

NVTC: Nonlinear Vector Transform Coding

Runsen Feng Zongyu Guo Weiping Li Zhibo Chen
 University of Science and Technology of China

{fengruns, guozy}@mail.ustc.edu.cn, {wpli, chenzhibo}@ustc.edu.cn

Abstract

In theory, vector quantization (VQ) is always better than scalar quantization (SQ) in terms of rate-distortion (R-D) performance [33]. Recent state-of-the-art methods for neural image compression are mainly based on nonlinear transform coding (NTC) with uniform scalar quantization, overlooking the benefits of VQ due to its exponentially increased complexity. In this paper, we first investigate on some toy sources, demonstrating that even if modern neural networks considerably enhance the compression performance of SQ with nonlinear transform, there is still an insurmountable chasm between SQ and VQ. Therefore, revolving around VQ, we propose a novel framework for neural image compression named Nonlinear Vector Transform Coding (NVTC). NVTC solves the critical complexity issue of VQ through (1) a multi-stage quantization strategy and (2) nonlinear vector transforms. In addition, we apply entropy-constrained VQ in latent space to adaptively determine the quantization boundaries for joint rate-distortion optimization, which improves the performance both theoretically and experimentally. Compared to previous NTC approaches, NVTC demonstrates superior rate-distortion performance, faster decoding speed, and smaller model size. Our code is available at <https://github.com/USTC-IMCL/NVTC>.

1. Introduction

Recent works based on nonlinear transform coding (NTC) [5] have achieved remarkable success in neural image compression [12, 34]. Unlike these traditional image codecs that employ linear transform such as discrete cosine transform (DCT), NTC is constructed with the nonlinear transform layers and optimized with data-driven techniques, where the modern neural networks present excellent capability in both encoding/decoding transform and entropy estimation [12, 20, 35]. Most NTC methods apply scalar quantization (SQ) to discretize the latent variables and use the additive uniform noise to approximate the quantization error during training [6]. However, in the era of 1990s, it

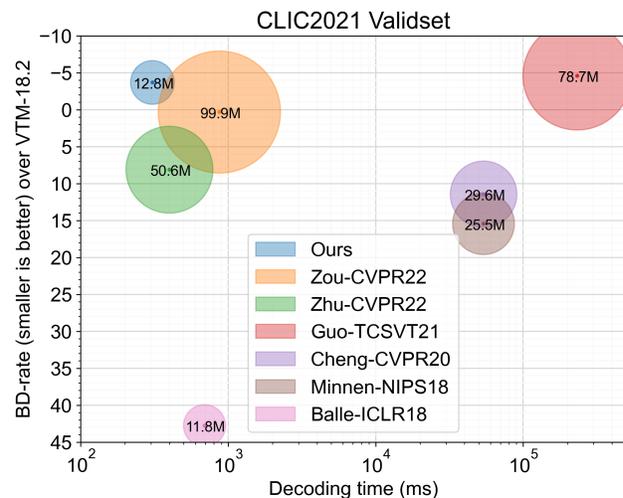


Figure 1. BD-rate vs. decoding time vs. model size on the CLIC2021 validation set [1].

has already been known that vector quantization, in spite of it exponentially increased complexity, is always better than SQ in terms of rate-distortion (RD) performance [33]. It inspires us to design a novel neural image compression model to fully leverages vector quantization.

Vector quantization (VQ) [18] is designed to map a continuous source distribution to a set of discrete vectors. The discrete nature of VQ has been successfully applied in generative models to avoid the “posterior collapse” issue, including these well-known image synthesis models such as VQVAE [40], VQGAN [16] and text-to-image synthesis models such as DALLE [36], latent diffusion [38]. However, if we go back to the basic requirement of quantization, we will find that VQ offers unique advantages in terms of rate-distortion performance, particularly the space-filling advantage and the memory advantage [33].

Given a source distribution, the goal of quantization (no matter SQ or VQ) is to determine the quantization centers and boundaries, and then assign indices to denote these separated quantization regions/cells. Combining these regions fills the whole space of source distribution. The space-filling advantage of VQ against SQ is related to the sphere

