mathbf{k}$  for  $i \in [d_z]$ . Among  $f_\theta(\mathbf{Z}_q^{\mathbf{k}})$  with all the possible  $\mathbf{k} \in [K]^{d_z}$ , the set of indices corresponding to the nearest neighbor of  $\mathbf{x}^{(n)}$  is denoted as  $\hat{\mathbf{k}}^{(n)}$  for  $n \in [N_{\text{data}}]$ , i.e.,  $\hat{\mathbf{k}}^{(n)} = \arg \min_{\mathbf{k} \in [K]^{d_z}} \|\mathbf{x}^{(n)} - f_\theta(\mathbf{Z}_q^{\mathbf{k}})\|_2$ . Now,  $P_\theta(\mathbf{Z}_q | \mathbf{x}^{(n)})$  can be evaluated by

$$\begin{aligned} P_\theta(\mathbf{Z}_q | \mathbf{x}^{(n)}) &\propto P_\theta(\mathbf{x}^{(n)} | \mathbf{Z}_q) P(\mathbf{Z}_q) \\ &\propto \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{x}^{(n)} - f_\theta(\mathbf{Z}_q)\|_2^2\right). \end{aligned} \quad (26)$$

Here, if  $\sigma^2 \rightarrow 0$ , then  $P_\theta(\mathbf{Z}_q | \mathbf{x}^{(n)})$  becomes the Kronecker delta function  $\delta_{\mathbf{k}, \hat{\mathbf{k}}^{(n)}}$ .

Next, define  $\mathcal{B}_k \subset \mathbb{R}^{d_b}$  to be a region such that  $k = \arg \min_{k'} \|\mathbf{z} - \mathbf{b}_{k'}\|_2$  for  $\mathbf{z} \in \mathcal{B}_k$ . With the arbitrarily complex  $g_\phi$ ,  $g_{\phi,i}(\mathbf{x}^{(n)}) \in \mathcal{B}_k$  can be achieved for all  $n \in [N_{\text{data}}]$ ,  $i \in [d_z]$  and  $k \in [K]$ , where  $g_{\phi,i}(\mathbf{x}^{(n)})$  indicates the  $i$ th coordinates of  $g_\phi(\mathbf{x}^{(n)})$ . The divergence  $\mathbb{E}_{p_{\text{data}}(\mathbf{x})} D_{\text{KL}}(Q_\omega(\mathbf{Z}_q | \mathbf{x}) \| P_\theta(\mathbf{Z}_q | \mathbf{x}))$  can be minimized to 0 with

$$\sigma_{\varphi^*}^2 \rightarrow 0 \quad \text{and} \quad (27a)$$

$$g_{\phi^*,i}(\mathbf{x}^{(n)}) \in \mathcal{B}_{\hat{k}_i^{(n)}} \quad \text{for all } i \in [d_z] \text{ and } n \in [N_{\text{data}}] \quad (27b)$$

since (27a) and (27b) lead to  $Q_\omega(\mathbf{Z}_q | \mathbf{x}) = \delta_{\mathbf{k}, \hat{\mathbf{k}}^{(n)}}$ . Here,  $\hat{k}_i^{(n)}$  denotes the  $i$ th coordinates of  $\hat{\mathbf{k}}^{(n)}$ .

Finally, we prove that if  $\sigma^2 \rightarrow 0$  and  $\mathbb{E}_{p_{\text{data}}(\mathbf{x})} D_{\text{KL}}(Q_{\omega^*}(\mathbf{Z}_q | \mathbf{x}) \| P_\theta(\mathbf{Z}_q | \mathbf{x})) = 0$ , then  $\sigma_{\varphi^*}^2 \rightarrow 0$ . Define  $p_1^{(n)}$  and  $q_1^{(n)}$  to be  $p_1 = P_\theta(\mathbf{Z}_q^{\hat{\mathbf{k}}^{(n)}} | \mathbf{x}^{(n)})$  and  $q_1 = Q_\omega(\mathbf{Z}_q^{\hat{\mathbf{k}}^{(n)}} | \mathbf{x}^{(n)})$ , respectively. Now, consider  $\sigma_{\varphi^*}^2 \not\rightarrow 0$  as  $\sigma^2 \rightarrow 0$ . It immediately follows that  $p_1 \rightarrow 1$  as  $\sigma^2 \rightarrow 0$ . Moreover, according to the assumption that  $p_{\text{data}}(\mathbf{x})$  has finite support and  $g_\phi$  is bounded, the distribution that  $\mathbf{Z} = g_\phi(\mathbf{x})$  follows with  $\mathbf{x} \sim p_{\text{data}}(\mathbf{x})$  has finite support as well, which leads to  $q_1 \neq 1$ . Here, the KL divergence for  $\mathbf{x}^{(n)}$  is bounded as

$$\begin{aligned} D_{\text{KL}}(Q_{\omega^*}(\mathbf{Z}_q | \mathbf{x}^{(n)}) \| P_\theta(\mathbf{Z}_q | \mathbf{x}^{(n)})) &= \sum_{\mathbf{k} \in [K]^{d_z}} Q_{\omega^*}(\mathbf{Z}_q^{\mathbf{k}} | \mathbf{x}^{(n)}) \left[ \log Q_{\omega^*}(\mathbf{Z}_q^{\mathbf{k}} | \mathbf{x}^{(n)}) - \log P_\theta(\mathbf{Z}_q^{\mathbf{k}} | \mathbf{x}^{(n)}) \right] \\ &= -H[Q_{\omega^*}(\mathbf{Z}_q | \mathbf{x}^{(n)})] - \sum_{\mathbf{k} \in [K]^{d_z}} Q_{\omega^*}(\mathbf{Z}_q^{\mathbf{k}} | \mathbf{x}^{(n)}) \log P_\theta(\mathbf{Z}_q^{\mathbf{k}} | \mathbf{x}^{(n)}) \\ &\geq -\log K - (1 - q_1) \log(1 - p_1) \end{aligned} \quad (28)$$

$$\geq -\log K - (1 - q_1) \log(1 - p_1) \quad (29)$$

from

$$H[Q_{\omega^*}(\mathbf{Z}_q | \mathbf{x}^{(n)})] \leq \log K \quad \text{and} \quad (30)$$

