

# Learning Scalable $\ell_\infty$ -constrained Near-lossless Image Compression via Joint Lossy Image and Residual Compression

Yuanchao Bai<sup>1</sup>, Xianming Liu<sup>1,2,\*</sup>, Wangmeng Zuo<sup>1,2</sup>, Yaowei Wang<sup>1</sup>, Xiangyang Ji<sup>3</sup>

<sup>1</sup>Peng Cheng Laboratory, <sup>2</sup>Harbin Institute of Technology, <sup>3</sup>Tsinghua University  
{baiych,wangyw}@pcl.ac.cn, {csxm,wmzuo}@hit.edu.cn, xyji@tsinghua.edu.cn

## Abstract

We propose a novel joint lossy image and residual compression framework for learning  $\ell_\infty$ -constrained near-lossless image compression. Specifically, we obtain a lossy reconstruction of the raw image through lossy image compression and uniformly quantize the corresponding residual to satisfy a given tight  $\ell_\infty$  error bound. Suppose that the error bound is zero, i.e., lossless image compression, we formulate the joint optimization problem of compressing both the lossy image and the original residual in terms of variational auto-encoders and solve it with end-to-end training. To achieve scalable compression with the error bound larger than zero, we derive the probability model of the quantized residual by quantizing the learned probability model of the original residual, instead of training multiple networks. We further correct the bias of the derived probability model caused by the context mismatch between training and inference. Finally, the quantized residual is encoded according to the bias-corrected probability model and is concatenated with the bitstream of the compressed lossy image. Experimental results demonstrate that our near-lossless codec achieves the state-of-the-art performance for lossless and near-lossless image compression, and achieves competitive PSNR while much smaller  $\ell_\infty$  error compared with lossy image codecs at high bit rates.

## 1. Introduction

Image compression is a ubiquitous technique in computer vision. For certain applications with stringent demands on image fidelity, such as medical imaging or image archiving, the most reliable choice is lossless image compression. However, the compression ratio of lossless compression is upper-bounded by Shannon’s source coding theorem [36], and is typically between 2:1 and 3:1 for practical lossless image codecs [50, 47, 37, 14, 5, 38]. To improve the compression performance while keeping the reliability

of the decoded images,  $\ell_\infty$ -constrained near-lossless image compression is developed [6, 15, 2, 47, 51] and standardized in traditional codecs, e.g., JPEG-LS [47] and CALIC [51]. Different from lossy image compression with Peak Signal-to-Noise Ratio (PSNR) or Multi-Scale Structural SIMilarity index (MS-SSIM) [45, 46] distortion measures,  $\ell_\infty$ -constrained near-lossless image compression requires the maximum reconstruction error of each pixel to be no larger than a given tight numerical bound. Lossless image compression is a special case of near-lossless image compression, when the tight error bound is zero.

With the fast development of deep neural networks (DNNs), learning-based lossy image compression [40, 3, 39, 41, 34, 24, 4, 31, 27, 22, 8, 7, 23, 25, 29] has achieved tremendous progress over the last four years. Most recent methods for lossy image compression adopt variational auto-encoder (VAE) architecture [17, 18] based on transform coding [11], where the rate is modeled by the entropy of quantized latent variables and the reconstruction distortion is measured by PSNR or MS-SSIM. Through end-to-end rate-distortion optimization, the state-of-the-art methods, such as [7], can achieve comparable performance with the latest compression standard Versatile Video Coding (VVC) [32]. However, the above transform coding scheme cannot be directly employed on near-lossless image compression, because it is difficult for DNN-based transforms to satisfy a tight bound on the maximum reconstruction error of each pixel, even without quantization. Invertible transforms, such as integer discrete flow [13] or wavelet-like transform [25], are possible solutions to lossless compression but not to general near-lossless compression.

In this paper, we propose a new joint lossy image and residual compression framework for learning near-lossless image compression, inspired by the traditional “lossy plus residual” coding scheme [10, 26, 2]. Specifically, we obtain a lossy reconstruction of the raw image through lossy image compression and uniformly quantize the corresponding residual to satisfy the given  $\ell_\infty$  error bound  $\tau$ . Suppose that the error bound  $\tau$  is zero, i.e., lossless image compression, we formulate the joint optimization problem of compress-

\*Corresponding author

ing both the lossy image and the original residual in terms of VAEs [17, 18], and solve it with end-to-end training. Note that our VAE model is novel, different from transform coding based VAEs [3, 39, 4] for simply lossy image compression or bits-back coding based VAEs [42, 19] for lossless image compression.

To achieve scalable near-lossless compression with error bound  $\tau > 0$ , we derive the probability model of the quantized residual by quantizing the learned probability model of the original residual at  $\tau = 0$ , instead of training multiple networks. Because residual quantization leads to the context mismatch between training and inference, we further propose a bias correction scheme to correct the bias of the derived probability model. An arithmetic coder [48] is adopted to encode the quantized residual according to the bias-corrected probability model. Finally, the near-lossless compressed image is stored including the bitstreams of the encoded lossy image and the quantized residual.

Our main contributions are summarized as follows:

- We propose a joint lossy image and residual compression framework to realize learning-based lossless and near-lossless image compression. The framework is interpreted as a VAE model and can be end-to-end optimized.
- We realize scalable compression by deriving the probability model of the quantized residual from the learned probability model of the original residual, instead of training multiple networks. A bias correction scheme further improves the compression performance.
- Our codec achieves the state-of-the-art performance for lossless and near-lossless image compression, and achieves competitive PSNR while much smaller  $\ell_\infty$  error compared with lossy image codecs at high bit rates.

## 2. Related Work

**Learning-based Lossy Image Compression.** Early learning-based methods [40, 41] for lossy image compression are based on recurrent neural networks (RNNs). Following [3, 39], most recent learned methods [34, 24, 4, 31, 27, 22, 8, 7, 23, 25, 29] adopt convolutional neural networks (CNNs) and can be interpreted as VAEs [17, 18] based on transform coding [11]. In our joint lossy image and residual compression framework, we take advantage of the advanced transform (network structures), quantization and entropy coding techniques in the existing learning-based methods for lossy image compression.

**Learning-based Lossless Image Compression.** Given the strong connections between lossless compression and unsupervised learning, auto-aggressive models [33, 43, 35], flow models [13], bits-back coding based VAEs [42, 19] and other specific models [28, 30, 25], are introduced to approximate the true distribution of raw images for entropy coding. Previous work [30] uses traditional BPG lossy image codec [5] to compress raw images and proposes a CNN to

further compress residuals, which is a special case of our framework. Beyond [30], our lossy image compressor and residual compressor are jointly optimized through end-to-end training and are interpretable as a VAE model.

**Near-lossless Image Compression.** Basically, previous methods for near-lossless image compression can be divided into three categories: 1) *pre-quantization*: adjusting raw pixel values to the  $\ell_\infty$  error bound, and then compressing the pre-processed images with lossless image compression, e.g., near-lossless WebP [14]; 2) *predictive coding*: predicting subsequent pixels based on previously encoded pixels, then quantizing prediction residuals to satisfy the  $\ell_\infty$  error bound, and finally compressing the quantized residuals, e.g., [6, 15], near-lossless JPEG-LS [47] and near-lossless CALIC [51]; 3) *lossy plus residual coding*: similar to 2), but replacing predictive coder with lossy image coder, and both the lossy image and the quantized residual are encoded, e.g., in [2].

In this paper, we propose a joint lossy image and residual compression framework for learning-based near-lossless image compression, inspired by lossy plus residual coding. *To the best of our knowledge, we propose the first deep-learning-based near-lossless image codec.* Recently, CNN-based soft-decoding methods [52, 53] are proposed to improve the reconstruction performance of near-lossless CALIC [51]. However, these methods should belong to image reconstruction rather than image compression.

