

CONTEXT-ADAPTIVE ENTROPY MODEL FOR END-TO-END OPTIMIZED IMAGE COMPRESSION

Jooyoung Lee, Seunghyun Cho & Seung-Kwon Beack

Broadcasting Media Research Laboratory
Electronics and Telecommunications Research Institute
Daejeon, Korea
{leejy1003, shcho, skbeack}@etri.re.kr

ABSTRACT

We propose a context-adaptive entropy model for use in end-to-end optimized image compression. Our model exploits two types of contexts, bit-consuming contexts and bit-free contexts, distinguished based upon whether additional bit allocation is required. Based on these contexts, we allow the model to more accurately estimate the distribution of each latent representation with a more generalized form of the approximation models, which accordingly leads to an enhanced compression performance. Based on the experimental results, the proposed method outperforms the traditional image codecs, such as BPG and JPEG2000, as well as other previous artificial-neural-network (ANN) based approaches, in terms of the peak signal-to-noise ratio (PSNR) and multi-scale structural similarity (MS-SSIM) index. The test code is publicly available at https://github.com/JooyoungLeeETRI/CA_Entropy_Model.

1 INTRODUCTION

Recently, artificial neural networks (ANNs) have been applied in various areas and have achieved a number of breakthroughs resulting from their superior optimization and representation learning performance. In particular, for various problems that are sufficiently straightforward that they can be solved within a short period of time by hand, a number of ANN-based studies have been conducted and significant progress has been made. With regard to image compression, however, relatively slow progress has been made owing to its complicated target problems. A number of works, focusing on the quality enhancement of reconstructed images, were proposed. For instance, certain approaches (Dong et al., 2015; Svoboda et al., 2016; Zhang et al., 2017) have been proposed to reduce artifacts caused by image compression, relying on the superior image restoration capability of an ANN. Although it is indisputable that artifact reduction is one of the most promising areas exploiting the advantages of ANNs, such approaches can be viewed as a type of post-processing, rather than image compression itself.

Regarding ANN-based image compression, the previous methods can be divided into two types. First, as a consequence of the recent success of generative models, some image compression approaches targeting the superior perceptual quality (Agustsson et al., 2018; Santurkar et al., 2018; Rippel & Bourdev, 2017) have been proposed. The basic idea here is that learning the distribution of natural images enables a very high compression level without severe perceptual loss by allowing the generation of image components, such as textures, which do not highly affect the structure or the perceptual quality of the reconstructed images. Although the generated images are very realistic, the acceptability of the machine-created image components eventually becomes somewhat application-dependent. Meanwhile, a few end-to-end optimized ANN-based approaches (Toderici et al., 2017; Johnston et al., 2018; Ballé et al., 2017; Theis et al., 2017; Ballé et al., 2018), without generative models, have been proposed. In these approaches, unlike traditional codecs comprising separate tools, such as prediction, transform, and quantization, a comprehensive solution covering all functions has been sought after using end-to-end optimization. Toderici et al. (2017)'s approach exploits a small number of latent binary representations to contain the compressed information in every step, and each step increasingly stacks the additional latent representations to achieve a progressive improvement in quality of the reconstructed images. Johnston et al. (2018) improved the compression

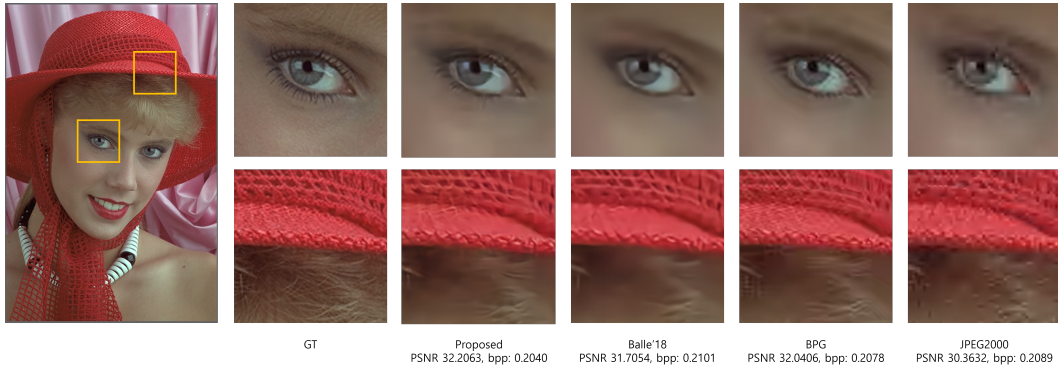


Figure 1: Comparison of sample test results including the ground truth, our method, Ballé et al. (2018)’s approach, BPG, and JPEG2000.