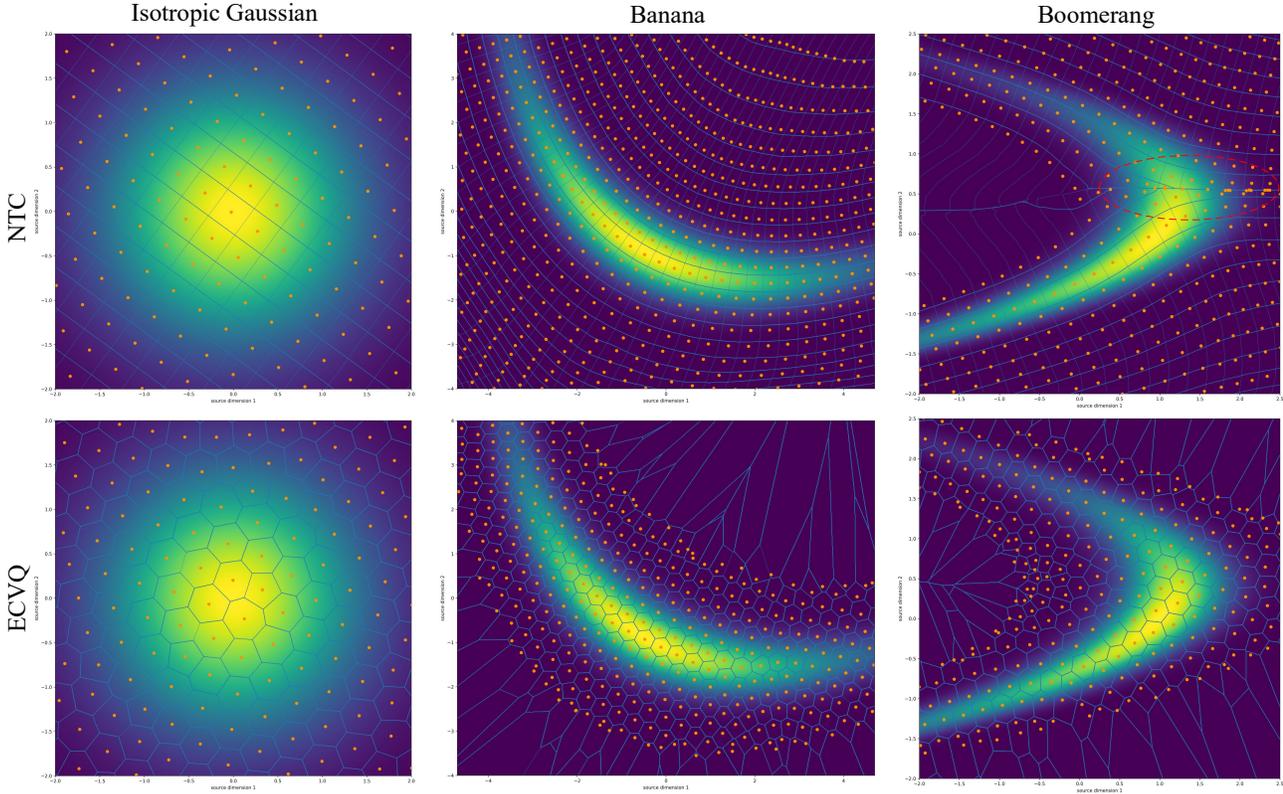


Figure 2. Quantization results for NTC (SQ with nonlinear transform) and ECVQ (entropy-constrained VQ) on 2-d distributions. Blue lines represent quantization boundaries and orange points represent quantization centers (codewords). “Isotropic Gaussian” refers to a 2-d isotropic Gaussian distribution. “Banana” and “Boomerang” are two different 2-d distributions. It is observed that NTC cannot achieve the space-filling advantage (i.e. learning hexagon-like quantization cells for 2-d sources) even on an isotropic Gaussian distribution. Moreover, NTC’s decorrelation capability is insufficient as source correlation becomes more complex. For example, quantization boundaries collide in the red circle of “Boomerang”, leading to a performance drop. The corresponding BD-PSNR results are shown in Table 2.

packing problem in geometry [14, 21, 41]. If we compare the quantization results of SQ and VQ, as shown in Figure 2, we will find that even for a simple isotropic Gaussian distribution, SQ with nonlinear transform cannot learn to approximate hexagon-like quantization cells, where the hexagon is the polytope with the best space-filling properties in 2-d space. Under the high-rate assumption, the gain of space-filling advantage is about 1.53 dB as the dimension approaches infinity [15, 33]. Following this conclusion, we experimentally provide the BD-PSNR results by comparing SQ with nonlinear transform to VQ on isotropic Gaussian distributions in Table 1. In addition, to reduce the redundancies of data distributions, existing NTC methods (SQ with nonlinear transform) rely on highly expensive nonlinear transform [12, 44, 45] and context-based autoregressive entropy models [20, 35]. However, different from NTC methods, VQ has superior decorrelation ability, which is known as the memory advantage of vector quantizers. This advantage is more obvious when quantizing complex source distributions, such as the Boomerang distribution in Figure 2 (especially in the red circle area).

In this paper, we build a novel framework that applies modern neural networks to leverage the space-filling advantages and memory advantages of VQ for image compression. We propose nonlinear vector transform coding (NVTC), which achieves encouraging rate-distortion performance with relatively low coding complexity. Specifically, as shown in Figure 3, we introduce three key points to design a practical VQ, including 1) a multi-stage product VQ rather than a single-stage VQ to reduce the exponentially increased complexity, 2) nonlinear vector transform rather than scalar transform to remove redundancy between sub-vectors with fewer parameters, and 3) entropy-constrained VQ rather than unconstrained VQ to achieve superior R-D optimality and joint optimization of latent-space VQ models.

For **the first point**, many well-known VQ variants have been proposed in recent decades, such as product VQ [22, 39], multi-stage VQ [23], tree-structured VQ [11] and lattice VQ [17]. Although tree-structured and lattice VQ offer fast encoding speeds, they do not reduce the storage complexity of codebooks or entropy-coding frequency tables. In

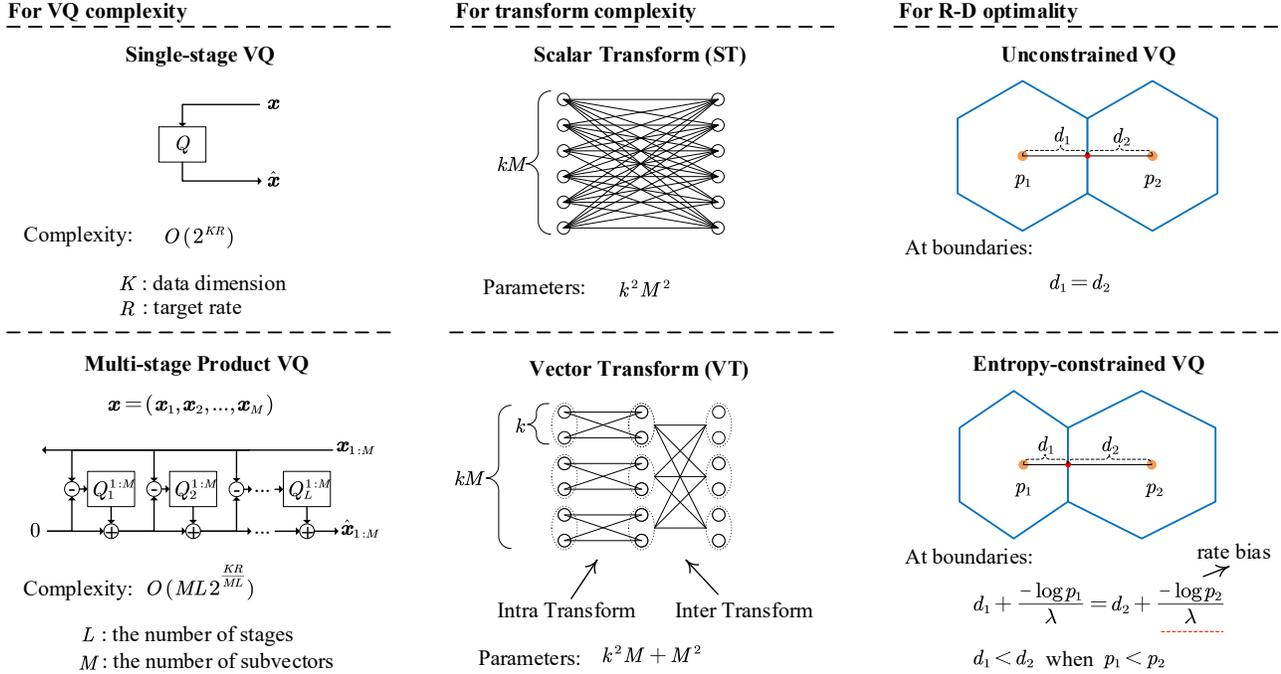


Figure 3. Three key points to design a practical vector quantizer. For VQ complexity (left), we suggest a hybrid VQ structure called multi-stage product VQ which reduces the VQ complexity from $O(2^{KR})$ to $O(ML2^{\frac{KR}{ML}})$. For transform complexity (middle), we use vector transform instead of scalar transform to remove inter-vector redundancy. For RD optimality (right), we find that ECVQ [13] is essential for the joint rate-distortion optimization, which is neglected in previous works [2, 31, 40, 43]