$$\sum_{\mathbf{k} \in [K]^{d_z}} Q_{\omega^*}(\mathbf{Z}_q^{\mathbf{k}} | \mathbf{x}^{(n)}) \log P_\theta(\mathbf{Z}_q^{\mathbf{k}} | \mathbf{x}^{(n)}) = Q_{\omega^*}(\mathbf{Z}_q^{\hat{\mathbf{k}}^{(n)}} | \mathbf{x}^{(n)}) \log P_\theta(\mathbf{Z}_q^{\hat{\mathbf{k}}^{(n)}} | \mathbf{x}^{(n)}) + \sum_{\substack{\mathbf{k} \in [K]^{d_z} \\ \mathbf{k} \neq \hat{\mathbf{k}}^{(n)}}} Q_{\omega^*}(\mathbf{Z}_q^{\mathbf{k}} | \mathbf{x}^{(n)}) \log P_\theta(\mathbf{Z}_q^{\mathbf{k}} | \mathbf{x}^{(n)}) \quad (31)$$

$$\leq q_1 \log p_1 + (1 - q_1) \log(1 - p_1) \quad (32)$$

$$\leq (1 - q_1) \log(1 - p_1) \quad (33)$$

The first term in (29) is finite since  $K$  is assumed to be finite. On the other hand, from  $1 - p_1 \rightarrow 0$  ( $\sigma^2 \rightarrow 0$ ) and  $q_1 \neq 1$ , it follows that the second term in (29) diverges to infinity as  $\sigma^2 \rightarrow 0$ . This leads to the infinite KL divergence and contradicts to the fact that  $\sigma_{\varphi^*}$  is a minimizer of  $\mathbb{E}_{p_{\text{data}}(\mathbf{x})} D_{\text{KL}}(Q_\omega(\mathbf{Z}_q | \mathbf{x}) \| P_\theta(\mathbf{Z}_q | \mathbf{x}))$ . Thus, we must have  $\sigma_{\varphi^*} \rightarrow 0$  as  $\sigma^2 \rightarrow 0$ .  $\square$

### D.3. Details of Section 3.3

We set the codebook capacity to  $(d_b, K) = (64, 128)$  for all the models. We implement the decoder and encoder using the following standard two-layer ConvNet architectures:

$$\begin{aligned}
 \mathbf{x} \in \mathbb{R}^{28 \times 28 \times 1} &\rightarrow \text{Conv}_{32}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (32, 14, 14)} \\
 &\rightarrow \text{Conv}_{64}^{(4 \times 4)} && \text{size: (64, 7, 7)} \\
 \mathbf{Z}_q \in \mathbf{B}^{7 \times 7} \subset \mathbb{R}^{64 \times 7 \times 7} &\rightarrow \text{ConvT}_{32}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (32, 14, 14)} \\
 &\rightarrow \text{ConvT}_1^{(4 \times 4)} \rightarrow \text{Sigmoid} && \text{size: (1, 28, 28)},
 \end{aligned}$$

where the notations for the architecture parts are listed in Table 7.

The Adam optimizer is used with the initial learning rate of 0.001. The temperature parameter in Gumbel–softmax is annealed with the same preset schedule as that in Appendix E.

### D.4. Case of vMF SQ-VAE

As for vMF SQ-VAE, we have the following proposition:

**Proposition 2.** *Assuming that  $p_{\text{data}}(\mathbf{x})$  has finite support, whereas  $g_\phi$  and  $\{\mathbf{b}_k\}_{k=1}^K$  are bounded. Let  $\omega^* = \{\phi^*, \varphi^*\}$  be a minimizer of  $\mathbb{E}_{p_{\text{data}}(\mathbf{x})} D_{\text{KL}}(Q_\omega(\mathbf{Z}_q|\mathbf{x}) \| P_\theta(\mathbf{Z}_q|\mathbf{x}))$  with fixed  $\theta, \kappa$ , and  $\{\mathbf{b}_k\}_{k=1}^K$ . If  $\kappa \rightarrow \infty$ , then  $\kappa_{\phi^*} \rightarrow \infty$ .*

This proposition is similar to Proposition 1. Thus, it can be proved similarly to the proof of Proposition 1 while noting the following:

$$\begin{aligned}
 P_\theta(\mathbf{Z}_q|\mathbf{x}^{(n)}) &\propto P_\theta(\mathbf{V}^{(n)}|\mathbf{Z}_q)P(\mathbf{Z}_q) \\
 &\propto \exp\left(\kappa \sum_{d=1}^D \mathbf{v}_d^{(n)\top} \tilde{f}_{\theta,d}(\mathbf{Z}_q)\right),
 \end{aligned} \tag{34}$$

where  $-1 \leq \mathbf{v}_d^{(n)\top} \tilde{f}_{\theta,d}(\mathbf{Z}_q) \leq 1$  since the vectors live on  $\mathcal{S}^{F-1}$ . Here, if the concentration  $\kappa$  approaches to  $\infty$ , then  $P_\theta(\mathbf{Z}_q|\mathbf{x}^{(n)})$  becomes the Kronecker delta function  $\delta_{\mathbf{k}, \hat{\mathbf{k}}^{(n)}}$ .

## E. Experimental Details

Throughout all the experiments, we apply the same annealing schedule as that used by Jang et al. (2017) for the temperature parameter of Gumbel–softmax, which is  $\tau = \exp(10^{-5} \cdot t)$ , where  $t$  denotes the global training step. We set the VQ-VAE hyperparameter  $\beta$  in (4) and weight decay  $\gamma$  in EMA to 0.25 and 0.99, respectively, as suggested in van den Oord et al. (2017). For SQ-VAEs, we use  $\mathbf{Z} = \mathbf{Z}_q$  instead of sampling  $\mathbf{Z} \sim p_\varphi(\mathbf{Z}|\hat{\mathbf{Z}}_q)$  in the Monte Carlo estimate of the expectations in (5) as in Algorithms 1 and 2, which stabilizes the estimation of Monte Carlo.

### E.1. Gaussian SQ-VAE on Image Datasets

#### E.1.1. DATASETS AND PREPROCESSING

**MNIST and Fashion-MNIST** They contain  $28 \times 28$  grayscale images, which are categorized into 10 classes. We use the default train/test split (60,000/10,000 samples) and further split 10,000 samples from the training set as the validation set.

**CIFAR10** CIFAR10 contains 10 classes of  $32 \times 32$  RGB images. We use the default train/test split (50,000/10,000 samples) and further split 10,000 samples from the training set as the validation set.

**CelebA** CelebA consists of 202,599 colored face images. We use the default train/valid/test split. We preprocess the images with a center cropping of  $140 \times 140$  and resize them to  $64 \times 64$  using the bilinear interpolation following these previous studies (Tolstikhin et al., 2018; Ghosh et al., 2020).



Table 7. Notations of network layers used on image datasets.