## 3. Scalable Near-lossless Image Compression

### 3.1. Overview of Compression Framework

Given a tight  $\ell_\infty$  bound  $\tau \in \{0, 1, 2, \dots\}$ , near-lossless codecs compress a raw image satisfying the following distortion measure:

$$D_{\text{null}}(\mathbf{x}, \hat{\mathbf{x}}) = \|\mathbf{x} - \hat{\mathbf{x}}\|_\infty = \max_{i,c} |x_{i,c} - \hat{x}_{i,c}| \leq \tau \quad (1)$$

where  $\mathbf{x}$  and  $\hat{\mathbf{x}}$  are a raw image and its near-lossless reconstructed counterpart.  $x_{i,c}$  and  $\hat{x}_{i,c}$  are the pixels of  $\mathbf{x}$  and  $\hat{\mathbf{x}}$ , respectively.  $i$  denotes the  $i$ -th spatial position in 2D raster-scan order, and  $c$  denotes the  $c$ -th channel.

In order to realize the  $\ell_\infty$  error bound (1), we propose a near-lossless image compression framework by integrating lossy image compression and residual compression. We first obtain a lossy reconstruction  $\tilde{\mathbf{x}}$  of the raw image  $\mathbf{x}$  through lossy image compression. Lossy image compression methods, such as traditional [44, 37, 14, 5] or learned methods [24, 4, 31, 27, 22, 8, 7, 23, 25], can achieve high PSNR results at relatively low bit rates, but it is difficult for these methods to ensure a tight error bound  $\tau$  of each pixel in  $\tilde{\mathbf{x}}$ . We then compute the residual  $\mathbf{r} = \mathbf{x} - \tilde{\mathbf{x}}$  and suppose that  $\mathbf{r}$  is quantized to  $\hat{\mathbf{r}}$ . Let  $\hat{\mathbf{x}} = \tilde{\mathbf{x}} + \hat{\mathbf{r}}$ , the reconstruction error  $\mathbf{x} - \hat{\mathbf{x}}$  is equivalent to the quantization error  $\mathbf{r} - \hat{\mathbf{r}}$  of  $\mathbf{r}$ .

Thus, we adopt a uniform residual quantizer whose bin size is  $2\tau + 1$  and quantized value is [47, 51]:

$$\hat{r}_{i,c} = \text{sgn}(r_{i,c})(2\tau + 1)\lfloor (|r_{i,c}| + \tau)/(2\tau + 1) \rfloor \quad (2)$$

where  $\text{sgn}(\cdot)$  denotes the sign function.  $r_{i,c}$  and  $\hat{r}_{i,c}$  are the elements of  $\mathbf{r}$  and  $\hat{\mathbf{r}}$ , respectively. With (2), we now have  $|r_{i,c} - \hat{r}_{i,c}| \leq \tau$  for each  $\hat{r}_{i,c}$  in  $\hat{\mathbf{r}}$ , satisfying the tight error bound (1). Finally, we encode  $\hat{\mathbf{r}}$  and concatenate it with the compressed  $\tilde{\mathbf{x}}$ , and send them to the decoder.

To compress both the lossy image  $\tilde{\mathbf{x}}$  and the quantized residual  $\hat{\mathbf{r}}$  effectively leads to a challenging joint optimization problem. In the following subsections, we first propose an end-to-end trainable VAE model [17, 18] to solve the joint optimization problem with  $\tau = 0$ . Then, we propose a scalable compression scheme with bias correction for  $\tau > 0$  based on the learned VAE model.

### 3.2. Joint Lossy Image & Residual Compression

Before addressing the joint compression of the lossy image  $\tilde{\mathbf{x}}$  and the quantized residual  $\hat{\mathbf{r}}$  with variable  $\tau$ , we first solve the special case of near-lossless image compression with  $\tau = 0$ , *i.e.*, lossless image compression.

**Problem Formulation.** Assuming that raw images are sampled from an unknown probability distribution  $p(\mathbf{x})$ , the compression performance of lossless image compression depends on how well we can approximate  $p(\mathbf{x})$  with an underlying model  $p_\theta(\mathbf{x})$ . We adopt the latent variable model which is formulated by a marginal distribution  $p_\theta(\mathbf{x}) = \int p_\theta(\mathbf{x}, \mathbf{y}) d\mathbf{y}$ .  $\mathbf{y}$  is an unobserved latent variable and  $\theta$  denote the parameters of this model.

Since directly learning the marginal distribution  $p_\theta(\mathbf{x})$  is typically intractable, one alternative way is to optimize the evidence lower bound (ELBO) via VAEs [17, 18]. By introducing an inference model  $q_\phi(\mathbf{y}|\mathbf{x})$  to approximate the posterior  $p_\theta(\mathbf{y}|\mathbf{x})$ , we can rewrite the logarithm of the marginal likelihood  $p_\theta(\mathbf{x})$  as:

$$\log p_\theta(\mathbf{x}) = \underbrace{\mathbb{E}_{q_\phi(\mathbf{y}|\mathbf{x})} \log \frac{p_\theta(\mathbf{x}, \mathbf{y})}{q_\phi(\mathbf{y}|\mathbf{x})}}_{ELBO} + \underbrace{\mathbb{E}_{q_\phi(\mathbf{y}|\mathbf{x})} \log \frac{q_\phi(\mathbf{y}|\mathbf{x})}{p_\theta(\mathbf{y}|\mathbf{x})}}_{D_{kl}(q_\phi(\mathbf{y}|\mathbf{x})||p_\theta(\mathbf{y}|\mathbf{x}))} \quad (3)$$

where  $D_{kl}(\cdot||\cdot)$  is the Kullback-Leibler (KL) divergence.  $\phi$  denote the parameters of the inference model  $q_\phi(\mathbf{y}|\mathbf{x})$ . Since  $D_{kl}(q_\phi(\mathbf{y}|\mathbf{x})||p_\theta(\mathbf{y}|\mathbf{x})) \geq 0$  and  $\log p_\theta(\mathbf{x}) \leq 0$ , ELBO is the lower bound of  $\log p_\theta(\mathbf{x})$ .

From (3), we can minimize the expectation of negative ELBO as a proxy for the expected codelength  $\mathbb{E}_{p(\mathbf{x})}[-\log p_\theta(\mathbf{x})]$ . In our compression framework, we first adopt lossy image compression based on transform coding [11], and the expectation of negative ELBO can be reformulated as follows:

$$\mathbb{E}_{p(\mathbf{x})} \mathbb{E}_{q_\phi(\hat{\mathbf{y}}|\mathbf{x})} [-\log p_\theta(\mathbf{x}|\hat{\mathbf{y}}) - \log p_\theta(\hat{\mathbf{y}})] \quad (4)$$

where  $\hat{\mathbf{y}}$  is quantized from continuous  $\mathbf{y}$  and  $\mathbf{y}$  is transformed from  $\mathbf{x}$ . Like [4], we relax the quantization of  $\mathbf{y}$  by adding noise from  $\mathcal{U}(-\frac{1}{2}, \frac{1}{2})$ , and assume  $q_\phi(\hat{\mathbf{y}}|\mathbf{x}) = \prod_{i,c} \mathcal{U}(y_{i,c} - \frac{1}{2}, y_{i,c} + \frac{1}{2})$ . Thus,  $\log q_\phi(\hat{\mathbf{y}}|\mathbf{x}) = 0$  is dropped from (4). For simply lossy image compression, such as [4, 31, 8, 7], the first and second terms of (4) are considered as the distortion loss and the rate loss, respectively. Only  $\hat{\mathbf{y}}$  needs to be encoded.