this paper, we suggest a hybrid VQ structure that incorporates both product VQ and multi-stage VQ, as shown in the left column of Figure 3. The quantization procedure comprises multiple stages, and each stage employs multiple independent low-dimension quantizers to compress the subvectors of the input vector. As the number of stages and subvectors increases, the proposed multi-stage product VQ exhibits a significant decrease in complexity.

While the intra-vector redundancy (i.e. the redundancy inside each subvector) can be removed by vector quantization, the inter-vector redundancy (i.e. the redundancy between subvectors) is still overlooked. Therefore, **our second point** focuses on efficiently eliminating inter-vector redundancy. Transform VQ [25] introduces a linear transform for decorrelation and performs product quantization on the transformed coefficients. Similar coding structures are observed in recent learning-based VQ methods [2, 43], which are improved by learnable nonlinear transform with superior decorrelation capabilities. However, the transform used in these works is designed to decorrelate scalar components, which is computationally inefficient for vector decorrelation. The intra-vector redundancy, which is intended to be removed by VQ, might be partially reduced in advance by the scalar transform. Therefore, certain parts of the scalar transform could be eliminated to improve com-

putational efficiency. Motivated by the linear vector transform [28, 30, 31], we propose a new VT variant that decouples a fully-connected scalar transform into two lightweight parts: intra-transform and inter-transform. In the middle of Figure 3, we provide a simple comparison between the scalar transform and the proposed vector transform. We further stack the single-layer VT to build a powerful nonlinear vector transform. The differences between our VT and the linear VT are discussed in Section 4.1.

Regarding **the third point**, we emphasize that the quantization process (either SQ or VQ) used in most previous methods [2, 5, 31, 40, 43] (including VQVAE) is not entropy-constrained, which is theoretically suboptimal for rate-distortion performance. In the right of Figure 3, we provide a quantization illustration of unconstrained VQ and entropy-constrained VQ (ECVQ [13]), where unconstrained VQ determines the quantization boundaries (blue line) using the nearest neighbor search. ECVQ introduces an additional rate bias $\frac{-\log p_i}{\lambda}$ in the quantization process, which shifts the quantization boundaries from the high-probability region to the low-probability region. In other words, ECVQ search the codewords with the best RD performance, instead of just the neighboring codewords. ECVQ provides an optimal VQ encoding process described in Section 2. With the help of latent-space ECVQ, we de-

sign a training strategy for joint RD optimization. Instead of manually controlling the RD trade-off by varying codebook size [43], our model can learn layer-adaptive bit allocation.

Our contributions can be summarized as 1) investigating on VQ advantages over SQ with nonlinear transform based on empirical results on some toy sources, 2) presenting a VQ-based coding scheme named nonlinear vector transform coding (NVTC) with three technical contributions that effectively leverages VQ while keeping complexity low, and 3) demonstrating that NVTC offers superior rate-distortion performance, faster decoding speed and smaller model size, compared with previous neural image codecs.

2. Related Works & Background

The task of lossy coding aims to map the k -dimensional Euclidean space \mathbb{R}^k of a data source into a set of vectors $\mathbf{c} = \{\mathbf{c}_i \in \mathbb{R}^k | i = 0, 1, \dots, N - 1\}$. At the encoder side, each sample vector \mathbf{x} can be represented by a reproduction vector $\hat{\mathbf{x}} = \mathbf{c}_i \in \mathbf{c}$. The vector index i is losslessly coded by an entropy model $P(i)$, which is an estimate of the probability distribution of the index. Both the vector set \mathbf{c} and entropy model P are known to the decoder. Thus, $\hat{\mathbf{x}}$ can be reconstructed by the entropy decoded index i at the decoder side. The average code length R , also called rate, is given by the cross entropy between the actual distribution and the estimated distribution: $R = \mathbb{E}_{\mathbf{x}}[-\log P(i)]$. And the average distortion is $\mathbb{E}_{\mathbf{x}}[d(\mathbf{x}, \hat{\mathbf{x}})]$ given a distortion metric d .

Unconstrained VQ. A vector quantizer [18] Q of dimension k and size N is defined as a mapping from \mathbb{R}^k into a finite set $\mathbf{c} = \{\mathbf{c}_i \in \mathbb{R}^k | i = 0, 1, \dots, N - 1\}$. The set \mathbf{c} here is called codebook and the vector \mathbf{c}_i is called codeword. Given an input vector \mathbf{x} , the encoder \mathcal{E} searches for the best codeword \mathbf{c}_i and outputs the corresponding index $i = \mathcal{E}(\mathbf{x}) = \arg \min_i d(\mathbf{x}, \mathbf{c}_i)$. With the input index i , the decoder \mathcal{D} looks up the codebook to produce the reproduction vector $\hat{\mathbf{x}} = \mathcal{D}(i) = \mathbf{c}_i$.

ECVQ. In unconstrained VQ, the input \mathbf{x} is assigned to the codeword \mathbf{c}_i with the lowest distortion $d(\mathbf{x}, \mathbf{c}_i)$. However, it is suboptimal for RD loss: $\mathbb{E}_{\mathbf{x}}[-\log P(i) + \lambda d(\mathbf{x}, \mathcal{D}(\mathcal{E}(\mathbf{x})))]$, where λ is the Lagrange multiplier that controls RD trade-off. Entropy-constrained vector quantization (ECVQ) [13] introduces this Lagrangian formulation into the encoder function: $\mathcal{E}(\mathbf{x}) = \arg \min_i [-\log P(i) + \lambda d(\mathbf{x}, \mathbf{c}_i)]$, which is exactly the RD loss of a sample. It means that ECVQ always produces optimal quantization results for each sample, given the codebook and entropy model. In the middle of Figure 3, we illustrate the differences between unconstrained VQ and ECVQ, where ECVQ shifts the quantization boundaries (blue line) from the high-probability region to the low-probability region.