Notation	Description
$\text{Conv}_n^{(4 \times 4)}$	2D Convolutional layer (channel= $n$ , kernel= $4 \times 4$ , stride= 2)
$\text{Conv}_n^{(3 \times 3)}$	2D Convolutional layer (channel= $n$ , kernel= $3 \times 3$ , stride= 1)
$\text{Conv}_n^{(1 \times 1)}$	2D Convolutional layer (channel= $n$ , kernel= $1 \times 1$ , stride= 1)
$\text{ConvT}_n^{(4 \times 4)}$	2D Transpose convolutional layer (channel= $n$ , kernel= $4 \times 4$ , stride= 2)
$\text{ConvT}_n^{(3 \times 3)}$	2D Transpose convolutional layer (channel= $n$ , kernel= $3 \times 3$ , stride= 1)
$\text{ResBlock}_n$	Resblock (ReLU $\rightarrow$ Conv $_n^{(3 \times 3)}$ $\rightarrow$ BatchBorm $\rightarrow$ ReLU $\rightarrow$ Conv $_n^{(1 \times 1)}$ $\rightarrow$ BatchBorm + identity mapping)

### E.1.2. MODEL DESCRIPTION AND TRAINING

We adopt the ConvResNets from the GitHub repository of DeepMind `sonnet/examples/vqvae_example.ipynb`<sup>7</sup>. These networks include convolutional layers, transpose convolutional layers, and ResBlocks. Their notations are listed in Table 7. We use the following network architectures for the encoders and decoders on MNIST and Fashion-MNIST:

$$\begin{aligned}
 \mathbf{x} \in \mathbb{R}^{28 \times 28} &\rightarrow \text{Conv}_{d_b/2}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: } (d_b/2, 14, 14) \\
 &\rightarrow \text{Conv}_{d_b}^{(4 \times 4)} && \text{size of } (d_b, 7, 7) \\
 &\rightarrow [\text{ResBlock}_{d_b}]_{\times N_{\text{resblock}}} && \text{size of } (d_b, 7, 7) \\
 \mathbf{Z}_q \in \mathbf{B}^{7 \times 7} \subset \mathbb{R}^{d_b \times 7 \times 7} &\rightarrow [\text{ResBlock}_{d_b}]_{\times N_{\text{resblock}}} && \text{size of } (d_b, 7, 7) \\
 &\rightarrow \text{ConvT}_{d_b/2}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: } (d_b/2, 14, 14) \\
 &\rightarrow \text{ConvT}_1^{(4 \times 4)} \rightarrow \text{Sigmoid} && \text{size: } (1, 28, 28),
 \end{aligned}$$

where  $N_{\text{resblock}}$  is set to 2 and 6 for MNIST and Fashion-MNIST, respectively. We use the following network architectures for CIFAR10 and CelebA:

$$\begin{aligned}
 \mathbf{x} \in \mathbb{R}^{32 \times 32 \times 3} &\rightarrow \text{Conv}_{d_b/2}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: } (d_b/2, w/2, h/2) \\
 &\rightarrow \text{Conv}_{d_b}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: } (d_b, w/4, h/4) \\
 &\rightarrow \text{Conv}_{d_b}^{(3 \times 3)} && \text{size of } (d_b, w/4, h/4) \\
 &\rightarrow [\text{ResBlock}_{d_b}]_{\times 6} && \text{size of } (d_b, w/4, h/4) \\
 \mathbf{Z}_q \in \mathbf{B}^{w/4 \times h/4} \subset \mathbb{R}^{d_b \times w/4 \times h/4} &\rightarrow [\text{ResBlock}_{d_b}]_{\times 6} && \text{size of } (d_b, w/4, h/4) \\
 &\rightarrow \text{ConvT}_{d_b/2}^{(3 \times 3)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: } (d_b, w/4, h/4) \\
 &\rightarrow \text{ConvT}_{d_b/2}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: } (d_b/2, w/2, h/2) \\
 &\rightarrow \text{ConvT}_3^{(4 \times 4)} \rightarrow \text{Sigmoid} && \text{size: } (3, w, h),
 \end{aligned}$$

where  $w$  and  $h$  denote the width and height of the target images, respectively.

We use the Adam optimizer (Kingma & Ba, 2015) with initial learning rates of 0.0003 and 0.001 for VQ-VAE and the other models, respectively. The learning rate will be halved every 3 epochs if the validation loss is not improving. We train 100 epochs with the minibatch size of 32 for MNIST, Fashion-MNIST, and CIFAR10 and 70 epochs for CelebA.

Regarding the codebook reset applied in Figure 5 with EMA, we adopt the same procedure as that adopted by Williams et al. (2020). At every 20th batch, the two codes that are most and least used in the recent 20 batches are found, which are denoted as  $\mathbf{b}_{\text{most}}$  and  $\mathbf{b}_{\text{least}}$ , respectively. If the usage of  $\mathbf{b}_{\text{least}}$  is less than 3% of that of  $\mathbf{b}_{\text{most}}$ , the position of  $\mathbf{b}_{\text{least}}$  is reset to  $\mathbf{b}_{\text{least}}^{\text{reset}} \sim \mathcal{N}(\mathbf{b}_{\text{most}}, s_{\text{reset}}^2 \mathbf{I})$  with  $s_{\text{reset}}^2 = 0.01$ .

We set the dimension of the latent space for the VAE on CelebA to 72 such that the number of bits used to represent the latent space is the same as that of the other models, i.e.,  $32 \text{ bits} \times 72 = 9 \text{ bits} \times 16 \times 16$ . We use the same architecture for VAE as that of the other models except appending a linear layer at the end of the encoder and doing the same at the beginning

<sup>7</sup>[https://github.com/deepmind/sonnet/blob/v2/examples/vqvae\\_example.ipynb](https://github.com/deepmind/sonnet/blob/v2/examples/vqvae_example.ipynb)