Beyond lossy image compression, we further take residual compression into consideration. For each  $\mathbf{x}$  and all  $(\tilde{\mathbf{x}}, \mathbf{r})$  pairs satisfying  $\tilde{\mathbf{x}} + \mathbf{r} = \mathbf{x}$ , we have  $p_\theta(\mathbf{x}|\hat{\mathbf{y}}) = \sum_{\tilde{\mathbf{x}}+\mathbf{r}=\mathbf{x}} p_\theta(\tilde{\mathbf{x}}, \mathbf{r}|\hat{\mathbf{y}}) \geq p_\theta(\tilde{\mathbf{x}}, \mathbf{r}|\hat{\mathbf{y}})$ . Hence, we substitute  $p_\theta(\tilde{\mathbf{x}}, \mathbf{r}|\hat{\mathbf{y}})$  for  $p_\theta(\mathbf{x}|\hat{\mathbf{y}})$ , leading to the upper bound of (4). Let  $p_\theta(\tilde{\mathbf{x}}, \mathbf{r}|\hat{\mathbf{y}}) = p_\theta(\tilde{\mathbf{x}}|\hat{\mathbf{y}}) \cdot p_\theta(\mathbf{r}|\tilde{\mathbf{x}}, \hat{\mathbf{y}})$ , we have:

$$\mathbb{E}_{p(\mathbf{x})} \mathbb{E}_{q_\phi(\hat{\mathbf{y}}|\mathbf{x})} \left[ \underbrace{-\log p_\theta(\tilde{\mathbf{x}}|\hat{\mathbf{y}})}_0 - \underbrace{\log p_\theta(\mathbf{r}|\tilde{\mathbf{x}}, \hat{\mathbf{y}})}_{R_r} - \underbrace{\log p_\theta(\hat{\mathbf{y}})}_{R_{\hat{\mathbf{y}}}} \right] \quad (5)$$

The lossy reconstruction  $\tilde{\mathbf{x}}$  is computed by the inverse transform of  $\hat{\mathbf{y}}$  and quantization. As the quantization is relaxed by adding uniform noise from  $\mathcal{U}(-\frac{1}{2}, \frac{1}{2})$ , we have  $\log p_\theta(\tilde{\mathbf{x}}|\hat{\mathbf{y}}) = 0$ . The second term  $R_r$  and the third term  $R_{\hat{\mathbf{y}}}$  of (5) are the rates of encoding  $\mathbf{r}$  and  $\hat{\mathbf{y}}$ , respectively.

Notice that no distortion loss of lossy image compression is specified in (5). Therefore, we can embed arbitrary lossy image compressors and optimize (5) to achieve lossless image compression. A special case is the previous lossless image compression method [30], in which the BPG lossy image compressor [5] with a learned quantization parameter classifier optimizes  $-\log p_\theta(\hat{\mathbf{y}})$  and a CNN-based residual compressor optimizes  $-\log p_\theta(\mathbf{r}|\tilde{\mathbf{x}})$ .

**Network Design and Optimization.** To realize (5), we propose our network architecture in Fig. 1(a). For lossy image compression, we employ the hyper-prior model proposed in [4], where side information  $\hat{\mathbf{z}}$  is extracted to model the distribution of  $\hat{\mathbf{y}}$ . We employ the residual and attention blocks similar to [7, 12] in the *encoder* and *decoder*. Detailed structure of the lossy image compressor is described in the supplementary material.  $R_{\hat{\mathbf{y}}}$  is thus extended by

$$R_{\hat{\mathbf{y}}, \hat{\mathbf{z}}} = \mathbb{E}_{p(\mathbf{x})} \mathbb{E}_{q_\phi(\hat{\mathbf{y}}, \hat{\mathbf{z}}|\mathbf{x})} [-\log p_\theta(\hat{\mathbf{y}}|\hat{\mathbf{z}}) - \log p_\theta(\hat{\mathbf{z}})] \quad (6)$$

where  $R_{\hat{\mathbf{y}}, \hat{\mathbf{z}}}$  is the cost of encoding both  $\hat{\mathbf{y}}$  and  $\hat{\mathbf{z}}$ .

Denote by  $\mathbf{u} = g_u(\hat{\mathbf{y}})$ , where the feature  $\mathbf{u}$  is extracted from  $\hat{\mathbf{y}}$  by  $g_u(\cdot)$ . Denote by  $\tilde{\mathbf{x}} = g_{\tilde{\mathbf{x}}}(\hat{\mathbf{y}})$ , where the lossy reconstruction  $\tilde{\mathbf{x}}$  is inversely transformed from  $\hat{\mathbf{y}}$  by  $g_{\tilde{\mathbf{x}}}(\cdot)$ .  $g_u$  and  $g_{\tilde{\mathbf{x}}}$  share the *decoder*, and only the last convolutional layers are different, as shown in Fig. 1(a). We interpret  $\mathbf{u}$  as the feature of the residual  $\mathbf{r}$  given  $\tilde{\mathbf{x}}$  and  $\hat{\mathbf{y}}$  ( $\mathbf{u}$  shares the same height and width with  $\mathbf{r}$  and has 64 channels). As  $\mathbf{r} = \mathbf{x} - \tilde{\mathbf{x}}$ , we can compute  $\mathbf{r}$  and further consider the causal context  $C_r$  of  $\mathbf{r}$  during encoding.  $R_r$  is thus extended by

$$R_r = \mathbb{E}_{p(\mathbf{x})} \mathbb{E}_{q_\phi(\hat{\mathbf{y}}, \hat{\mathbf{z}}|\mathbf{x})} [-\log p_\theta(\mathbf{r}|\mathbf{u}, C_r)] \quad (7)$$

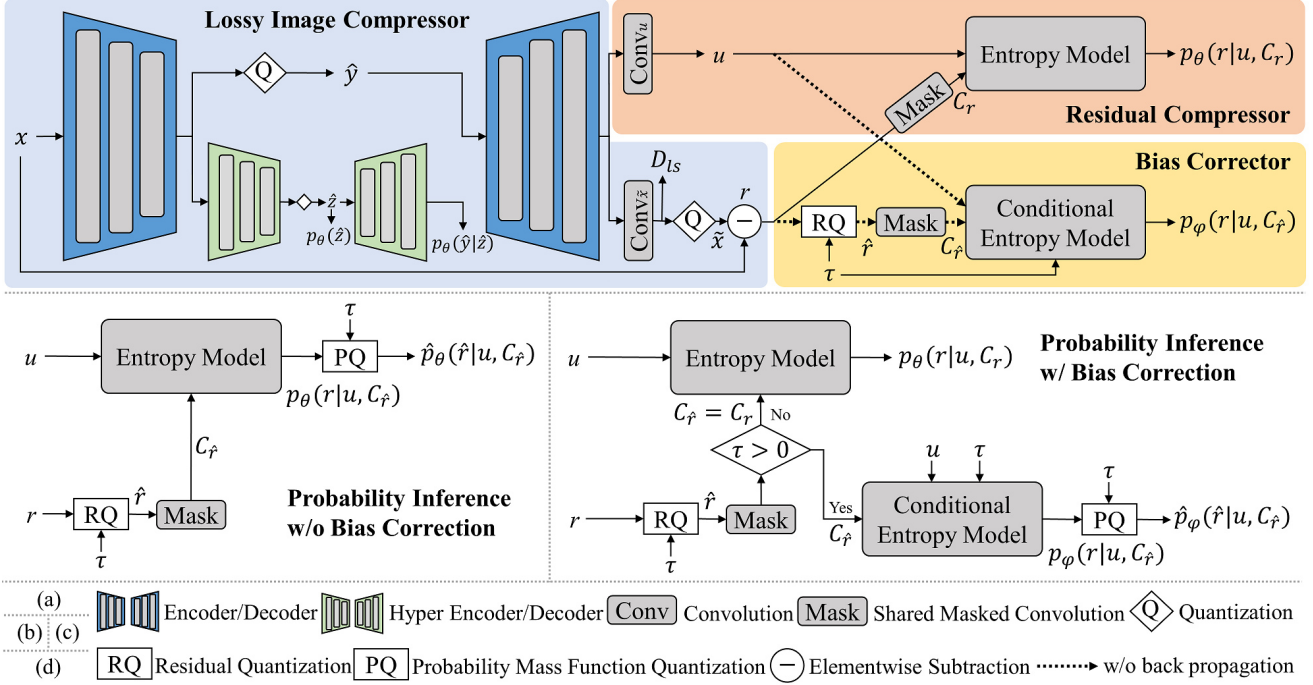


Figure 1. (a) Overview of network architecture. (b) Probability inference of the quantized residual  $\hat{r}$  without bias correction. (c) Probability inference of the quantized residual  $\hat{r}$  with bias correction. (d) Legend with descriptions for notations in (a)-(c).