NTC. Nonlinear transform coding (NTC) [5] is a nonlinear extension of transform coding [19], which simplifies

vector quantization as scalar quantization and decorrelates data sources using transform and entropy model. NTC is the most popular coding structure in recent methods for learned image compression. The state-of-art NTC methods improve the components of transform coding by designing powerful nonlinear transforms [12, 44, 45], complex forward-adaptive or backward-adaptive entropy models [7, 20, 34, 35], and differentiable scalar quantization [3, 42]. In our paper, we will use NTC to denote SQ with nonlinear transform.

Latent-space VQ. Prior works such as transform VQ [25] and vector transform coding [28, 30, 31] (VTC) first propose to vector quantize the transform coefficients decades ago. However, due to the high complexity of VQ and the weak decorrelation capabilities of linear transform, they draw less attention to the researchers compared to transform coding [19]. Recently, VQ is known to be effective to address the ‘‘posterior collapse’’ issue in many generative models [16, 36, 40]. In these works, VQ is integrated into an autoencoder network for quantizing the bottleneck latents. Latent-space VQ is also considered in learned image compression [2, 43]. [2] propose a soft relation of VQ for end-to-end optimization, which is gradually annealed to real VQ during training. [43] propose to replace the univariate prior with vectorized prior in existing NTC methods, where a probabilistic vector quantizer is proposed to estimate the prior parameters.

	2d	4d	8d	16d
BD-PSNR (dB)	-0.15	-0.31	-0.47	-0.71

Table 1. BD-PSNR (smaller is worse) of NTC over ECVQ on the isotropic Gaussian distributions with different dimensions.

	IG	Banana	Boomerang
BD-PSNR (dB)	-0.15	-0.21	-0.58

Table 2. BD-PSNR (smaller is worse) of NTC over ECVQ on the 2-d distributions in Figure 2. IG means Isotropic Gaussian.

3. VQ Advantages on Toy Sources

Under high-rate assumptions, vector quantization has three advantages over scalar quantization: space-filling advantage, shape advantage and memory advantage [33]. The first advantage comes from the space-filling properties of different polytopes which is independent of the source distribution, making it a fundamental advantage of VQ over scalar quantization. The second advantage depends solely on the shape of the marginal density function, regardless of the correlation of vector components. The third one reveals the decorrelation capability, which is the major advantage

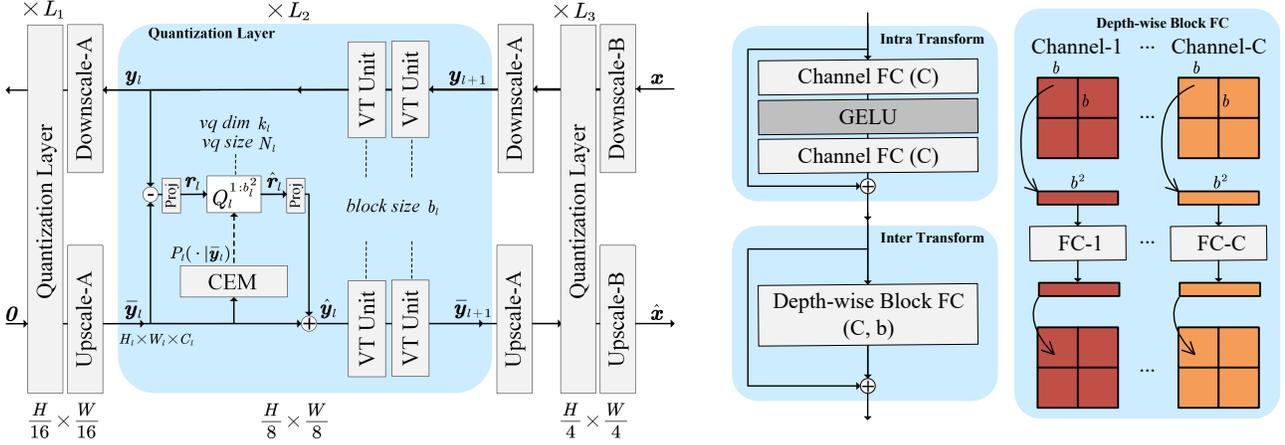


Figure 4. (Left) Overview of our proposed model for image compression. VT means vector transform and CEM means conditional entropy model. Proj is a linear projection layer. CEM is described in Section 4.2. (Right) VT unit with depth-wise block FC. FC refers to a fully-connected layer, and Channel FC is a fully-connected layer along the channel axis. See Section 4.1 for more details.

when dealing with highly correlated source data. The shape advantage could be obtained by accurate entropy modeling. Therefore, in this section, we focus on investigating whether NTC effectively achieves the space-filling and memory advantages. More results can be found in the supplementary.

To verify the space-filling advantage, we first conduct experiments on isotropic Gaussian distributions, where the memory advantage of VQ does not exist. On the left of Figure 2 we visualize the quantization results on a 2d isotropic Gaussian. It can be observed that NTC has rectangular-like quantization cells, different from hexagon-like cells in ECVQ. Hexagon has better space-filling efficiency than rectangular in 2d space. The gain of space-filling advantage increases as the source dimension grows higher. Under the high-rate assumption, the approximate gain of space-filling advantage is 1.53 dB as the dimension approaches infinity [15, 33]. We also conduct experiments to verify the empirical rate-distortion performance on 4d, 8d and 16d isotropic Gaussian. As shown in Table 1, the BD-PSNR [9] gain of ECVQ is surprisingly 0.71 dB on 16d isotropic Gaussian, which is about 50% of the estimated gain (i.e. 1.53 dB) on infinite-dimensional distributions. Considering that VQ complexity grows exponentially with dimension, a 16-d vector quantizer is sufficient to achieve the space-filling advantage under the current complexity constraints.

For memory advantage, we evaluate NTC and ECVQ on several 2d correlated sources. In Figure 2 we show the results on two distributions: one is banana-like distribution named “Banana” that comes from [5]; the other is named “Boomerang” which has a larger bending angle than “Banana”. We find that the BD-PSNR gains of ECVQ on these distributions are more significant than the gain of space-filling advantage alone (i.e. 0.15 dB on 2d isotropic Gaussian). Even if the local transformation of NTC is roughly or-

thogonalized [5], there are some failure cases in the highly nonlinear density area (see red circle in Figure 2). In a word, the potential gain of memory advantage is also significant when the source density is complex.