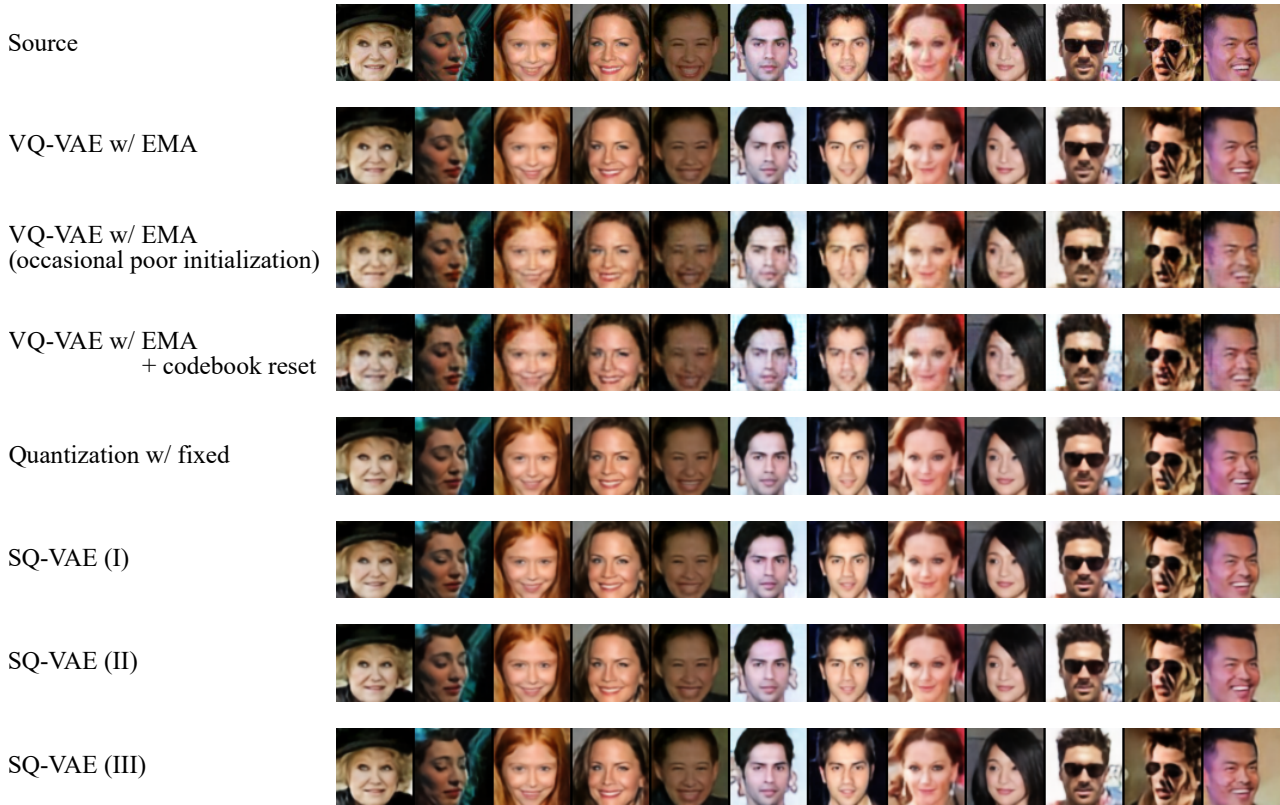


Figure 6. Reconstructed samples of CelebA  $64 \times 64$ .

of the decoder. This is to adjust the dimension of latent variables.

### E.1.3. RECONSTRUCTED AND GENERATED SAMPLES ON CELEBA $64 \times 64$

We show examples of reconstructed images and images generated with the learned approximated prior in Figures 6 and 7. We adopt PixelCNN (van den Oord et al., 2016) as the estimator of the prior while we believe the quality of synthetic samples of SQ-VAE can be improved by using other stronger autoregressive estimators (Child et al., 2019). We observe that the quality of the samples from VQ-VAE varies depending on random seeds, as shown in Figure 6.

## E.2. Gaussian SQ-VAE on Speech Dataset

### E.2.1. DATASETS AND PREPROCESSING

**VCTK** VCTK version 0.80 (Veaux et al., 2017) is a speech dataset of 109 English speakers. Each speaker reads out about 400 sentences, which were mostly selected from a newspaper. We use 90% of the samples as the training set, and the remaining 10% as the test set.

We preprocess waveforms following the approach of the GitHub repository kan-bayashi/ParallelWaveGAN<sup>8</sup>, which is one of the most popular repositories related to speech generation and provides implementations of several types of vocoder such as Parallel WaveGAN (Yamamoto et al., 2020) and MelGAN (Kumar et al., 2019). The details of the preprocessing are as follows.

1. We resample 48000 Hz signals to 24000 Hz.
2. We extract 80-dimensional Mel spectrogram features with a 50.0 ms Hann window, 12.5 ms frame shift, 2048-point FFT, and 80 Hz and 7600 Hz frequency cutoffs.
3. We convert the amplitude spectrograms to dB-scaled spectrograms and clip the bins below  $-200.0$  dB.
4. We perform dimension-wise standardization.

<sup>8</sup><https://github.com/kan-bayashi/ParallelWaveGAN>

Table 8. Notations of network layers on speech datasets.

Notation	Description
$\text{Conv}_n^{(3)}$	1D Convolutional layer (channel= $n$ , kernel= 3, stride= 1)
$\text{Conv}_n^{(4)}$	1D Convolutional layer (channel= $n$ , kernel= 4, stride= 2)
Jitter	Time-jitter regularization (Chorowski et al., 2019) operation
$\text{Concat}(\dots)$	Concatenation operation
$\text{Embedding}_n$	Lookup table (dimension= $n$ )
Interpolate	Nearest-neighbor interpolation operation
$\text{BiGRU}_n$	Bidirectional gated recurrent unit (Cho et al., 2014) layer (dimension= $n$ )

**ZeroSpeech 2019** ZeroSpeech 2019 English (Dunbar et al., 2019) is a multi-speaker corpus sampled at 16000 Hz. This corpus was originally made for a speech task called acoustic unit discovery. It consists of four subsets: Train Voice Dataset (directory `train/voice`), Train Unit Dataset (directory `train/unit`), Train Parallel Dataset (directories `train/parallel/source` and `train/parallel/voice`), and Test Dataset (directory `test`). We use the `.wav` files under the directories `train/voice` and `train/unit` for training data, following van Niekerk et al. (2020), and the `.wav` files under the directory `train/parallel/voice` for testing.

We preprocess waveforms following the approach of van Niekerk et al. (2020). The details of the preprocessing are as follows:

1. We scale the maximum amplitude of each audio signal to 0.999.
2. We pre-emphasize the scaled audio signals with a first-order autoregressive filter:  $y[n] - 0.97 * y[n - 1]$ .
3. We extract 80-dimensional Mel-spectrogram features with a 25.0 ms Hann window, 10.0 ms frame shift, 2048-point FFT, and 50 Hz and 8000 Hz frequency cutoffs.
4. We convert the amplitudes to decibels and clip those bins that are 80.0 dB lower than the maximum.
5. We rescale the dB-scaled Mel-spectrogram by dividing it by 80.0.