performance by enhancing operation methods of the networks developed by Toderici et al. (2017). Although Toderici et al. (2017); Johnston et al. (2018) provided novel frameworks suitable to quality control using a single trained network, the increasing number of iteration steps to obtain higher image quality can be a burden to certain applications. In contrast to the approaches developed by Toderici et al. (2017) and Johnston et al. (2018), which extract binary representations with as high an entropy as possible, Ballé et al. (2017), Theis et al. (2017), and Ballé et al. (2018) regard the image compression problem as being how to retrieve discrete latent representations having as low an entropy as possible. In other words, the target problem of the former methods can be viewed as how to include as much information as possible in a fixed number of representations, whereas the latter is simply how to reduce the expected bit-rate when a sufficient number of representations are given, assuming that the low entropy corresponds to small number of bits from the entropy coder. To solve the second target problem, Ballé et al. (2017), Theis et al. (2017), and Ballé et al. (2018) adopt their own entropy models to approximate the actual distributions of the discrete latent representations. More specifically, Ballé et al. (2017) and Theis et al. (2017) proposed novel frameworks that exploit the entropy models, and proved their performance capabilities by comparing the results with those of conventional codecs such as JPEG2000. Whereas Ballé et al. (2017) and Theis et al. (2017) assume that each representation has a fixed distribution, Ballé et al. (2018) introduced an input-adaptive entropy model that estimates the scale of the distribution for each representation. This idea is based on the characteristics of natural images in which the scales of the representations vary together in adjacent areas. They provided test results that outperform all previous ANN-based approaches, and reach very close to those of BPG (Bellard, 2014), which is known as a subset of HEVC (ISO/IEC 23008-2, ITU-T H.265), used for image compression.

One of the principle elements in end-to-end optimized image compression is the trainable entropy model used for the latent representations. Because the actual distributions of latent representations are unknown, the entropy models provide the means to estimate the required bits for encoding the latent representations by approximating their distributions. When an input image x is transformed into a latent representation y and then uniformly quantized into \hat{y} , the simple entropy model can be represented by $p_{\hat{y}}(\hat{y})$, as described by Ballé et al. (2018). When the actual marginal distribution of \hat{y} is denoted as $m(\hat{y})$, the rate estimation, calculated through cross entropy using the entropy model, $p_{\hat{y}}(\hat{y})$, can be represented as shown in equation (1), and can be decomposed into the actual entropy of \hat{y} and the additional bits owing to a mismatch between the actual distributions and their approximations. Therefore, decreasing the rate term R during the training process allows the entropy model $p_{\hat{y}}(\hat{y})$ to approximate $m(\hat{y})$ as closely as possible, and let the other parameters transform x into y properly such that the actual entropy of \hat{y} becomes small.

$$R = \mathbb{E}_{\hat{y} \sim m}[-\log_2 p_{\hat{y}}(\hat{y})] = H(m) + D_{KL}(m||p_{\hat{y}}). \quad (1)$$

In terms of KL-divergence, R is minimized when $p_{\hat{y}}(\hat{y})$ becomes perfectly matched with the actual distribution $m(\hat{y})$. This means that the compression performance of the methods essentially depends on the capacity of the entropy model. To enhance the capacity, we propose a new entropy model

that exploits two types of contexts, bit-consuming and bit-free contexts, distinguished according to whether additional bit allocation is required. Utilizing these two contexts, we allow the model to more accurately estimate the distribution of each latent representation through the use of a more generalized form of the entropy models, and thus more effectively reduce the spatial dependencies among the adjacent latent representations. Figure 1 demonstrates a comparison of the compression results of our method to those of other previous approaches. The contributions of our work are as follows:

- We propose a new context-adaptive entropy model framework that incorporates the two different types of contexts.
- We provide the test results that outperform the widely used conventional image codec BPG in terms of the PSNR and MS-SSIM.
- We discuss the directions of improvement in the proposed methods in terms of the model capacity and the level of the contexts.

Note that we follow a number of notations given by Ballé et al. (2018) because our approach can be viewed as an extension of their work, in that we exploit the same rate-distortion (R-D) optimization framework. The rest of this paper is organized as follows. In Section 2, we introduce the key approaches of end-to-end optimized image compression and propose the context-adaptive entropy model. Section 3 demonstrates the structure of the encoder and decoder models used, and the experimental setup and results are then given in section 4. Finally, in Section 5, we discuss the current state of our work and directions for improvement.

2 END-TO-END OPTIMIZATION BASED ON CONTEXT-ADAPTIVE ENTROPY MODELS

2.1 PREVIOUS ENTROPY MODELS

Since they were first proposed by Ballé et al. (2017) and Theis et al. (2017), entropy models, which approximate the distribution of discrete latent representations, have noticeably improved the image compression performance of ANN-based approaches. Ballé et al. (2017) assumes the entropy models of the latent representations as non-parametric models, whereas Theis et al. (2017) adopted a Gaussian scale mixture model composed of six weighted zero-mean Gaussian models per representation. Although they assume different forms of entropy models, they have a common feature in that both concentrate on learning the distributions of the representations without considering input adaptivity. In other words, once the entropy models are trained, the trained model parameters for the representations are fixed for any input during the test time. Ballé et al. (2018), in contrast, introduced a novel entropy model that adaptively estimates the scales of the representations based on input. They assume that the scales of the latent representations from the natural images tend to move together within an adjacent area. To reduce this redundancy, they use a small amount of additional information by which the proper scale parameters (standard deviations) of the latent representations are estimated. In addition to the scale estimation, Ballé et al. (2018) have also shown that when the prior probability density function (PDF) for each representation in a continuous domain is convolved with a standard uniform density function, it approximates the prior probability mass function (PMF) of the discrete latent representation, which is uniformly quantized by rounding, much more closely. For training, a uniform noise is added to each latent representation so as to fit the distribution of these noisy representations into the mentioned PMF-approximating functions. Using these approaches, Ballé et al. (2018) achieved a state-of-the-art compression performance, close to that of BPG.