4. Nonlinear Vector Transform Coding

The overview of the proposed method is shown in Figure 4. Following the structure of multi-stage product VQ in Figure 3, NVTC comprises multiple quantization layers, where quantization is performed sequentially from the leftmost layer to the rightmost layer. Motivated by the hierarchical structure in VAE-based methods [7, 37], we integrate Downscale/Upscale layers between the quantization layers to facilitate global-to-local quantization. In other words, NVTC first quantizes the global vector representation derived from big image patches and then quantizes the local vector representation derived from small image patches.

Let’s take a deeper look at the quantization layer l . Note that N_l is the VQ size, k_l is the VQ dimension and b_l is the block size. The vector representation \mathbf{y}_l has a spatial size of $H_l \times W_l$ and a channel size of C_l . Each channel tensor is considered as a subvector. The vector transform and product vector quantizers $Q_l^{1:b_l^2}$ are shared across different spatial blocks of size $b_l \times b_l$, and for each block, b_l^2 sub-vectors are processed.

Given \mathbf{y}_{l+1} from layer $l + 1$, \mathbf{y}_l is obtained by $\mathbf{y}_l = f_l(\mathbf{y}_{l+1})$, where f_l is vector transform layer that consists of two vector transform units with a block size b_l . With the output $\bar{\mathbf{y}}_l$ of the layer $l - 1$, the residual vectors $\mathbf{r}_l = \text{Lin1}(\mathbf{y}_l - \bar{\mathbf{y}}_l)$ are compressed by an entropy-constrained product vector quantizer $Q_l^{1:b_l^2}$ with the conditional vector prior $P_l(\cdot|\bar{\mathbf{y}}_l)$. Lin1 is a linear projection layer that changes the tensor shape from $H_l \times W_l \times C_l$ to $H_l \times W_l \times k_l$.

Given the dequantized residual \hat{r}_l , the reconstructed vector representations are obtained by $\hat{y}_l = \text{Lin2}(\hat{r}_l) + \bar{y}_l$ with a shape of $H_l \times W_l \times C_l$, where Lin2 is a linear projection layer. Then \hat{y}_l is transformed to $\bar{y}_{l+1} = f_l(\hat{y}_l)$ and fed into the next quantization layer $l + 1$.

4.1. Nonlinear Vector Transform

Vector transform (VT) is not simple nonlinear transform. Although VQ is always superior to scalar quantization, the high complexity hinders its practical application (such as the huge number of codewords). Therefore, the method of VT is designed to cooperate with VQ, which ensures the complexity is acceptable. In our initial experiments, we found that directly using the linear vector transform [29–31] led to spatial aliasing and a significant performance drop. We identified two primary reasons: 1) linear transform (e.g., block-DCT) cannot handle nonlinear inter-vector correlations, and 2) intra-vector transform is crucial for obtaining effective vector representations. Therefore, we propose a novel vector transform unit comprising two components: intra transform for nonlinear representation learning and inter transform for inter-vector decorrelation. As shown in the right of Figure 4, the intra transform consists of two fully-connected (FC) layers along the channel axis and one activation function. The inter transform contains a Depth-wise Block FC layer that can be regarded as a block-DCT combined with channel-wise learnable transform matrices.

4.2. Conditional Entropy Model

For a single-stage VQ of dimension k with codebook size N , the entropy model P of the codeword index i is a discrete probability distribution $P(i) = p_i$, where $p_i > 0$ and $\sum_{i=1}^N p_i = 1$. Following previous work [2, 5], we parameterize p_i with the Softmax function and unnormalized logits $\mathbf{w} = (w_1, w_2, \dots, w_N)$:

$$p_i = \frac{e^{-w_i}}{\sum_{j=1}^N e^{-w_j}}. \quad (1)$$

For a multi-stage product VQ like NVTC, we propose a conditional entropy model (CEM) to further remove the inter-vector redundancy. Specifically, at layer l , CEM estimates the spatial-adaptive logits $\mathbf{w}_l \in \mathbb{R}^{H_l \times W_l \times N_l}$ by using the output $\bar{y}_l \in \mathbb{R}^{H_l \times W_l \times C_l}$ of layer $l - 1$ as:

$$\mathbf{w}_l = \Phi(\bar{y}_l), \quad (2)$$

where Φ is a neural network described in the supplementary material. Then the conditional probability $P_l(\cdot|\bar{y}_l)$ is calculated by the Softmax function.

4.3. RD optimization with Latent-Space ECVQ

Due to the gradient of VQ being almost zero everywhere, it is hard to end-to-end optimize both vector quantization and nonlinear transform for the rate-distortion trade-off.

Some recent works for generative modeling like VQ-VAE [16, 40] use the straight-through estimator (STE) to pass gradients from the decoder transform to the encoder transform, and introduce additional losses to optimize the codebook. Recent works [2, 43] for image compression utilize Softmax or Gumbel Softmax to approximate the VQ gradients. However, the joint RD trade-off is rarely considered in these works. For example, the rate-distortion trade-off in [43] is controlled by varying codebook size, where a larger codebook size corresponds to a higher bitrate. In this paper, similar to VQVAE [40], we also use STE to pass gradients through VQ. Then we propose a new rate-distortion loss to optimize the methods with latent-space ECVQ. The single-layer loss is as follows:

$$\mathcal{L} = \lambda \mathbb{E}_{\mathbf{x}}[d(\mathbf{x}, \hat{\mathbf{x}})] + \mathbb{E}_{\mathbf{x}}[-\log P(I)] + \beta \mathbb{E}_{\mathbf{x}}[d_1(\mathbf{y}, \hat{\mathbf{y}})] \quad (3)$$

where \mathbf{y} is the latent to be quantized, I is the codeword index from ECVQ, d is the pixel-space distortion metric, and d_1 is the latent-space VQ distortion metric. λ controls the rate-distortion trade-off. β controls the trade-off between d and d_1 . The encoding process of ECVQ has been described in Section 2. For multi-layer NVTC, the loss is as follows:

$$\mathcal{L} = \lambda \mathbb{E}_{\mathbf{x}}[d(\mathbf{x}, \hat{\mathbf{x}})] + \sum_l \mathbb{E}_{\mathbf{x}}[-\log P_l(I_l|\bar{y}_l)] + \beta \sum_l \mathbb{E}_{\mathbf{x}}[d_1(r_l, \hat{r}_l)] \quad (4)$$

Algorithm 1 Progressive Initialization of NVTC

Require: Input \mathbf{x} ; The vector quantizers $\{Q_l\}_{l=1}^L$; Encoder transform $f_{1:L}$; Decoder Transform $g_{1:L}$; Iteration milestones $T_{1:L}$, ($T_i \leq T_j$ when $i < j$).