### E.2.2. MODEL DESCRIPTION AND TRAINING

We adopt the model proposed by van Niekerk et al. (2020)<sup>9</sup> as a baseline and replace its RNN-based vocoder with a projection layer. The encoder consists of a stack of five convolutional layers, which downsamples the input by 2, and the latent representation is quantized with 512 codes ( $K = 512$ ). The decoder aims to reconstruct the normalized log-Mel spectrogram, conditioned on both the quantized latent representation and a speaker embedding. Time-jitter regularization (Chorowski et al., 2019) is applied with a replacement probability of 0.5.

$$\begin{aligned}
 x \in \mathbb{R}^{80 \times N} &\rightarrow \text{Conv}_{768}^{(3)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: } (768, N) \\
 &\rightarrow \text{Conv}_{768}^{(3)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: } (768, N) \\
 &\rightarrow \text{Conv}_{768}^{(4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: } (768, N/2) \\
 &\rightarrow \text{Conv}_{768}^{(3)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: } (768, N/2) \\
 &\rightarrow \text{Conv}_{768}^{(3)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: } (768, N/2) \\
 &\rightarrow \text{FC}_{64} && \text{size: } (64, N/2), \\
 z \in \mathbb{R}^{64 \times N/2} &\rightarrow \text{Jitter} && \text{size: } (64, N/2) \\
 &\rightarrow \text{Concat}(\text{Embedding}_{64}) && \text{size: } (128, N/2) \\
 &\rightarrow \text{Interpolate} && \text{size: } (128, N) \\
 &\rightarrow \text{BiGRU}_{128} && \text{size: } (256, N) \\
 &\rightarrow \text{BiGRU}_{128} && \text{size: } (256, N) \\
 &\rightarrow \text{FC}_{80} && \text{size: } (80, N).
 \end{aligned}$$

The notations of the layers are listed in Table 8, and  $N$  denotes the number of frames of an input data.

Following van Niekerk et al. (2020), the model is trained on a minibatch of 52 segments, each 32 frames long. We use the Adam optimizer with an initial learning rate of  $4 \cdot 10^{-4}$ , and the learning rate is halved after 300k and 400k steps. The network

<sup>9</sup><https://github.com/bshall/ZeroSpeech>

Table 9. Evaluation on VCTK and ZeroSpeech 2019. The MSE (dB<sup>2</sup>) of sample reconstruction is evaluated using the test set.

Model	MSE (dB <sup>2</sup> )	
	VCTK	ZeroSpeech 2019
VQ-VAE w/ EMA ( $\sigma^2 = 10^{-2}$ )	31.25 ± 0.40	35.20 ± 1.21
VQ-VAE w/ EMA ( $\sigma^2 = 10^{-1}$ )	30.89 ± 0.46	34.33 ± 1.57
VQ-VAE w/ EMA ( $\sigma^2 = 10^0$ )	29.59 ± 0.25	40.40 ± 1.24
VQ-VAE w/ EMA ( $\sigma^2 = 10^1$ )	36.92 ± 0.95	79.30 ± 42.95
Gaussian SQ-VAE (I)	25.52 ± 0.08	33.17 ± 1.11
Gaussian SQ-VAE (III)	25.94 ± 0.22	34.35 ± 1.07
Gaussian SQ-VAE (IV)	<b>24.68</b> ± 0.21	<b>32.32</b> ± 0.88

Table 10. Evaluation on ZeroSpeech 2019. The average ABX score is evaluated on the test set.

Model	ABX score ↓
VQ-VAE w/ EMA ( $\sigma^2 = 10^{-2}$ )	22.36 ± 0.14
VQ-VAE w/ EMA ( $\sigma^2 = 10^{-1}$ )	22.53 ± 0.33
VQ-VAE w/ EMA ( $\sigma^2 = 10^0$ )	23.78 ± 0.74
VQ-VAE w/ EMA ( $\sigma^2 = 10^1$ )	35.40 ± 6.73
Gaussian SQ-VAE (I)	<b>22.11</b> ± 0.62
Gaussian SQ-VAE (III)	22.13 ± 0.29
Gaussian SQ-VAE (IV)	24.46 ± 0.66

is trained for a total of 500k steps.

We do not apply SQ-VAE (II) in this evaluation because of the variable length property of speech data and the different manipulations of speech signals between training and inference. SQ-VAE (II) has one parameter  $\sigma_\varphi^2(\mathbf{x})$  for one trimmed speech signal (32 frames long) in training but has to deal with longer signals in inference.  $\sigma_\varphi^2(\mathbf{x})$  is calculated based on the content of a trimmed signal during training regardless of the content of the whole signal. That means there is a discrepancy between training and inference with SQ-VAE (II). This discrepancy does not exist when it is applied to image datasets, where the size of images is fixed. On the other hand, SQ-VAE (III) has a parameter  $\sigma_{\varphi,i}^2(\mathbf{x})$  for every 2 consecutive frames regardless of the signal length. This parameterization is applicable even when signals are trimmed to 32 frames long during training but are not trimmed during inference. The same is true for SQ-VAE (IV).

### E.2.3. DETAILS OF EXPERIMENTAL RESULTS

The MSEs of VQ-VAE models with various  $\sigma^2$  values are shown in Table 9.

### E.2.4. RECONSTRUCTED SAMPLES

As a sample of demonstration, we randomly select two speech signals from our test split of VCTK and show their reconstructed log-Mel spectrograms from VQ-VAE and SQ-VAE in Figures 8 and 9, respectively.

### E.2.5. ACOUSTIC UNIT DISCOVERY

We compare the performance of SQ-VAE with that of VQ-VAE in the acoustic unit discovery task, which evaluates the feasibility of a representation to discriminate phonetic units. VQ-VAE is a popular approach used for this task (Chorowski et al., 2019; Eloff et al., 2019; van Niekerk et al., 2020; Tjandra et al., 2020).

We follow the evaluation scheme of ZeroSpeech 2019 (Dunbar et al., 2019) and use the minimal pair ABX discriminability test (Schatz et al., 2013) for comparison. This test asks whether a triphone X is more similar to triphone A than triphone B. Here, A and X are instances of the same triphone (e.g., “beg”), whereas B differs in the middle phone (e.g., “bag”). Moreover, A and B are selected from the same speaker, but X is selected from a different speaker. An ABX score is reported as an aggregated error rate over all pairs of triphones in the test set. The lower the ABX score, the better the performance. We compute the ABX scores of the trained models using Test Dataset (the .wav files under the directory test).

Surprisingly, as shown by the results in Table 10, SQ-VAE is on par with VQ-VAE in this experiment, with no statistically significant difference. In addition, SQ-VAE (IV) performs slightly worse than SQ-VAE (I) or SQ-VAE (III) here, although SQ-VAE (IV) is better in terms of MSE as shown in Table 3. We plan to perform further analysis on the learned representations in our future work.

### E.3. vMF SQ-VAE on Vision Dataset

#### E.3.1. DATASETS AND PREPROCESSING

**CelebAHQ-Mask** CelebAHQ-Mask is a dataset of colored face images with segmentation maps, which have 19 categories including all facial components and accessories such as skin, nose, eyes, eyebrows, ears, mouth, lip, hair, hat, eyeglass, earring, necklace, neck, and cloth.