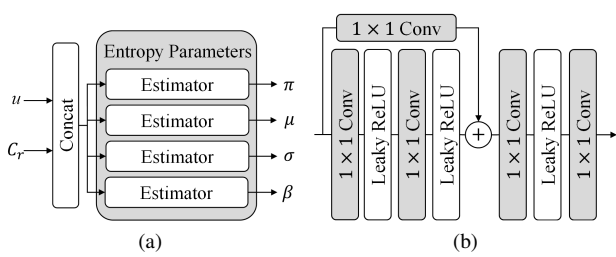


Figure 2. Entropy parameter estimation sub-network. Given  $u$  and  $C_r$ , (a) estimates parameters of discrete logistic mixture likelihoods corresponding to *entropy model* in Fig. 1(a). (b) is the detailed structure of *estimator* in (a). All  $1 \times 1$  convolutional layers except the last layer have 128 channels. The last convolutional layer has  $3 \cdot K$  channels.

where the causal context  $C_{r_i} \in C_r$  is shared by  $r_{i,c}$  of all channels.  $C_{r_i}$  is extracted from  $r_{<i,c}$  by a  $5 \times 5$  masked convolutional layer [31] with 64 channels.  $r_{<i,c}$  denote the residuals encoded before  $r_{i,c}$ . For RGB images with three channels, we have

$$p_\theta(\mathbf{r}|u, C_r) = \prod_i p_\theta(r_{i,1}, r_{i,2}, r_{i,3}|u_i, C_{r_i}) \quad (8)$$

We further model the probability mass function (PMF) of the residual  $\mathbf{r}$  with discrete logistic mixture likelihood [35, 28, 30]. We use a channel autoregression over  $r_{i,1}, r_{i,2}, r_{i,3}$

and reformulate  $p_\theta(r_{i,1}, r_{i,2}, r_{i,3}|u_i, C_{r_i})$  as

$$p_\theta(r_{i,1}, r_{i,2}, r_{i,3}|u_i, C_{r_i}) = p_\theta(r_{i,1}|u_i, C_{r_i}) \cdot p_\theta(r_{i,2}|r_{i,1}, u_i, C_{r_i}) \cdot p_\theta(r_{i,3}|r_{i,1}, r_{i,2}, u_i, C_{r_i}) \quad (9)$$

As shown in Fig. 2, we introduce a sub-network to estimate entropy parameters, including mixture weights  $\pi_{i,c}^k$ , means  $\mu_{i,c}^k$ , variances  $\sigma_{i,c}^k$  and mixture coefficients  $\beta_{i,t}$ . We choose a mixture of  $K = 5$  logistic distributions.  $k$  denotes the index of the  $k$ -th logistic distribution.  $t$  denotes the channel index of  $\beta$ . The channel autoregression over  $r_{i,1}, r_{i,2}, r_{i,3}$  is implemented by updating the means using:

$$\begin{aligned} \tilde{\mu}_{i,1}^k &= \mu_{i,1}^k, & \tilde{\mu}_{i,2}^k &= \mu_{i,2}^k + \beta_{i,1} \cdot r_{i,1}, \\ \tilde{\mu}_{i,3}^k &= \mu_{i,3}^k + \beta_{i,2} \cdot r_{i,1} + \beta_{i,3} \cdot r_{i,2} \end{aligned} \quad (10)$$

With  $\pi_{i,c}^k$ ,  $\tilde{\mu}_{i,c}^k$  and  $\sigma_{i,c}^k$ , we have

$$p_\theta(r_{i,c}|r_{i,<c}, u_i, C_{r_i}) \sim \sum_{k=1}^K \pi_{i,c}^k \text{logistic}(\tilde{\mu}_{i,c}^k, \sigma_{i,c}^k) \quad (11)$$

where  $\text{logistic}(\cdot)$  is the logistic distribution. For discrete  $r_{i,c}$ , we evaluate  $p_\theta(r_{i,c}|r_{i,<c}, u_i, C_{r_i})$  as [35, 28, 30]:

$$\sum_{k=1}^K \pi_{i,c}^k \left[ S\left(\frac{r_{i,c}^+ - \tilde{\mu}_{i,c}^k}{\sigma_{i,c}^k}\right) - S\left(\frac{r_{i,c}^- - \tilde{\mu}_{i,c}^k}{\sigma_{i,c}^k}\right) \right] \quad (12)$$

where  $S(\cdot)$  denotes the sigmoid function.  $r_{i,c}^+ = r_{i,c} + 0.5$  and  $r_{i,c}^- = r_{i,c} - 0.5$ .






Figure 3. PMF quantization corresponding to residual quantization (2) with  $\tau = 1$ . Each red number is the value of the quantized residual  $\hat{r}_{i,c}$ . The probability of each quantized value is the sum of the probabilities of the values in the same bin.

Besides rate terms  $R_{\hat{y},\hat{z}}$  and  $R_{\mathbf{r}}$ , we add a distortion term  $D_{ls}(\mathbf{x}, \tilde{\mathbf{x}}) = \mathbb{E}_{p(\mathbf{x})} \mathbb{E}_{i,c} (x_{i,c} - \tilde{x}_{i,c})^2$  to minimize the mean square error (MSE) between raw image  $\mathbf{x}$  and lossy reconstruction  $\tilde{\mathbf{x}}$ . As discussed in [4], minimizing MSE loss is equivalent to learning a lossy image compressor that fits residual  $\mathbf{r}$  to a zero-mean factorized Gaussian distribution. However, the discrepancy between the real distribution of  $\mathbf{r}$  and the Gaussian distribution could be large. Thus, a more sophisticated entropy model is proposed to encode  $\mathbf{r}$ . The full loss function for learning near-lossless image compression with  $\tau = 0$ , *i.e.*, lossless image compression, is

$$\mathcal{L}(\theta, \phi) = R_{\hat{y},\hat{z}} + R_{\mathbf{r}} + \lambda \cdot D_{ls} \quad (13)$$

where  $\lambda$  controls the “rate-distortion” trade-off. When  $\lambda = 0$ ,  $\tilde{\mathbf{x}}$  becomes a latent variable without any constraints. It leads to the best lossless compression performance but is not suitable for our near-lossless image compression with  $\tau > 0$ . We choose  $\lambda = 0.03$  empirically in our codec. Experiments with different  $\lambda$ 's are conducted in Sec. 4.4.

### 3.3. Probability Inference of Quantized Residuals

We next propose a scheme to achieve *scalable* near-lossless image compression with  $\tau > 0$ . For variable  $\ell_\infty$  error bound  $\tau$ , we keep the lossy reconstruction  $\tilde{\mathbf{x}}$  fixed and quantize the residual  $\mathbf{r}$  to  $\hat{\mathbf{r}}$  by (2). Instead of training multiple networks for variably quantized  $\hat{\mathbf{r}}$ , we derive the probability model of  $\hat{\mathbf{r}}$  from the learned probability model of  $\mathbf{r}$  at  $\tau = 0$ . The resulting cost of encoding  $\hat{\mathbf{r}}$ , denoted by  $R_{\hat{\mathbf{r}}}$ , is reduced significantly with the increase of  $\tau$ .