- 1: **for** $t \leftarrow 1$ to T_{init} **do**
 - 2: $\bar{y}_1 \leftarrow \mathbf{0}$; transform \mathbf{x} into $\mathbf{y}_{1:L}$ using $f_{1:L}$
 - 3: **for** $l \leftarrow 1$ to L **do**
 - 4: **if** $t > T_l$ **then**
 - 5: $\hat{y}_l \leftarrow Q_l(\mathbf{y}_l - \bar{y}_l) + \bar{y}_l$
 - 6: **else**
 - 7: $\hat{y}_l \leftarrow \bar{y}_l$
 - 8: **end if**
 - 9: $\bar{y}_{l+1} \leftarrow g_l(\hat{y}_l)$
 - 10: **end for**
 - 11: $\hat{\mathbf{x}} \leftarrow \bar{y}_{L+1}$; update model using Equation 4.
 - 12: **if** $t \bmod \delta T = 0$ **then**
 - 13: reinitialize low-frequency codewords.
 - 14: **end if**
 - 15: **end for**
-

In practice, the initialization of VQ codebooks has a significant impact on training efficiency [4, 24]. We find this

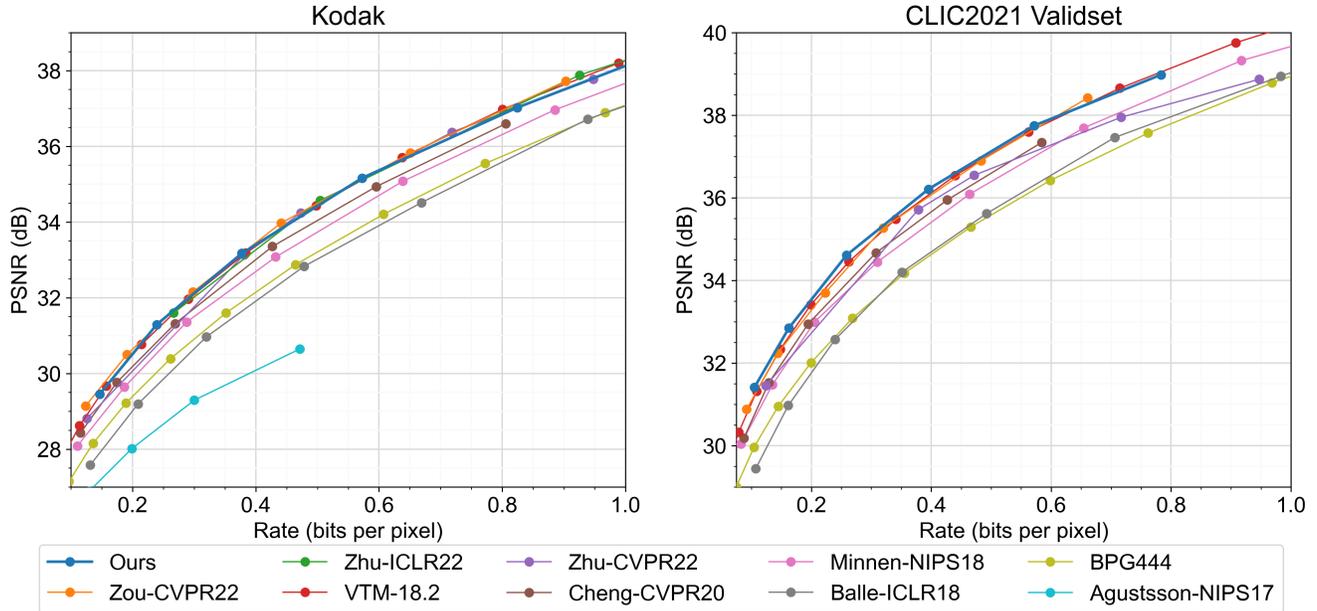


Figure 5. Rate-distortion results on kodak (left) and CLIC2021 valid set (right). The proposed method achieves state-of-the-art compression performance.

issue becomes more serious in the training of multi-stage VQ. Therefore, instead of random initialization, we propose a progressive initialization method shown in Algorithm 1. For simplicity, we omit the notation of product VQ, linear projection and Downscale/Upscale layers. During the initialization process, ECVQ is replaced with VQ, and the conditional entropy model is removed. The vector quantizers are initialized sequentially. To make full use of every codeword, we randomly reinitialize low-frequency codewords using high-frequency codewords for every δT iterations.

5. Experiments

5.1. Experiment Setup

Datasets. The image compression models were trained on COCO [32] train2017 set, which contains 118,287 images. Training Images are randomly cropped into 256×256 patches. For evaluation, we use Kodak dataset [27] which contains 24 images with 768×512 pixels. We also evaluate our models on CLIC2021 validset [1] with 41 high-resolution images with a size of about 2000×1000 . Details regarding the toy sources are provided in the supplementary.

Metric. We measure the quality of reconstructed images using the peak signal-to-noise ratio (PSNR) in RGB colorspace. The rate is measured by bits per pixel (bpp) for evaluation and bits per dimension (bpd) for training. The distortion metrics d and d_1 in Equation 4 are the mean squared error (MSE). We also use BD-rate/BD-PSNR [9] to calculate the average rate-distortion performance gain.

Implementation details. Our image codec consists of 16 quantization layers, where $(L_1, L_2, L_3) =$

$(6, 6, 4)$. The detailed configurations are described in the supplementary. We train six models with $\lambda \in \{64, 128, 256, 512, 1024, 2048\}$ for different RD points. β is set equivalent to λ . We initialize the models using Algorithm 1 with 0.6M iterations and then optimize the full models with 1.7M iterations. We use the Adam optimizer [26] with a batch size of 8. The initial learning rate is set as 10^{-4} , which is reduced to 10^{-5} for the final 0.2M iterations.