**Gray-CelebA** We obtain Gray-CelebA by converting CelebA to grayscale. As in CelebA, we preprocess the grayscaled images with center cropping of  $140 \times 140$  and resized to  $64 \times 64$  using the bilinear interpolation.

#### E.3.2. SETUP

**CelebAHQ-Mask** We set the codebook capacity to  $(d_b, K) = (64, 64)$  for all the discrete models. In SQ-VAE, we introduce a hypersphere  $\mathcal{S}^{19-1}$  as a space for projecting data categories. We set the projections of the categories  $(\mathbf{w}_l)_{l=1}^L$  to one-hot vectors, where all  $\mathbf{w}_l$  reside on  $\mathcal{S}^{18}$  and are orthogonal to each others.

For CIFAR10 and CelebA, we use the following networks for the encoder and decoder, respectively.

$$\begin{aligned}
 \mathbf{x} \in \mathbb{R}^{64 \times 64 \times 19} &\rightarrow \text{Conv}_{32}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (32, 32, 32)} \\
 &\rightarrow \text{Conv}_{64}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (64, 16, 16)} \\
 &\rightarrow \text{Conv}_{64}^{(4 \times 4)} && \text{size of (64, 8, 8)} \\
 &\rightarrow [\text{ResBlock}_{64}]_{\times 2} && \text{size of (64, 8, 8)} \\
 \mathbf{Z}_q \in \mathbf{B}^{8 \times 8} \subset \mathbb{R}^{64 \times 8 \times 8} &\rightarrow [\text{ResBlock}_{64}]_{\times 2} && \text{size of (64, 8, 8)} \\
 &\rightarrow \text{ConvT}_{64}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (64, 16, 16)} \\
 &\rightarrow \text{ConvT}_{32}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (32, 32, 32)} \\
 &\rightarrow \text{ConvT}_{19}^{(4 \times 4)} && \text{size: (19, 64, 64)}.
 \end{aligned}$$

In vMF SQ-VAE, the outputs of the encoder and decoder are normalized along the channel axis so that the normalized vectors are on the hypersphere  $\mathcal{S}^{64-1}$  and  $\mathcal{S}^{19-1}$ , respectively.

In this experiment, the optimizer, the initial learning rate and the learning rate scheduling are the same as those in Appendix E.1. We run 70 epochs with a minibatch size of 32.

In VAE, we set the dimension of the latent space to 12 such that the number of bits representing the latent space is the same as that of the other models, i.e.,  $32 \text{ bit} \times 12 = 6 \text{ bit} \times 8 \times 8$ . We use the same architecture for VAE as those of the other models except adding a linear layer to the end of the encoder.

**MNIST and Gray-CelebA** We set the codebook capacity  $(d_b, K)$  to  $(64, 128)$  and  $(64, 512)$  on MNIST and Gray-CelebA, respectively. We introduce a unit circle  $\mathcal{S}^1$  as a space for projecting data categories and set  $(\mathbf{w}_l)_{l=1}^L$  as  $\mathbf{w}_l = [\cos(\alpha_l), \sin(\alpha_l)]^\top$  with  $\alpha_l = \frac{\pi}{L}l$ , where all  $\mathbf{w}_l$  reside on  $\mathcal{S}^1$ . The common network architecture for MNIST and Gray-CelebA is as follows:

$$\begin{aligned}
 \mathbf{x} \in \mathbb{R}^{28 \times 28 \times 256} &\rightarrow \text{Conv}_{32}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (32, 14, 14)} \\
 &\rightarrow \text{Conv}_{64}^{(4 \times 4)} && \text{size of (64, 7, 7)} \\
 &\rightarrow [\text{ResBlock}_{64}]_{\times 2} && \text{size of (64, 7, 7)} \\
 \mathbf{Z}_q \in \mathbf{B}^{7 \times 7} \subset \mathbb{R}^{64 \times 7 \times 7} &\rightarrow [\text{ResBlock}_{64}]_{\times 2} && \text{size of (64, 7, 7)} \\
 &\rightarrow \text{ConvT}_{32}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (32, 14, 14)} \\
 &\rightarrow \text{ConvT}_c^{(4 \times 4)} && \text{size: (c, 28, 28)},
 \end{aligned}$$

and

$$\begin{aligned}
 \mathbf{x} \in \mathbb{R}^{64 \times 64 \times 256} &\rightarrow \text{Conv}_{32}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (32, 32, 32)} \\
 &\rightarrow \text{Conv}_{64}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (64, 16, 16)} \\
 &\rightarrow \text{Conv}_{64}^{(4 \times 4)} && \text{size of (64, 8, 8)} \\
 &\rightarrow [\text{ResBlock}_{64}]_{\times 6} && \text{size of (64, 8, 8)} \\
 \mathbf{Z}_q \in \mathbf{B}^{8 \times 8} \subset \mathbb{R}^{64 \times 8 \times 8} &\rightarrow [\text{ResBlock}_{64}]_{\times 6} && \text{size of (64, 8, 8)} \\
 &\rightarrow \text{ConvT}_{64}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (64, 16, 16)} \\
 &\rightarrow \text{ConvT}_{32}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (32, 32, 32)} \\
 &\rightarrow \text{ConvT}_c^{(4 \times 4)} && \text{size: (c, 64, 64),}
 \end{aligned}$$

where  $c$  is the the number of output channel, which is 2 for vMF SQ-VAE and 256 for other models. In vMF SQ-VAE, the outputs of the encoder and decoder are normalized along the channel axis so that the normalized vectors are on hyperspheres  $\mathcal{S}^{64-1}$  and  $\mathcal{S}^{2-1}$ , respectively.

In this experiment, the optimizer, the initial learning rate and the learning rate scheduling are the same as those in Appendix E.1. We run 100 and 70 epochs on MNIST and Gray-CelebA, respectively, with a minibatch size of 32.

We set the latent space dimension of VAE to 11 on MNIST and 18 on Gray-CelebA. In this way, the number of bits representing the latent space becomes almost the same as that of the other models, i.e.,  $32 \text{ bit} \times 11 \approx 7 \text{ bit} \times 7 \times 7$  and  $32 \text{ bit} \times 18 = 9 \text{ bit} \times 8 \times 8$ .