Given  $\tau$  and the learned PMF  $p_\theta(r_{i,c}|r_{i,<c}, u_i, C_{r_i})$ ,  $\hat{p}_\theta(\hat{r}_{i,c}|r_{i,<c}, u_i, C_{r_i})$  of quantized  $\hat{r}_{i,c}$  can be computed by the following PMF quantization:

$$\hat{p}_\theta(\hat{r}_{i,c}|r_{i,<c}, u_i, C_{r_i}) = \sum_{v=\hat{r}_{i,c}-\tau}^{\hat{r}_{i,c}+\tau} p_\theta(v|r_{i,<c}, u_i, C_{r_i}) \quad (14)$$

An illustrative example is shown in Fig. 3. Together with (8) and (9), we can derive the probability model  $\hat{p}_\theta(\hat{\mathbf{r}}|\mathbf{u}, C_{\hat{\mathbf{r}}})$  of  $\hat{\mathbf{r}}$ , which is ideal given the learned  $p_\theta(\mathbf{r}|\mathbf{u}, C_{\mathbf{r}})$  of  $\mathbf{r}$ .

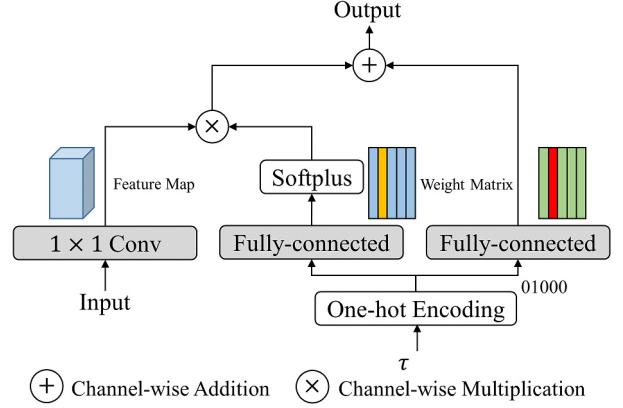


Figure 4. Conditional convolutional layer for bias correction. Different outputs can be produced conditioned on  $\tau \in \{1, 2, \dots, 5\}$ .

However, encoding  $\hat{\mathbf{r}}$  with  $\hat{p}_\theta(\hat{\mathbf{r}}|\mathbf{u}, C_{\hat{\mathbf{r}}})$  results in undecodable bitstreams, since the original residual  $\mathbf{r}$  is unknown to the decoder.  $\hat{p}_\theta(\hat{r}_{i,c}|r_{i,<c}, u_i, C_{r_i})$  cannot be evaluated without  $r_{i,<c}$  and causal context  $C_{r_i}$ . Instead, we evaluate PMF using the quantized residual  $\hat{\mathbf{r}}$ . Because of the mismatch between training (with  $\mathbf{r}$ ) and inference (with  $\hat{\mathbf{r}}$ ) phases, it leads to a biased PMF  $p_\theta(r_{i,c}|\hat{r}_{i,<c}, u_i, C_{\hat{r}_i})$  of  $r_{i,c}$ . We then evaluate  $\hat{p}_\theta(\hat{r}_{i,c}|\hat{r}_{i,<c}, u_i, C_{\hat{r}_i})$  with (14) and derive  $\hat{p}_\theta(\hat{\mathbf{r}}|\mathbf{u}, C_{\hat{\mathbf{r}}})$  for the encoding of  $\hat{\mathbf{r}}$ . The above probability inference scheme is sketched in Fig. 1(b).

### 3.4. Bias Correction

Because of the potential discrepancy between the ideal  $\hat{p}_\theta(\hat{\mathbf{r}}|\mathbf{u}, C_{\hat{\mathbf{r}}})$  and the biased  $\hat{p}_\theta(\hat{\mathbf{r}}|\mathbf{u}, C_{\hat{\mathbf{r}}})$ , encoding  $\hat{\mathbf{r}}$  with  $\hat{p}_\theta(\hat{\mathbf{r}}|\mathbf{u}, C_{\hat{\mathbf{r}}})$  degrades the compression performance. We further propose a bias correction scheme to close the gap between the ideal  $\hat{p}_\theta(\hat{\mathbf{r}}|\mathbf{u}, C_{\hat{\mathbf{r}}})$  and the biased  $\hat{p}_\theta(\hat{\mathbf{r}}|\mathbf{u}, C_{\hat{\mathbf{r}}})$ , while the resulting bitstreams are still decodable.

The components of the proposed bias corrector are illustrated in Fig. 1(a). The masked convolutional layer in the bias corrector is shared with that in the residual compressor. The *conditional entropy model* has the same structure as the *entropy model* illustrated in Fig. 2, but replaces the convolutional layers with the conditional convolutional layers [43, 8] illustrated in Fig. 4.

During training, we generate random  $\tau \in \{1, 2, \dots, N\}$  and quantize  $\mathbf{r}$  to  $\hat{\mathbf{r}}$  with (2). We choose  $N = 5$  in this paper. Given  $\mathbf{u}$  and the extracted context  $C_{\hat{\mathbf{r}}}$ , we use the *conditional entropy model* to estimate  $-\log p_\varphi(\mathbf{r}|\mathbf{u}, C_{\hat{\mathbf{r}}})$  conditioned on different  $\tau$ , and minimize

$$\mathcal{L}(\varphi) = \mathbb{E}_{p(\mathbf{x})} \mathbb{E}_{q_\phi(\hat{y}, \hat{z}|\mathbf{x})} \left[ \log \frac{p_\theta(\mathbf{r}|\mathbf{u}, C_{\mathbf{r}})}{p_\varphi(\mathbf{r}|\mathbf{u}, C_{\hat{\mathbf{r}}})} \right] \quad (15)$$

where  $\varphi$  denote the parameters of the *conditional entropy model*.  $-\log p_\theta(\mathbf{r}|\mathbf{u}, C_{\mathbf{r}})$  is estimated by the *entropy model* in the residual compressor.  $\mathcal{L}(\varphi)$  can be considered as

an approximate KL-divergence or relative entropy [9] between  $p_{\theta}(\mathbf{r}|\mathbf{u}, C_{\mathbf{r}})$  and  $p_{\varphi}(\mathbf{r}|\mathbf{u}, C_{\hat{\mathbf{r}}})$ . Because the *entropy model* receives the context  $C_{\mathbf{r}}$  extracted from the original residual  $\mathbf{r}$  and is only trained with  $\tau = 0$ ,  $p_{\theta}(\mathbf{r}|\mathbf{u}, C_{\mathbf{r}})$  approximates the true distribution  $p(\mathbf{r}|\tilde{\mathbf{x}}, \hat{\mathbf{y}})$  better than  $p_{\varphi}(\mathbf{r}|\mathbf{u}, C_{\hat{\mathbf{r}}})$ . Thus,  $-\log p_{\theta}(\mathbf{r}|\mathbf{u}, C_{\mathbf{r}})$  is the lower bound of  $-\log p_{\varphi}(\mathbf{r}|\mathbf{u}, C_{\hat{\mathbf{r}}})$  on average.

The probability inference scheme with bias correction is sketched in Fig. 1(c). For  $\tau = 0$ , we choose the *entropy model* to estimate  $p_{\theta}(\mathbf{r}|\mathbf{u}, C_{\mathbf{r}})$  to encode  $\mathbf{r}$ . For  $\tau \in \{1, 2, \dots, 5\}$ , we choose the *conditional entropy model* to estimate  $p_{\varphi}(\mathbf{r}|\mathbf{u}, C_{\hat{\mathbf{r}}})$  and derive  $\hat{p}_{\varphi}(\hat{\mathbf{r}}|\mathbf{u}, C_{\hat{\mathbf{r}}})$  with (14) to encode  $\hat{\mathbf{r}}$ . As  $\hat{p}_{\varphi}(\hat{\mathbf{r}}|\mathbf{u}, C_{\hat{\mathbf{r}}})$  approximates the ideal  $\hat{p}_{\theta}(\hat{\mathbf{r}}|\mathbf{u}, C_{\mathbf{r}})$  better than the biased  $\hat{p}_{\theta}(\hat{\mathbf{r}}|\mathbf{u}, C_{\hat{\mathbf{r}}})$ , the compression performance can be improved. Since evaluating  $\hat{p}_{\varphi}(\hat{\mathbf{r}}|\mathbf{u}, C_{\hat{\mathbf{r}}})$  is independent of  $\mathbf{r}$ , the resulting bitstreams are decodable. We analyze the efficacy of the proposed bias correction scheme in Sec. 4.4.