5.2. Performance

RD curves. We compare our method with two traditional coding methods: BPG and VVC [10] (VTM-18.2), and six learning-based methods including Zou-CVPR22 [45], Zhu-ICLR22 [44], Zhu-CVPR22 [43], Cheng-CVPR20 [12], Balle-ICLR18 [7] and Agustsson-NIPS17 [2]. Among them, Zhu-CVPR22 [43] and Agustsson-NIPS17 [2] are two latent-space VQ methods. As shown in Figure 5, our method is on par with previous methods on Kodak but surpasses them on CLIC 2021 validset. Even if our method is trained on the COCO dataset with a resolution of about 600×400 , the generalization performance on the high-resolution CLIC dataset is superior to the previous latent-space VQ method [43].

BD-rate vs. model complexity. In Figure 1, we compare the BD-rate, decoding time, and model parameters of different methods on CLIC 2021 validset. The BD-rate is computed over VTM 18.2, and the reported time is the average decoding time per image across the entire dataset. For a fair comparison, we evaluate the model complexity at high rate points (larger than 0.5 bpp on CLIC) except for Zhu-

		Time (ms)				
		NVT	CEM	VQ	EC	Total
Kodak	Encode	48	30	12	24	114
	Decode	23	13	4	28	68
CLIC	Encode	270	121	63	89	543
	Decode	154	38	14	101	307

Table 3. Encoding & decoding time. NVT refers to the nonlinear vector transform including VT Units and Downscale/Upscale layers. CEM is the conditional entropy model. EC refers to entropy encoding/decoding with one CPU core. GPU is GTX 1080Ti.

CVPR22¹ [43] and Guo-TCSVT21² [20]. All methods run with one GTX 1080Ti GPU and one CPU core, based on their official code or public reimplementations [8]. It can be observed that our method has the best rate-distortion-complexity performance. **More importantly**, compared to the recent VQ-based approach Zhu-CVPR2022 [43] that relies on computationally expensive scalar transform and unconstrained VQ, our method has a significantly reduced model size (about 5× smaller), faster decoding speed and better RD performance. The BD-rate of our method against Zhu-CVPR2022 is -10.9% on CLIC.

In Table 3, we provide detailed results regarding the encoding/decoding time of our method. It is worth noting that the encoding time is approximately twice the decoding time, which mainly stems from two factors: 1) the encoder needs to run with most of the VT units, where the decoder only has to execute the decoder VT units; 2) The computational gap between ECVQ encoding and decoding. We provide a more detailed explanation in the supplementary.

5.3. Ablation Study and Analysis

	BD-rate	Decoding time (ms)	Model parameters (M)
Full model	0.0%	307	12.8
A1: w/o VT	7.3%	405	49.6
A2: w/o ECVQ	23.9%	268	11.5
A3: w/o init	17.1%	307	12.8

Table 4. Ablation study on CLIC 2021 validset. BD-rate is computed over the full models (smaller is better).

Vector transform. To investigate the impact of the proposed VT, we replace the VT unit with the convolutional resblock used in previous image coding methods [12, 43]. The result in Table 4 (“A1: w/o VT”) shows that VT provides a better decorrelation ability with less complexity.

Latent-space ECVQ. As mentioned before, compared to ECVQ, the unconstrained VQ is theoretically suboptimal

¹The authors only release one low-rate model. Therefore we report the estimated results based on the description in their paper.

²The results are provided by the authors.

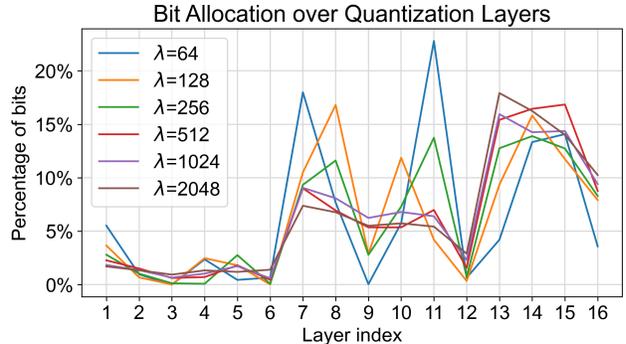


Figure 6. Layer-adaptive bit allocation results. We report the average percentage of bits of different quantization layers on Kodak.

in source space. To verify the effectiveness of ECVQ in latent space, we replace ECVQ with the unconstrained VQ in previous works [2, 43], and retrain our model. We manually adjust the codebook size to control the rate-distortion trade-off. As shown in Table 4 (“A2: w/o ECVQ”), unconstrained VQ brings a significant performance drop.

Progressive initialization. To evaluate the impact of the progressive initialization in Algorithm 1, we optimize the full model with randomly initialized parameters. As shown in Table 4 (“A3: w/o init”), random initialization causes a large performance drop. We find that only a part of codewords is activated for quantization, and the other codewords are ignored by ECVQ due to their extremely low probability. ECVQ with random initialization is more likely to make codewords “dead”, restricting its performance at high rate.

Layer-adaptive bit allocation. Note that our model consists of 16 quantization layers, which are joint optimized for the target RD trade-off. To investigate the bit allocation of our model, we show the percentage of bits of each layer in Figure 6. It can be observed that the model allocates more bits to high-resolution layers as λ increases.

For more results, analysis and visualization, please refer to the supplementary.

6. Conclusion

In this paper, we demonstrate the superiority of VQ over SQ with nonlinear transform based on empirical results on some toy sources. We propose a novel VQ-based coding scheme named nonlinear vector transform coding to make better use of VQ while keeping complexity low. Specifically, three technical points are proposed, including 1) a multi-stage quantization strategy, 2) nonlinear vector transforms and 3) RD optimization with entropy-constrained VQ in latent space. Our method and these analyses shed light in utilizing VQ for better neural network-based compression.

Acknowledgments

This work was supported in part by NSFC under Grant U1908209, 62021001.