2.2 SPATIAL DEPENDENCIES OF THE LATENT VARIABLES

The latent representations, when transformed through a convolutional neural network, essentially contain spatial dependencies because the same convolutional filters are shared across the spatial regions, and natural images have various factors in common within adjacent regions. Ballé et al. (2018) successfully captured these spatial dependencies and enhanced the compression performance by input-adaptively estimating standard deviations of the latent representations. Taking a step forward, we generalize the form of the estimated distribution by allowing, in addition to the standard




Figure 2: Examples of latent representations and their normalized versions for the two cases (the first two images show the results in which only the standard deviations are estimated using side information, whereas the last two images show the results in which the mu and standard deviation are estimated using our method). For a clear demonstration, the latent representations having the highest covariance between the spatially adjacent variables are extracted. Left: the latent representations of $\hat{\mathbf{y}}$ from the first case. Middle left: the normalized versions of $\hat{\mathbf{y}}$ from the first case, divided by the estimated standard deviation. Middle right: the latent variable of $\hat{\mathbf{y}}$ from the second case. Right: the normalized versions of $\hat{\mathbf{y}}$ from the second case, shifted and divided by the estimated mu and the standard deviation.

deviation, the mean estimation utilizing the contexts. For instance, assuming that certain representations tend to have similar values within a spatially adjacent area, when all neighborhood representations have a value of 10, we can intuitively guess that, for the current representation, the chance of having a value equal or similar to 10 is relatively high. This simple estimation will consequently reduce the entropy. Likewise, our method utilizes the given contexts for estimating the mean, as well as the standard deviation, of each latent representation. Note that Toderici et al. (2017), Johnston et al. (2018), and Rippel & Bourdev (2017) also apply context-adaptive entropy coding by estimating the probability of each binary representation. However, these context-adaptive entropy-coding methods can be viewed as separate components, rather than one end-to-end optimization component because their probability estimation does not directly contribute to the rate term of the R-D optimization framework. Figure 2 visualizes the latent variables $\hat{\mathbf{y}}$ and their normalized versions of the two different approaches, one estimating only the standard deviation parameters and the other estimating both the mu and standard deviation parameters with the two types of mentioned contexts. The visualization shows that the spatial dependency can be removed more effectively when the mu is estimated along with the given contexts.

2.3 CONTEXT-ADAPTIVE ENTROPY MODEL

The optimization problem described in this paper is similar with Ballé et al. (2018), in that the input \mathbf{x} is transformed into \mathbf{y} having a low entropy, and the spatial dependencies of \mathbf{y} are captured into $\hat{\mathbf{z}}$. Therefore, we also use four fundamental parametric transform functions: an analysis transform $g_a(\mathbf{x}; \phi_g)$ to transform \mathbf{x} into a latent representation \mathbf{y} , a synthesis transform $g_s(\hat{\mathbf{y}}; \theta_g)$ to reconstruct image $\hat{\mathbf{x}}$, an analysis transform $h_a(\hat{\mathbf{y}}; \phi_h)$ to capture the spatial redundancies of $\hat{\mathbf{y}}$ into a latent representation \mathbf{z} , and a synthesis transform $h_s(\hat{\mathbf{z}}; \theta_h)$ used to generate the contexts for the model estimation. Note that h_s does not estimate the standard deviations of the representations directly as in Ballé et al. (2018)’s approach. In our method, instead, h_s generates the context \mathbf{c}' , one of the two types of contexts for estimating the distribution. These two types of contexts are described in this section.

Ballé et al. (2018) analyzed the optimization problem from the viewpoint of the variational autoencoder (Kingma & Welling (2014); Rezende et al. (2014)), and showed that the minimization of the KL-divergence is the same problem as the R-D optimization of image compression. Basically, we follow the same concept; however, for training, we use the discrete representations on the conditions instead of the noisy representations, and thus the noisy representations are only used as the inputs to the entropy models. Empirically, we found that using discrete representations on the conditions show better results, as shown in appendix 6.2. These results might come from removing the mis-

matches of the conditions between the training and testing, thereby enhancing the training capacity by limiting the affect of the uniform noise only to help the approximation to the probability mass functions. We use the gradient overriding method with the identity function, as in Theis et al. (2017), to deal with the discontinuities from the uniform quantization. The resulting objective function used in this paper is given in equation (2). The total loss consists of two terms representing the rates and distortions, and the coefficient λ controls the balance between the rate and distortion during the R-D optimization. Note that λ is not an optimization target, but a manually configured condition that determines which to focus on between rate and distortion:

$$\mathcal{L} = R + \lambda D \quad (2)$$

$$\text{with } R = \mathbb{E}_{\mathbf{x} \sim p_{\mathbf{x}}} \mathbb{E}_{\tilde{\mathbf{y}}, \tilde{\mathbf{z}} \sim q} \left[-\log p_{\tilde{\mathbf{y}}|\tilde{\mathbf{z}}}(\tilde{\mathbf{y}} | \tilde{\mathbf{z}}) - \log p_{\tilde{\mathbf{z}}}(\tilde{\mathbf{z}}) \right],$$

$$D = \mathbb{E}_{\mathbf{x} \sim p_{\mathbf{x}}} \left[-\log p_{\mathbf{x}|\hat{\mathbf{y}}}(\mathbf{x} | \hat{\mathbf{y}}) \right]$$

Here, the noisy representations of $\tilde{\mathbf{y}}$ and $\tilde{\mathbf{z}}$ follow the standard uniform distribution, the mean values of which are \mathbf{y} and \mathbf{z} , respectively, when \mathbf{y} and \mathbf{z} are the result of the transforms g_a and h_a , respectively. Note that the input to h_a is $\hat{\mathbf{y}}$, which is a uniformly quantized representation of \mathbf{y} , rather than the noisy representation $\tilde{\mathbf{y}}$. Q denotes the uniform quantization function, for which we simply use a rounding function:

$$q(\tilde{\mathbf{y}}, \tilde{\mathbf{z}} | \mathbf{x}, \phi_g, \phi_h) = \prod_i \mathcal{U}(\tilde{y}_i | y_i - \frac{1}{2}, y_i + \frac{1}{2}) \cdot \prod_j \mathcal{U}(\tilde{z}_j | z_j - \frac{1}{2}, z_j + \frac{1}{2}) \quad (3)$$

$$\text{with } \mathbf{y} = g_a(\mathbf{x}; \phi_g), \hat{\mathbf{y}} = Q(\mathbf{y}), \mathbf{z} = h_a(\hat{\mathbf{y}}; \phi_h).$$

The rate term represents the expected bits calculated using the entropy models of $p_{\tilde{\mathbf{y}}|\tilde{\mathbf{z}}}$ and $p_{\tilde{\mathbf{z}}}$. Note that $p_{\tilde{\mathbf{y}}|\tilde{\mathbf{z}}}$ and $p_{\tilde{\mathbf{z}}}$ are eventually the approximations of $p_{\hat{\mathbf{y}}|\hat{\mathbf{z}}}$ and $p_{\hat{\mathbf{z}}}$, respectively. Equation (4) represents the entropy model for approximating the required bits for $\hat{\mathbf{y}}$. The model is based on the Gaussian model, which not only has the standard deviation parameter σ_i , but also the mu parameter μ_i . The values of μ_i and σ_i are estimated from the two types of given contexts based on the function f , the distribution estimator, in a deterministic manner. The two types of contexts, bit-consuming and bit-free contexts, for estimating the distribution of a certain representation are denoted as \mathbf{c}'_i and \mathbf{c}''_i . E' extracts \mathbf{c}'_i from \mathbf{c}' , the result of transform h_s . In contrast to \mathbf{c}'_i , no additional bit allocation is required for \mathbf{c}''_i . Instead, we simply utilize the known (already entropy-coded or decoded) subset of $\hat{\mathbf{y}}$, denoted as $\langle \hat{\mathbf{y}} \rangle$. Here, \mathbf{c}''_i is extracted from $\langle \hat{\mathbf{y}} \rangle$ by the extractor E'' . We assume that the entropy coder and the decoder sequentially process \hat{y}_i in the same specific order, such as with raster scanning, and thus $\langle \hat{\mathbf{y}} \rangle$ given to the encoder and decoder can always be identical when processing the same \hat{y}_i . A formal expression of this is as follows:

$$p_{\tilde{\mathbf{y}}|\tilde{\mathbf{z}}}(\tilde{\mathbf{y}} | \tilde{\mathbf{z}}, \theta_h) = \prod_i \left(\mathcal{N}(\mu_i, \sigma_i^2) * \mathcal{U}(-\frac{1}{2}, \frac{1}{2}) \right) (\tilde{y}_i) \quad (4)$$

$$\text{with } \mu_i, \sigma_i = f(\mathbf{c}'_i, \mathbf{c}''_i),$$

$$\mathbf{c}'_i = E'(h_s(\tilde{\mathbf{z}}; \theta_h), i),$$

$$\mathbf{c}''_i = E''(\langle \hat{\mathbf{y}} \rangle, i),$$

$$\tilde{\mathbf{z}} = Q(\mathbf{z})$$

In the case of $\hat{\mathbf{z}}$, a simple entropy model is used. We assumed that the model follows zero-mean Gaussian distributions which have a trainable σ . Note that $\hat{\mathbf{z}}$ is regarded as side information and it contributes a very small amount of the total bit-rate, as described by Ballé et al. (2018), and thus we use this simpler version of the entropy model, rather than a more complex model, for end-to-end optimization over all parameters of the proposed method:

$$p_{\tilde{\mathbf{z}}}(\tilde{\mathbf{z}}) = \prod_j \left(\mathcal{N}(0, \sigma_j^2) * \mathcal{U}(-\frac{1}{2}, \frac{1}{2}) \right) (\tilde{z}_j) \quad (5)$$




Figure 3: Encoder-decoder model of the proposed method. The left block represents the encoder side, whereas the right block represents the decoder side. The small icons between them represent the entropy-coded bitstreams. EC and ED represent entropy coding and entropy decoding, and $U|Q$ represents the addition of uniform noise to y or a uniform quantization of y . Noisy representations are used only for training as inputs to the entropy models, and are illustrated with the dotted lines.