**Training Strategy.** Bias corrector is trained together with the lossy image compressor and the residual compressor, but minimizing (15) only updates the parameters of the *conditional entropy model* as shown in Fig. 1(a). The masked convolutional layer is shared with that in the residual compressor, thus can be updated by minimizing (13). This leads to three advantages: 1) achieving the target *conditional entropy model*; 2) circumventing the relaxation of residual quantization; 3) avoiding degrading the estimation of  $p_{\theta}(\mathbf{r}|\mathbf{u}, C_{\mathbf{r}})$  caused by training with mixed  $\tau$ .

## 4. Experiments

### 4.1. Experimental Setup

**Training.** We train on DIV2K high resolution training dataset [1] consisting of 800 2K resolution color images. Although DIV2K is originally built for image super-resolution task, it contains large numbers of high-quality images that is suitable for training our near-lossless codec. The images are first cropped into 27958 samples with the size of  $256 \times 256$ . To augment training data and prevent overfitting on noise, we use bicubic method to downscale each sample with a random factor selected from  $[0.6, 1.0]$  and further randomly crop the downsampled sample to  $128 \times 128$  during training. Our near-lossless codec is optimized for 400 epochs using Adam [16] with minibatches of size 16. The learning rate is set to  $1 \times 10^{-4}$  for the first 350 epochs and decays to  $1 \times 10^{-5}$  for the remaining 50 epochs.

**Evaluation.** We evaluate our near-lossless codec on four datasets: 1) Kodak dataset [20] consists of 24 uncompressed  $768 \times 512$  color images. 2) DIV2K high resolution validation dataset [1] consists of 100 2K resolution color images. 3) CLIC professional validation dataset (CLIC.p) [49] consists of 41 color images taken by professional photographers. 4) CLIC mobile validation dataset (CLIC.m) [49]

Codec	Kodak	DIV2K	CLIC.p	CLIC.m
PNG	4.35	4.23	3.93	3.93
JPEG-LS	3.16	2.99	2.82	2.53
CALIC	3.18	3.07	2.87	2.59
JPEG2000	3.19	3.12	2.93	2.71
WebP	3.18	3.11	2.90	2.73
BPG	3.38	3.28	3.08	2.84
FLIF	<b>2.90</b>	2.91	2.72	<b>2.48</b>
L3C	–	3.09	2.94	2.64
RC	–	3.08	2.93	2.54
Ours	3.04	<b>2.81</b>	<b>2.66</b>	2.51

Table 1. Compression performance (bpsp) of the proposed near-lossless image codec with  $\tau = 0$ , compared to other lossless image codecs on Kodak, DIV2K, CLIC.p and CLIC.m datasets.

Codec	$\tau^*$	Kodak	DIV2K	CLIC.p	CLIC.m
JPEG-LS	1	2.90	2.62	2.34	2.44
	2	2.30	2.07	1.80	1.89
	4	1.68	1.53	1.28	1.35
CALIC	1	2.75	2.45	2.18	2.28
	2	2.14	1.88	1.62	1.70
	4	1.51	1.31	1.07	1.13
WebP nll	1	2.41	2.45	2.26	2.11
	2	2.01	2.04	1.89	1.85
	4	1.82	1.83	1.73	1.75
Ours w/ bc	1	1.84	1.72	1.60	1.53
	2	1.38	1.29	1.15	1.10
	4	0.92	0.89	0.74	0.72

Table 2. Compression performance (bpsp) of the proposed near-lossless image codec (w/ bias correction) with  $\tau > 0$ , compared to near-lossless JPEG-LS, near-lossless CALIC and near-lossless WebP on Kodak, DIV2K, CLIC.p and CLIC.m datasets. \*The error bounds of near-lossless WebP are powers of two.

consists of 61 color images taken using mobile phones. Images in CLIC.p and CLIC.m are typically 2K resolution but some of them are of small sizes.

### 4.2. Performance Evaluation at $\tau = 0$

We first evaluate the compression performance of the proposed near-lossless image codec at  $\tau = 0$ , measured by *bits per subpixel* (bpsp). Each RGB pixel has 3 subpixels. We compare with seven traditional lossless image codecs, *i.e.*, PNG, JPEG-LS [47], CALIC [50], JPEG2000 [37], WebP [14], BPG [5], FLIF [38], and two the state-of-the-art learned lossless image codecs, *i.e.*, L3C [28] and RC [30]. Unlike ours, L3C [28] and RC [30] are trained on a much larger dataset (300,000 images from the Open Images dataset [21]). We report the compression performance of L3C [28] and RC [30] published by their authors.

As reported in Table 1, our codec achieves the best performance on DIV2K validation dataset, which shares the




Figure 5. Rate-distortion curves of the proposed near-lossless image codec (w/ bias correction) with  $\tau > 0$  compared to other near-lossless image codecs and lossy image codecs on Kodak dataset. (a) bps- $\ell_\infty$  error curves. (b) bps-PSNR curves.




Figure 6. Bit rates of each image compressed by the proposed near-lossless image codec with  $\tau = \{0, 1, \dots, 5\}$  on Kodak dataset.

same domain with the training dataset. Our codec achieves the best performance on CLIC.p, and the second best performance on Kodak and CLIC.m, only secondary to FLIF. The results demonstrate that the jointly learned lossy image compressor and residual compressor with  $\tau = 0$  can effectively model  $p_\theta(\mathbf{x})$  and can be generalized to various domains of natural images. This provides a solid foundation for the near-lossless image compression with  $\tau > 0$ .

### 4.3. Performance Evaluation at $\tau > 0$

We next evaluate the compression performance of the proposed near-lossless image codec (with bias correction) at  $\tau > 0$ . We compare with near-lossless JPEG-LS [47], near-lossless CALIC [51] and near-lossless WebP (WebP nll) [14]. Besides, we compare with six traditional lossy image codecs, *i.e.*, JPEG [44], JPEG2000 [37], WebP [14], BPG [5], lossy FLIF [38] and VVC [32]. Recent learned lossy image codecs published online are trained at relatively low bit rates ( $\leq 2$  bpp  $\approx 0.67$  bpsp on Kodak). We reimplement and train two recent learned lossy image codecs, *i.e.*, Ballé-hyper [4] and Minnen-joint [31], at high bit rates ( $\geq 0.8$  bpsp on Kodak) for comparison.




Figure 7. Compression performance of the proposed near-lossless image codec without bias correction, with bias correction and with the ideal  $\hat{p}_\theta(\hat{\mathbf{r}}|\mathbf{u}, C_r)$  on Kodak dataset.

Near-lossless WebP simply adjusts pixel values to the  $\ell_\infty$  error bound  $\tau$  and compresses the pre-processed images losslessly. Near-lossless JPEG-LS and CALIC adopt predictive coding schemes, and encode the prediction residuals quantized by (2). The predictors and the probability models used in near-lossless JPEG-LS and CALIC are hand-crafted, which are not effective enough. More effectively, our near-lossless codec is based on jointly trained lossy image compressor and residual compressor. We employ (2) to realize the error bound  $\tau$  and the probability models of the quantized residuals can be derived from the learned residual compressor with  $\tau = 0$ . Therefore, our codec outperforms near-lossless JPEG-LS, CALIC and WebP by a wide margin, as shown in Table 2.