## F. Experiments on CelebA HQ $256 \times 256$

To demonstrate that our SQ-VAE is applicable to larger scale datasets, we apply Gaussian SQ-VAE (I) to CelebA HQ  $256 \times 256$ . We train Gaussian SQ-VAE (I) using the following network architecture:

$$\begin{aligned}
 \mathbf{x} \in \mathbb{R}^{256 \times 256 \times 3} &\rightarrow \text{Conv}_{16}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (16, 128, 128)} \\
 &\rightarrow \text{Conv}_{32}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (32, 64, 64)} \\
 &\rightarrow \text{Conv}_{64}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (64, 32, 32)} \\
 &\rightarrow \text{Conv}_{64}^{(3 \times 3)} && \text{size of (64, 32, 32)} \\
 &\rightarrow [\text{ResBlock}_{64}]_{\times 6} && \text{size of (64, 32, 32)} \\
 \mathbf{Z}_q \in \mathbf{B}^{32 \times 32} \subset \mathbb{R}^{64 \times 32 \times 32} &\rightarrow [\text{ResBlock}_{64}]_{\times 6} && \text{size of (64, 32, 32)} \\
 &\rightarrow \text{ConvT}_{64}^{(3 \times 3)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (64, 32, 32)} \\
 &\rightarrow \text{ConvT}_{32}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (32, 64, 64)} \\
 &\rightarrow \text{ConvT}_{16}^{(4 \times 4)} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} && \text{size: (16, 128, 128)} \\
 &\rightarrow \text{ConvT}_3^{(4 \times 4)} \rightarrow \text{Sigmoid} && \text{size: (3, 256, 256),}
 \end{aligned}$$

We follow the same experimental setup as that on CelebA  $64 \times 64$  except for the network architecture. We show examples of reconstructed images of the models in Figure 10(b).

To ease the difficulty of training a feasible prior, we adapt the hierarchical structure which is popular in the recent autoencoder-based works (Razavi et al., 2019; Dhariwal et al., 2020; Williams et al., 2020). We follow a similar procedure to that of Dhariwal et al. (2020) by stacking two SQ-VAE models. The architecture of the first SQ-VAE is described as above. The second SQ-VAE has two extra  $\text{Conv}^{(4 \times 4)}$  layers at the end of its encoder and the beginning of its decoder. Also, its latent space is  $\mathbf{B}^{8 \times 8}$ . We use two PixelSNAIL (Chen et al., 2018) models to act as the prior and the upsampler between the latent codes of the two SQ-VAEs. As a result, the second PixelSNAIL is conditioned by the latent code of the first SQ-VAE. We show images generated with the learned approximated prior in Figure 10(c). As we have demonstrated in Section 5 that SQ-VAE enables to learn good discrete latent features and produce superior reconstructed samples, we believe the quality of synthetic samples of SQ-VAE can be improved via other stronger autoregressive estimators.




(a) VQ-VAE



(b) VQ-VAE (with occasional poor initialization)



(c) Quantization w/ fixed  $\sigma_q^2$



(d) SQ-VAE (I)



(e) SQ-VAE (II)



(f) SQ-VAE (III)

Figure 7. Generated samples of CelebA  $64 \times 64$ .

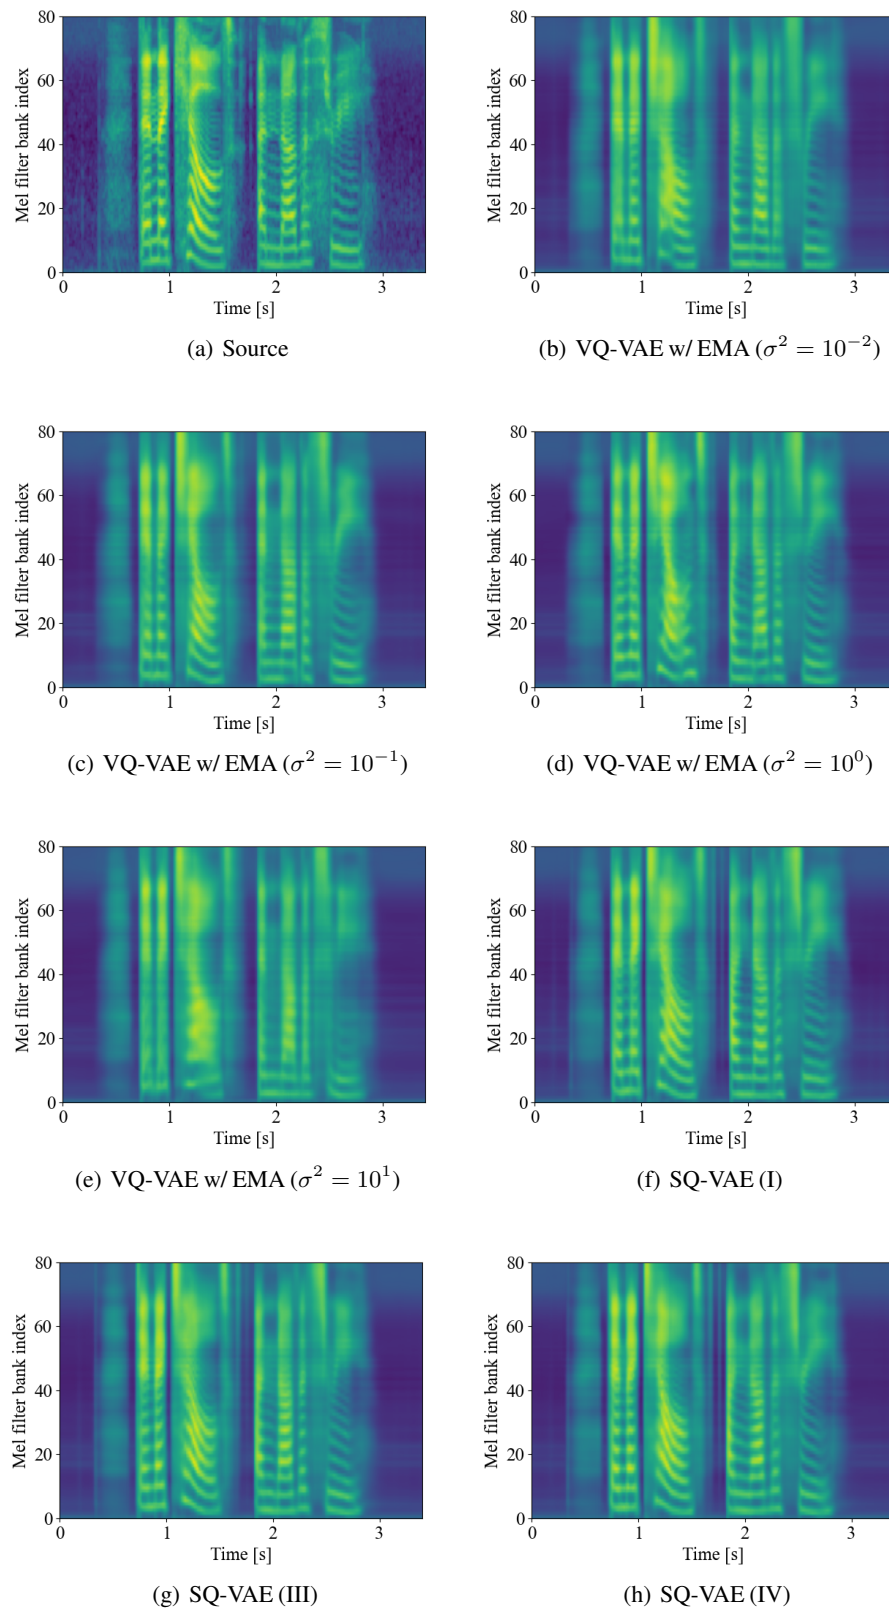


Figure 8. Reconstructed log-Mel spectrograms of p323\_064 in VCTK.