Note that actual entropy coding or decoding processes are not necessarily required for training or encoding because the rate term is not the amount of real bits, but an estimation calculated from the entropy models, as mentioned previously. We calculate the distortion term using the mean squared error (MSE)¹, assuming that $p_{x|\hat{y}}$ follows a Gaussian distribution as a widely used distortion metric.

3 ENCODER-DECODER MODEL

This section describes the basic structure of the proposed encoder-decoder model. On the encoder side, an input image is transformed into latent representations, quantized, and then entropy-coded using the trained entropy models. In contrast, the decoder first applies entropy decoding with the same entropy models used for the encoder, and reconstructs the image from the latent representations, as illustrated in figure 3. It is assumed that all parameters that appear in this section were already trained. The structure of the encoder-decoder model basically includes g_a and g_s in charge of the transform of x into y and its inverse transform, respectively. The transformed y is uniformly quantized into \hat{y} by rounding. Note that, in the case of approaches based on the entropy models, unlike traditional codecs, tuning the quantization steps is usually unnecessary because the scales of the representations are optimized together through training. The other components between g_a and g_s carry out the role of entropy coding (or decoding) with the shared entropy models and underlying context preparation processes. More specifically, the entropy model estimates the distribution of each \hat{y}_i individually, in which μ_i and σ_i are estimated with the two types of given contexts, c_i' and c_i'' . Among these contexts, c_i' can be viewed as side information, which requires an additional bit allocation. To reduce the required bit-rate for carrying c_i' , the latent representation z , transformed from \hat{y} , is quantized and entropy-coded by its own entropy model, as specified in section 2.3. On the other hand, c_i'' is extracted from $\langle \hat{y} \rangle$, without any additional bit allocation. Note that $\langle \hat{y} \rangle$ varies as the entropy coding or decoding progresses, but is always identical for processing the same \hat{y}_i in both the encoder and decoder, as described in 2.3. The parameters of h_s and the entropy models are simply shared by both the encoder and the decoder. Note that the inputs to the entropy models during training are the noisy representations, as illustrated with the dotted line in figure 3, to allow the entropy model to approximate the probability mass functions of the discrete representations.

¹We also provide supplemental experiment results in which an MS-SSIM (Wang et al. (2003)) based distortion term is used for optimization. To calculate the distortion term, we multiplied 1-MS-SSIM by 3000.

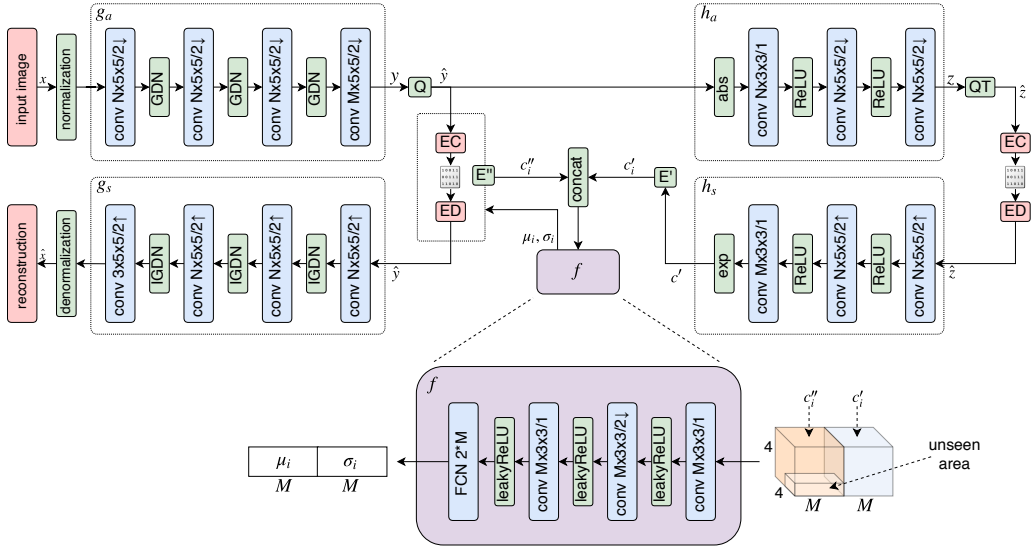


Figure 4: Implementation of the proposed method. We basically use the convolutional autoencoder structure, and the distribution estimator f is also implemented using convolutional neural networks. The notations of the convolutional layer follow Ballé et al. (2018): the number of filters \times filter height \times filter width / the downscale or upscale factor, where \uparrow and \downarrow denote the up and downscaling, respectively. For up or downscaling, we use the transposed convolution. For the networks, input images are normalized into a scale between -1 and 1.