The rate-distortion performance of all codecs on Kodak dataset is reported in Fig. 5. In terms of  $\ell_\infty$  error defined in (1), our codec consistently yields the best results among all codecs, as shown in Fig. 5(a). Besides  $\ell_\infty$  error, we also compare PSNR of all codecs, as shown in Fig. 5(b). Our codec is better than or comparable with JPEG, near-lossless JPEG-LS, near-lossless CALIC, WebP, near-lossless WebP, BPG (4:2:0) and lossy FLIF. Our codec is surpassed by JPEG2000, BPG (4:4:4), VVC, Ballé-hyper and Minnen-joint. Overall, our codec achieves competitive performance at bit rates higher than 0.8 bpsp, although PSNR is not our optimization objective.

### 4.4. Ablation Study

**Scalability.** In Fig. 6, we show the bit rates of each image compressed by the proposed near-lossless image codec with  $\tau = \{0, 1, \dots, 5\}$  on Kodak dataset.  $R_{\hat{\mathbf{y}}, \hat{\mathbf{z}}}$ , on average, accounts for about 18% of  $R_{\hat{\mathbf{y}}, \hat{\mathbf{z}}} + R_r$  at  $\tau = 0$ . With the increase of  $\tau$ , the bit rate  $R_{\hat{\mathbf{r}}}$  of the quantized residual  $\hat{\mathbf{r}}$  is significantly reduced. Especially the  $\tau = 1$  mode saves about 40% bit rates compared with the  $\tau = 0$  lossless image compression mode.

**Bias Correction.** In Fig. 7, we demonstrate the efficacy of bias correction at  $\tau > 0$ . Because of the discrepancy between the ideal  $\hat{p}_\theta(\hat{\mathbf{r}}|\mathbf{u}, C_r)$  and the biased  $\hat{p}_\theta(\hat{\mathbf{r}}|\mathbf{u}, C_r)$ , encoding  $\hat{\mathbf{r}}$  with  $\hat{p}_\theta(\hat{\mathbf{r}}|\mathbf{u}, C_r)$  (without bias correction) degrades the compression performance. Instead, we encode  $\hat{\mathbf{r}}$




Figure 8. Visualizing the effects of different  $\lambda$ 's on the lossy reconstruction  $\tilde{\mathbf{x}}$  and the near-lossless (nll) reconstruction  $\hat{\mathbf{x}}$ . The images are best viewed in full size on screen.

$\lambda$	Total Rate	$R_{\hat{y}, \hat{z}}$	$R_r$	$\tilde{\mathbf{x}}$ (PSNR)
0	2.93	0.05	2.88	24.09
0.01	3.00	0.42	2.58	37.55
0.03	3.04	0.55	2.49	39.29
0.05	3.06	0.62	2.44	39.91

Table 3. The effects of different  $\lambda$ 's on the learned near-lossless image codec with  $\tau = 0$  on Kodak dataset.

with  $\hat{p}_\varphi(\hat{\mathbf{r}}|\mathbf{u}, C_{\hat{\mathbf{r}}})$  (with bias correction) resulting in lower bit rates. With the increase of  $\tau$ , the PMF of  $\mathbf{r}$  is quantized by larger bins and becomes coarser. Thus, the compression performance with bias correction approaches the “ideal”. However, the gap between the compression performance without bias correction and the “ideal” remains large, as the discrepancy between  $\hat{p}_\theta(\hat{\mathbf{r}}|\mathbf{u}, C_r)$  and  $\hat{p}_\theta(\hat{\mathbf{r}}|\mathbf{u}, C_{\hat{\mathbf{r}}})$  is also magnified with the increasing  $\tau$ . Note that the compression performance of our codec without bias correction is also better than near-lossless JPEG-LS, CALIC and WebP.

**Discussion on different  $\lambda$ 's.** The  $\lambda$  in (13) adapts the rate of losslessly encoding the raw image  $\mathbf{x}$  and the distortion of the lossy reconstruction  $\tilde{\mathbf{x}}$ . In Table 3, we evaluate the effects of different  $\lambda \in \{0, 0.01, 0.03, 0.05\}$  on Kodak dataset. With the increase of  $\lambda$ , the PSNR of  $\tilde{\mathbf{x}}$  is improved but the bps of encoding  $\mathbf{x}$  also becomes higher. If we simply do lossless image compression ( $\tau = 0$ ), we should choose  $\lambda = 0$ . The  $\tilde{\mathbf{x}}$  is learned by the network without any hand-crafted constraints, leading to the best compression performance.

However,  $\lambda = 0$  is not suitable for our near-lossless codec with  $\tau > 0$ . As the average PSNR of the  $\tilde{\mathbf{x}}$  is 24.09

dB, the magnitude of  $\mathbf{r} = \mathbf{x} - \tilde{\mathbf{x}}$  is large. To quantize such large residuals with large  $\tau$  results in blocking artifacts in the near-lossless reconstruction  $\hat{\mathbf{x}}$ . As shown in Fig. 8(b),  $\tilde{\mathbf{x}}$  with  $\lambda = 0$  is blurry. If  $\tau = 5$ , the quantization bin size is large, *i.e.*,  $2\tau + 1 = 11$ , resulting in blocking artifacts in  $\hat{\mathbf{x}}$  in Fig. 8(d). For  $\lambda = 0.01$ , we also observe small artifacts in  $\hat{\mathbf{x}}$  caused by quantization with large  $\tau$ , as shown in Fig. 8(i). In our codec, we choose  $\lambda = 0.03$ , which is the best trade-off between compression performance and visual quality in the experiments. As shown in Fig. 8(e) and Fig. 8(j), we observe no artifacts in  $\hat{\mathbf{x}}$  compared with  $\lambda = 0, 0.01$ .

## 5. Conclusion

In this paper, we propose a joint lossy image and residual compression framework to learn a scalable  $\ell_\infty$ -constrained near-lossless image codec. With our codec, a raw image is first compressed by the lossy image compressor, and the corresponding residual is then quantized to satisfy variable  $\ell_\infty$  error bounds. To compress the variably quantized residuals, the probability models of the quantized residuals are derived from the learned probability model of the original residual, which is realized by the residual compressor in conjunction with the bias corrector. Experimental results demonstrate the state-of-the-art performance of our near-lossless image codec.

## Acknowledgement

This work was supported by National Natural Science Foundation of China under Grants 61922027, 61827804 and U20B2052, National Key Research and Development Project under Grant 2019YFE0109600, China Postdoctoral Science Foundation under Grant 2020M682826.