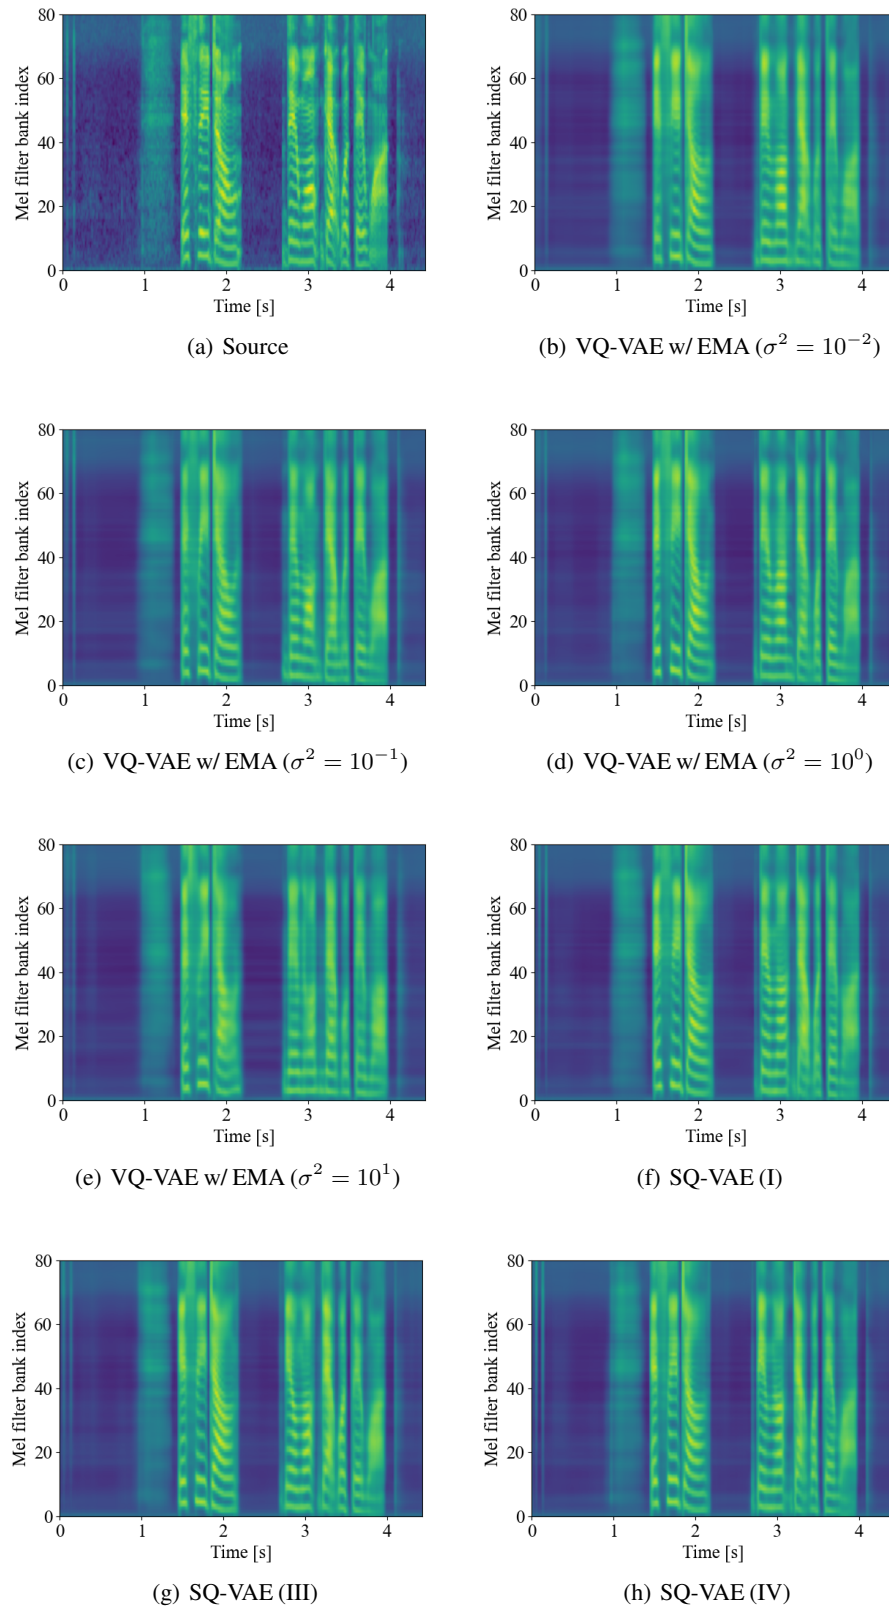


Figure 9. Reconstructed log-Mel spectrograms of p341\_285 in VCTK.

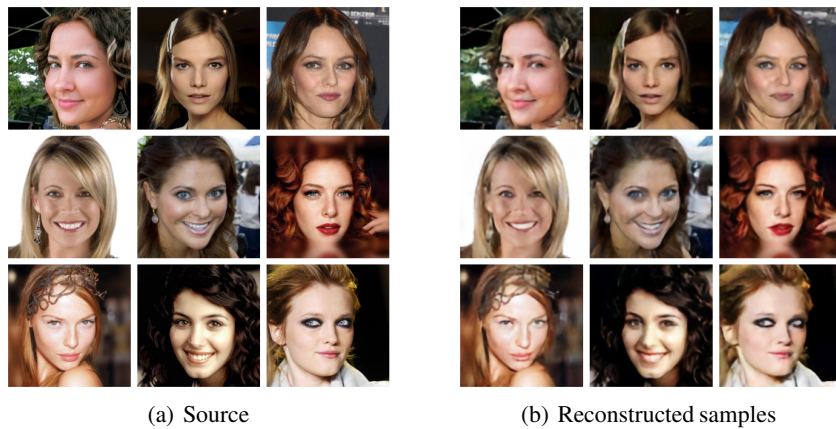


Figure 10. Samples of CelebAHQ  $256 \times 256$  from SQ-VAE (I).