4 EXPERIMENTS

4.1 IMPLEMENTATION

We use a convolutional neural networks to implement the analysis transform and the synthesis transform functions, g_a , g_s , h_a , and h_s . The structures of the implemented networks follow the same structures of Ballé et al. (2018), except that we use the exponentiation operator instead of an absolute operator at the end of h_s . Based on Ballé et al. (2018)’s structure, we added the components to estimate the distribution of each \hat{y}_i , as shown in figure 4. Herein, we represent a uniform quantization (round) as "Q," entropy coding as "EC," and entropy decoding as "ED." The distribution estimator is denoted as f , and is also implemented using the convolutional layers which takes channel-wise concatenated c_i' and c_i'' as inputs and provides estimated μ_i and σ_i as results. Note that the same c_i' and c_i'' are shared for all \hat{y}_i s located at the same spatial position. In other words, we let E' extract all spatially adjacent elements from c' across the channels to retrieve c_i' and likewise let E'' extract all adjacent known elements from $\langle \hat{y} \rangle$ for c_i'' . This could have the effect of capturing the remaining correlations among the different channels. In short, when M is the total number of channels of \hat{y} , we let f estimate all M distributions of \hat{y}_i s, which are located at the same spatial position, using only a single step, thereby allowing the total number of estimations to be reduced. Furthermore, the parameters of f are shared for all spatial positions of \hat{y} , and thus only one trained f per λ is necessary to process any sized images. In the case of training, however, collecting the results from the all spatial positions to calculate the rate term becomes a significant burden, despite the simplifications mentioned above. To reduce this burden, we designate a certain number (32 and 16 for the base model and the hybrid model, respectively) of random spatial points as the representatives per training step, to calculate the rate term easily. Note that we let these random points contribute solely to the rate term, whereas the distortion is still calculated over all of the images.

Because \mathbf{y} is a three-dimensional array in our implementation, index i can be represented as three indexes, k , l , and m , representing the horizontal index, the vertical index, and the channel index, respectively. When the current position is given as (k, l, m) , E' extracts $c'_{[k-2\dots k+1],[l-3\dots l],[1\dots M]}$ as c_i' , and E'' extracts $\langle \hat{y} \rangle_{[k-2\dots k+1],[l-3\dots l],[1\dots M]}$ as c_i'' , when $\langle \hat{y} \rangle$ represents the known area of \hat{y} . Note that we filled in the unknown area of $\langle \hat{y} \rangle$ with zeros, to maintain the dimensions of $\langle \hat{y} \rangle$

identical to those of $\hat{\mathbf{y}}$. Consequently, $\mathbf{c}''_{i \in [3..4], 4, [1..M]}$ are always padded with zeros. To keep the dimensions of the estimation results to the inputs, the marginal areas of \mathbf{c}' and $\langle \hat{\mathbf{y}} \rangle$ are also set to zeros. Note that when training or encoding, \mathbf{c}''_i can be extracted using simple $4 \times 4 \times M$ windows and binary masks, thereby enabling parallel processing, whereas a sequential reconstruction is inevitable for decoding.

Another implementation technique used to reduce the implementation cost is combining the lightweight entropy model with the proposed model. The lightweight entropy model assumes that the representations follow a zero-mean Gaussian model with the estimated standard deviations, which is very similar with Ballé et al. (2018)’s approach. We utilize this hybrid approach for the top four cases, in bit-rate descending order, of the nine λ configurations, based on the assumption that for the higher-quality compression, the number of sparse representations having a very low spatial dependency increases, and thus a direct scale estimation provides sufficient performance for these added representations. For implementation, we separate the latent representation \mathbf{y} into two parts, \mathbf{y}_1 and \mathbf{y}_2 , and two different entropy models are applied for them. Note that the parameters of g_a , g_s , h_a , and h_s are shared, and all parameters are still trained together. The detailed structure and experimental settings are described in appendix 6.1.

The number of parameters N and M are set to 128 and 192, respectively, for the five λ configurations for lower bit-rates, whereas 2-3 times more parameters, described in appendix 6.1, are used for the four λ configurations for higher bit-rates. Tensorflow and Python were used to setup the overall network structures, and for the actual entropy coding and decoding using the estimated model parameters, we implemented an arithmetic coder and decoder, for which the source code of the "Reference arithmetic coding" project² was used as the base code.