## References

- [1] Eirikur Agustsson and Radu Timofte. NTIRE 2017 challenge on single image super-resolution: dataset and study. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, July 2017. 6
- [2] Rashid Ansari, Nasir D. Memon, and Ersan Ceran. Near-lossless image compression techniques. *Journal of Electronic Imaging*, 7(3):486–494, 1998. 1, 2
- [3] Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end optimized image compression. In *Int. Conf. Learn. Represent.*, 2017. 1, 2
- [4] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. In *Int. Conf. Learn. Represent.*, 2018. 1, 2, 3, 5, 7
- [5] Fabrice Bellard. BPG image format. <https://bellard.org/bpg/>. 1, 2, 3, 6, 7
- [6] Keshi Chen and Tenkasi V Ramabadran. Near-lossless compression of medical images through entropy-coded dpcm. *IEEE Transactions on Medical Imaging*, 13(3):538–548, 1994. 1, 2
- [7] Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7939–7948, 2020. 1, 2, 3
- [8] Yoojin Choi, Mostafa El-Khamy, and Jungwon Lee. Variable rate deep image compression with a conditional autoencoder. In *Int. Conf. Comput. Vis.*, pages 3146–3154, 2019. 1, 2, 3, 5
- [9] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, USA, 2006. 6
- [10] Sharaf E Elnahas. Data compression with applications to digital radiology. *Washington Univ., St. Louis, MO (USA)*, 1986. 1
- [11] Vivek K Goyal. Theoretical foundations of transform coding. *IEEE Sign. Process. Magazine*, 18(5):9–21, 2001. 1, 2, 3
- [12] Zongyu Guo, Yaojun Wu, Runsen Feng, Zhizheng Zhang, and Zhibo Chen. 3-d context entropy model for improved practical image compression. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 116–117, 2020. 3
- [13] Emiel Hoogeboom, Jorn Peters, Rianne van den Berg, and Max Welling. Integer discrete flows and lossless compression. In *Adv. Neural Inform. Process. Syst.*, pages 12134–12144, 2019. 1, 2
- [14] WebP image format. <https://developers.google.com/speed/webp/>. 1, 2, 6, 7
- [15] Ligang Ke and Michael W Marcellin. Near-lossless image compression: minimum-entropy, constrained-error dpcm. *IEEE Trans. Image Process.*, 7(2):225–228, 1998. 1, 2
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Int. Conf. Learn. Represent.*, 2015. 6
- [17] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *Int. Conf. Learn. Represent.*, 2014. 1, 2, 3
- [18] Diederik P. Kingma and Max Welling. An introduction to variational autoencoders. *Foundations and Trends in Machine Learning*, 12(4):307–392, 2019. 1, 2, 3
- [19] Friso H Kingma, Pieter Abbeel, and Jonathan Ho. Bit-swap: Recursive bits-back coding for lossless compression with hierarchical latent variables. In *Int. Conf. Mach. Learn.*, 2019. 2
- [20] Eastman Kodak. Kodak lossless true color image suite (photocd pcd0992), 1993. <http://r0k.us/graphics/kodak/>. 6
- [21] Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Andreas Veit, et al. Openimages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from <https://github.com/openimages>*, 2(3):2–3, 2017. 6
- [22] Jooyoung Lee, Seunghyun Cho, and Seung-Kwon Beack. Context-adaptive entropy model for end-to-end optimized image compression. In *Int. Conf. Learn. Represent.*, May 2019. 1, 2
- [23] Mu Li, Wangmeng Zuo, Shuhang Gu, Jane You, and David Zhang. Learning content-weighted deep image compression. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2020. 1, 2
- [24] M. Li, W. Zuo, S. Gu, D. Zhao, and D. Zhang. Learning convolutional networks for content-weighted image compression. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3214–3223, 2018. 1, 2
- [25] H. Ma, D. Liu, N. Yan, H. Li, and F. Wu. End-to-end optimized versatile image compression with wavelet-like transform. *IEEE Trans. Pattern Anal. Mach. Intell.*, pages 1–1, 2020. 1, 2
- [26] Paul W Melnychuck and Majid Rabbani. Survey of lossless image coding techniques. In *Digital Image Processing Applications*, volume 1075, pages 92–100. International Society for Optics and Photonics, 1989. 1
- [27] F. Mentzer, E. Agustsson, M. Tschannen, R. Timofte, and L. V. Gool. Conditional probability models for deep image compression. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4394–4402, 2018. 1, 2
- [28] F. Mentzer, E. Agustsson, M. Tschannen, R. Timofte, and L. Van Gool. Practical full resolution learned lossless image compression. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10621–10630, 2019. 2, 4, 6
- [29] Fabian Mentzer, George Toderici, Michael Tschannen, and Eirikur Agustsson. High-fidelity generative image compression. In *Adv. Neural Inform. Process. Syst.*, 2020. 1, 2
- [30] Fabian Mentzer, Luc Van Gool, and Michael Tschannen. Learning better lossless compression using lossy compression. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 2, 3, 4, 6
- [31] David Minnen, Johannes Ballé, and George D Toderici. Joint autoregressive and hierarchical priors for learned image compression. In *Adv. Neural Inform. Process. Syst.*, pages 10771–10780, 2018. 1, 2, 3, 4, 7
- [32] Jens-Rainer Ohm and Gary J Sullivan. Versatile video coding—towards the next generation of video compression. In *Picture Coding Symposium*, 2018. 1, 7

- [33] Aäron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *Int. Conf. Mach. Learn.*, pages 1747–1756. JMLR.org, 2016. [2](#)
- [34] Oren Rippel and Lubomir Bourdev. Real-time adaptive image compression. In *Int. Conf. Mach. Learn.*, volume 70, pages 2922–2930, 2017. [1, 2](#)
- [35] Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P Kingma. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. In *Int. Conf. Learn. Represent.*, 2017. [2, 4](#)
- [36] Claude E Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948. [1](#)
- [37] Athanassios Skodras, Charilaos Christopoulos, and Touradj Ebrahimi. The jpeg 2000 still image compression standard. *IEEE Sign. Process. Magazine*, 18(5):36–58, 2001. [1, 2, 6, 7](#)
- [38] Jon Sneyers and Pieter Wuille. Flif: Free lossless image format based on maniac compression. In *IEEE Int. Conf. Image Process.*, pages 66–70. IEEE, 2016. [1, 6, 7](#)
- [39] Lucas Theis, Wenzhe Shi, Andrew Cunningham, and Ferenc Huszár. Lossy image compression with compressive autoencoders. In *Int. Conf. Learn. Represent.*, 2017. [1, 2](#)
- [40] George Toderici, Sean M. O’Malley, Sung Jin Hwang, Damien Vincent, David Minnen, Shumeet Baluja, Michele Covell, and Rahul Sukthankar. Variable rate image compression with recurrent neural networks. In *Int. Conf. Learn. Represent.*, 2016. [1, 2](#)
- [41] George Toderici, Damien Vincent, Nick Johnston, Sung Jin Hwang, David Minnen, Joel Shor, and Michele Covell. Full resolution image compression with recurrent neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5306–5314, 2017. [1, 2](#)
- [42] James Townsend, Tom Bird, and David Barber. Practical lossless compression with latent variables using bits back coding. In *Int. Conf. Learn. Represent.*, 2019. [2](#)
- [43] Aäron van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, and Alex Graves. Conditional image generation with pixelcnn decoders. In *Adv. Neural Inform. Process. Syst.*, pages 4790–4798, 2016. [2, 5](#)
- [44] Gregory K Wallace. The jpeg still picture compression standard. *IEEE Trans. Consumer Electronics*, 38(1):xviii–xxxiv, 1992. [2, 7](#)
- [45] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4):600–612, 2004. [1](#)
- [46] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multi-scale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402. Ieee, 2003. [1](#)
- [47] Marcelo J Weinberger, Gadiel Seroussi, and Guillermo Sapiro. The loco-i lossless image compression algorithm: Principles and standardization into jpeg-ls. *IEEE Trans. Image Process.*, 9(8):1309–1324, 2000. [1, 2, 3, 6, 7](#)
- [48] Ian H Witten, Radford M Neal, and John G Cleary. Arithmetic coding for data compression. *Communications of the ACM*, 30(6):520–540, 1987. [2](#)
- [49] Workshop and Challenge on Learned Image Compression. <https://www.compression.cc/challenge/>. [6](#)
- [50] X. Wu and N. Memon. Context-based, adaptive, lossless image coding. *IEEE Trans. Communications*, 45(4):437–444, 1997. [1, 6](#)
- [51] Wu Xiaolin and P. Bao.  $l_\infty$  constrained high-fidelity image compression via adaptive context modeling. *IEEE Trans. Image Process.*, 9(4):536–542, 2000. [1, 2, 3, 7](#)
- [52] X. Zhang and X. Wu. Near-lossless  $l_\infty$ -constrained image decompression via deep neural network. In *Data Compression Conference*, pages 33–42, 2019. [2](#)
- [53] X. Zhang and X. Wu. Ultra high fidelity deep image decompression with  $l_\infty$ -constrained compression. *IEEE Trans. Image Process.*, 30:963–975, 2021. [2](#)