References

- [1] Clic2021 challenge on learned image compression. <http://clic.compression.cc/2021/>, 2021. 1, 7
- [2] Eirikur Agustsson, Fabian Mentzer, Michael Tschannen, Lukas Cavigelli, Radu Timofte, Luca Benini, and Luc V Gool. Soft-to-hard vector quantization for end-to-end learning compressible representations. *Advances in neural information processing systems*, 30, 2017. 3, 4, 6, 7, 8
- [3] Eirikur Agustsson and Lucas Theis. Universally quantized neural compression. *Advances in neural information processing systems*, 33:12367–12376, 2020. 4
- [4] David Arthur and Sergei Vassilvitskii. k-means++: The advantages of careful seeding. Technical report, Stanford, 2006. 6
- [5] Johannes Ballé, Philip A Chou, David Minnen, Saurabh Singh, Nick Johnston, Eirikur Agustsson, Sung Jin Hwang, and George Toderici. Nonlinear transform coding. *IEEE Journal of Selected Topics in Signal Processing*, 15(2):339–353, 2020. 1, 3, 4, 5, 6
- [6] Johannes Ballé, Valero Laparra, and Eero P. Simoncelli. End-to-end optimized image compression. In *5th International Conference on Learning Representations, ICLR 2017*, 2017. 1
- [7] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*, 2018. 4, 5, 7
- [8] Jean Bégaint, Fabien Racapé, Simon Feltman, and Akshay Pushparaja. Compressai: a pytorch library and evaluation platform for end-to-end compression research. *arXiv preprint arXiv:2011.03029*, 2020. 8
- [9] Gisle Bjontegaard. Calculation of average psnr differences between rd-curves. *VCEG-M33*, 2001. 5, 7
- [10] Benjamin Bross, Ye-Kui Wang, Yan Ye, Shan Liu, Jianle Chen, Gary J Sullivan, and Jens-Rainer Ohm. Overview of the versatile video coding (vvc) standard and its applications. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(10):3736–3764, 2021. 7
- [11] Andrés Buzo, A Gray, RM Gray, and John Markel. Speech coding based upon vector quantization. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 28(5):562–574, 1980. 2
- [12] Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7939–7948, 2020. 1, 2, 4, 7, 8
- [13] Philip A Chou, Tom Lookabaugh, and Robert M Gray. Entropy-constrained vector quantization. *IEEE Transactions on acoustics, speech, and signal processing*, 37(1):31–42, 1989. 3, 4
- [14] Henry Cohn, Abhinav Kumar, Stephen Miller, Danylo Radchenko, and Maryna Viazovska. The sphere packing problem in dimension 24. *Annals of Mathematics*, 185(3):1017–1033, 2017. 2
- [15] John Conway and Neil Sloane. Voronoi regions of lattices, second moments of polytopes, and quantization. *IEEE transactions on information theory*, 28(2):211–226, 1982. 2, 5
- [16] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021. 1, 4, 6
- [17] Allen Gersho. Asymptotically optimal block quantization. *IEEE Transactions on information theory*, 25(4):373–380, 1979. 2
- [18] Allen Gersho and Robert M Gray. *Vector quantization and signal compression*. 1992. 1, 4
- [19] Vivek K Goyal. Theoretical foundations of transform coding. *IEEE Signal Processing Magazine*, 18(5):9–21, 2001. 4
- [20] Zongyu Guo, Zhizheng Zhang, Runsen Feng, and Zhibo Chen. Causal contextual prediction for learned image compression. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(4):2329–2341, 2021. 1, 2, 4, 8
- [21] Thomas C Hales. The sphere packing problem. *Journal of Computational and Applied Mathematics*, 44(1):41–76, 1992. 2
- [22] Herve Jegou, Matthijs Douze, and Cordelia Schmid. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33(1):117–128, 2010. 2
- [23] Biing-Hwang Juang and A Gray. Multiple stage vector quantization for speech coding. In *ICASSP’82. IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 7, pages 597–600. IEEE, 1982. 2
- [24] Ioannis Katsavounidis, C-C Jay Kuo, and Zhen Zhang. A new initialization technique for generalized lloyd iteration. *IEEE Signal processing letters*, 1(10):144–146, 1994. 6
- [25] Robert A King and Nasser M Nasrabadi. Image coding using vector quantization in the transform domain. *Pattern Recognition Letters*, 1(5-6):323–329, 1983. 3, 4
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7
- [27] Eastman Kodak. Kodak Lossless True Color Image Suite (PhotoCD PCD0992). <http://r0k.us/graphics/kodak/>, 1993. 7
- [28] Weiping Li. Vector transform and image coding. *IEEE Transactions on Circuits and Systems for Video Technology*, 1(4):297–307, 1991. 3, 4
- [29] Weiping Li. On vector transformation. *IEEE Transactions on Signal Processing*, 41(11):3114–3126, 1993. 6
- [30] Weiping Li and Y-Q Zhang. New insights and results on transform domain vq of images. In *1993 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 5, pages 609–612. IEEE, 1993. 3, 4, 6
- [31] Weiping Li and Ya-Qin Zhang. Vector-based signal processing and quantization for image and video compression. *Proceedings of the IEEE*, 83(2):317–335, 1995. 3, 4, 6
- [32] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 7

- [33] Tom D Lookabaugh and Robert M Gray. High-resolution quantization theory and the vector quantizer advantage. *IEEE Transactions on Information Theory*, 35(5):1020–1033, 1989. [1](#), [2](#), [4](#), [5](#)
- [34] David Minnen, Johannes Ballé, and George D Toderici. Joint autoregressive and hierarchical priors for learned image compression. *Advances in neural information processing systems*, 31, 2018. [1](#), [4](#)
- [35] David Minnen and Saurabh Singh. Channel-wise autoregressive entropy models for learned image compression. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 3339–3343. IEEE, 2020. [1](#), [2](#), [4](#)
- [36] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021. [1](#), [4](#)
- [37] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32, 2019. [5](#)
- [38] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. [1](#)
- [39] ML Sabin and R Gray. Product code vector quantizers for waveform and voice coding. *IEEE transactions on acoustics, speech, and signal processing*, 32(3):474–488, 1984. [2](#)
- [40] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. [1](#), [3](#), [4](#), [6](#)
- [41] Maryna S Viazovska. The sphere packing problem in dimension 8. *Annals of Mathematics*, pages 991–1015, 2017. [2](#)
- [42] Yibo Yang, Robert Bamler, and Stephan Mandt. Improving inference for neural image compression. *Advances in Neural Information Processing Systems*, 33:573–584, 2020. [4](#)
- [43] Xiaosu Zhu, Jingkuan Song, Lianli Gao, Feng Zheng, and Heng Tao Shen. Unified multivariate gaussian mixture for efficient neural image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17612–17621, 2022. [3](#), [4](#), [6](#), [7](#), [8](#)
- [44] Yin hao Zhu, Yang Yang, and Taco Cohen. Transformer-based transform coding. In *International Conference on Learning Representations*, 2021. [2](#), [4](#), [7](#)
- [45] Renjie Zou, Chunfeng Song, and Zhaoxiang Zhang. The devil is in the details: Window-based attention for image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17492–17501, 2022. [2](#), [4](#), [7](#)