4.2 EXPERIMENTAL ENVIRONMENTS

We optimized the networks using two different types of distortion terms, one with MSE and the other with MS-SSIM. For each distortion type, the average bits per pixel (BPP) and the distortion, PSNR and MS-SSIM, over the test set are measured for each of the nine λ configurations. Therefore, a total of 18 networks are trained and evaluated within the experimental environments, as explained below:

- For training, we used 256×256 patches extracted from 32,420 randomly selected YFCC100m (Thomee et al. (2016)) images. We extracted one patch per image, and the extracted regions were randomly chosen. Each batch consists of eight images, and 1M iterations of the training steps were conducted, applying the ADAM optimizer (Kingma & Ba (2015)). We set the initial learning rate to 5×10^{-5} , and reduced the rate by half every 50,000 iterations for the last 200,000 iterations. Note that, in the case of the four λ configurations for high bpp, in which the hybrid entropy model is used, 1M iterations of pre-training steps were conducted using the learning rate of 1×10^{-5} . Although we previously indicated that the total loss is the sum of R and λD for a simple explanation, we tuned the balancing parameter λ in a similar way as Theis et al. (2017), as indicated in equation (6). We used the λ parameters ranging from 0.01 to 0.5.

$$\mathcal{L} = \frac{\lambda}{W_y \cdot H_y \cdot 256} R + \frac{1 - \lambda}{1000} D. \quad (6)$$

- For the evaluation, we measured the average BPP and average quality of the reconstructed images in terms of the PSNR and MS-SSIM over 24 PNG images of the Kodak PhotoCD image dataset (Kodak, 1993). Note that we represent the MS-SSIM results in the form of decibels, as in Ballé et al. (2018), to increase the discrimination.

4.3 EXPERIMENTAL RESULTS

We compared the test results with other previous methods, including traditional codecs such as BPG and JPEG2000, as well as previous ANN-based approaches such as Theis et al. (2017) and Ballé

²Nayuki, "Reference arithmetic coding," <https://github.com/nayuki/Reference-arithmetic-coding>, Copy-right © 2018 Project Nayuki. (MIT License).

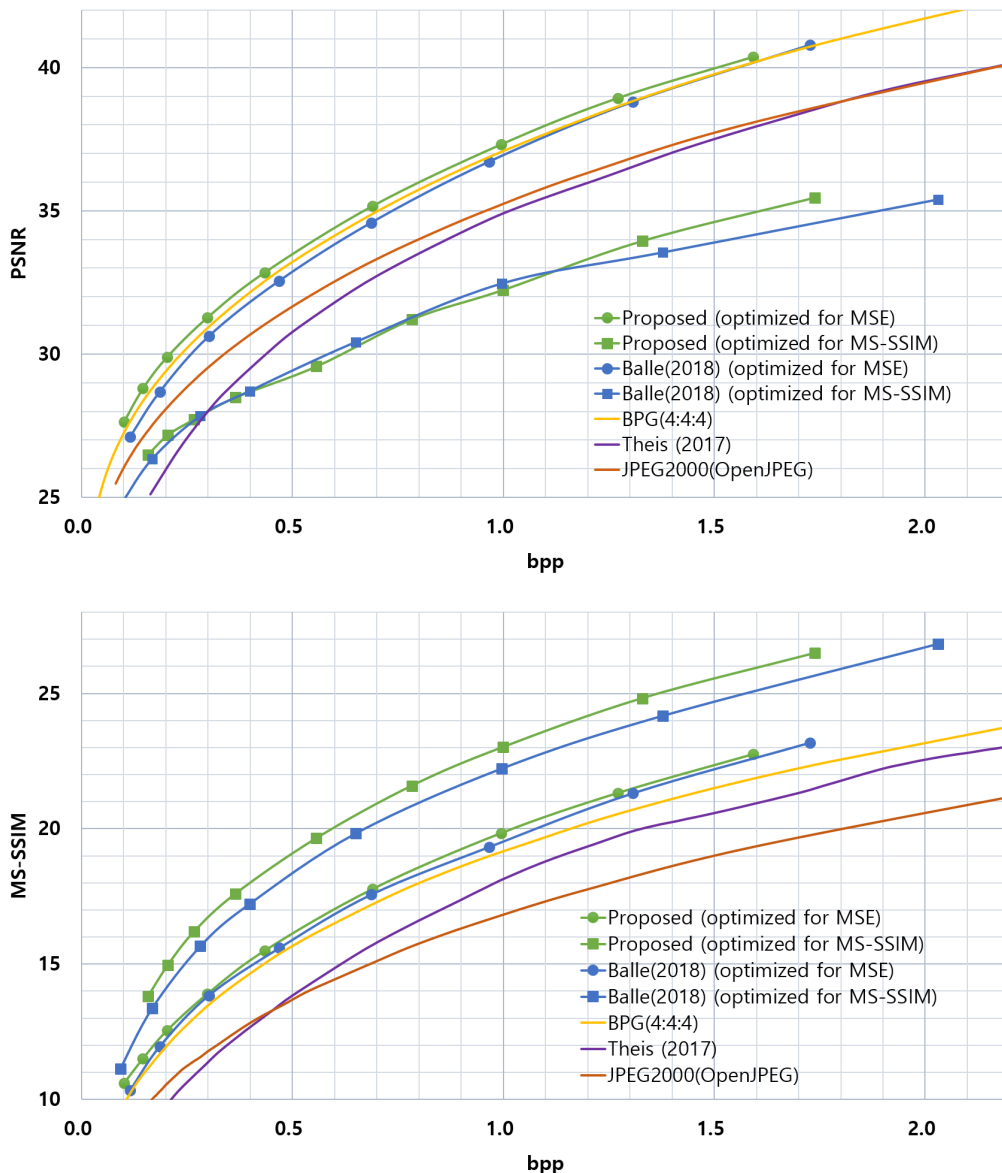


Figure 5: Rate–distortion curves of the proposed method and competitive methods. The top plot represents the PSNR values as a result of changes in bpp, whereas the bottom plot shows MS-SSIM values in the same manner. Note that MS-SSIM values are converted into decibels ($-10 \log_{10}(1 - \text{MS-SSIM})$) for differentiating the quality levels, in the same manner as in Ballé et al. (2018).

et al. (2018). Because two different quality metrics are used, the results are presented with two separate plots. As shown in figure 5, our methods outperform all other previous methods in both metrics. In particular, our models not only outperform Ballé et al. (2018)’s method, which is believed to be a state-of-the-art ANN-based approach, but we also obtain better results than the widely used conventional image codec, BPG.

More specifically, the compression gains in terms of the BD-rate of PSNR over JPEG2000, Ballé et al. (2018)’s approach (MSE-optimized), and BPG are 34.08%, 11.97%, and 6.85%, respectively. In the case of MS-SSIM, we found wider gaps of 68.82%, 13.93%, and 49.68%, respectively. Note that we achieved significant gains over traditional codecs in terms of MS-SSIM, although this might be because the dominant target metric of the traditional codec developments have been the PSNR.

In other words, they can be viewed as a type of MSE-optimized codec. Even when setting aside the case of MS-SSIM, our results can be viewed as one concrete evidence supporting that ANN-based image compression can outperform the existing traditional image codecs in terms of the compression performance. Supplemental image samples are provided in appendix 6.3.

5 DISCUSSION

Based on previous ANN-based image compression approaches utilizing entropy models (Ballé et al., 2017; Theis et al., 2017; Ballé et al., 2018), we extended the entropy model to exploit two different types of contexts. These contexts allow the entropy models to more accurately estimate the distribution of the representations with a more generalized form having both mean and standard deviation parameters. Based on the evaluation results, we showed the superiority of the proposed method. The contexts we utilized are divided into two types. One is a sort of free context, containing the part of the latent variables known to both the encoder and the decoder, whereas the other is the context, which requires additional bit allocation. Because the former is a generally used context in a variety of codecs, and the latter was already verified to help compression using Ballé et al. (2018)’s approach, our contributions are not the contexts themselves, but can be viewed as providing a framework of entropy models utilizing these contexts.

Although the experiments showed the best results in the ANN-based image compression domain, we still have various studies to conduct to further improve the performance. One possible way is generalizing the distribution models underlying the entropy model. Although we enhanced the performance by generalizing the previous entropy models, and have achieved quite acceptable results, the Gaussian-based entropy models apparently have a limited expression power. If more elaborate models, such as the non-parametric models of Ballé et al. (2018) or Oord et al. (2016), are combined with the context-adaptivity proposed in this paper, they would provide better results by reducing the mismatch between the actual distributions and the approximation models. Another possible way is improving the level of the contexts. Currently, our methods only use low-level representations within very limited adjacent areas. However, if the sufficient capacity of the networks and higher-level contexts are given, a much more accurate estimation could be possible. For instance, if an entropy model understands the structures of human faces, in that they usually have two eyes, between which a symmetry exists, the entropy model could approximate the distributions more accurately when encoding the second eye of a human face by referencing the shape and position of the first given eye. As is widely known, various generative models (Goodfellow et al., 2014; Radford et al., 2016; Zhao et al., 2017) learn the distribution $p(x)$ of the images within a specific domain, such as human faces or bedrooms. In addition, various in-painting methods (Pathak et al., 2016; Yang et al., 2017; Yeh et al., 2017) learn the conditional distribution $p(x | context)$ when the viewed areas are given as *context*. Although these methods have not been developed for image compression, hopefully such high-level understandings can be utilized sooner or later. Furthermore, the contexts carried using side information can also be extended to some high-level information such as segmentation maps or any other information that helps with compression. Segmentation maps, for instance, may be able to help the entropy models estimate the distribution of a representation discriminatively according to the segment class the representation belongs to.

Traditional codecs have a long development history, and a vast number of hand-crafted heuristics have been stacked thus far, not only for enhancing compression performance, but also for compromising computational complexities. Therefore, ANN-based image compression approaches may not provide satisfactory solutions as of yet, when taking their high complexity into account. However, considering its much shorter history, we believe that ANN-based image compression has much more potential and possibility in terms of future extension. Although we remain a long way from completion, we hope the proposed context-adaptive entropy model will provide an useful contribution to this area.

ACKNOWLEDGMENTS

This work was supported by Institute for Information & communications Technology Promotion(IITP) grant funded by the Korea government (MSIT) (No. 2017-0-00072, Development of Audio/Video Coding and Light Field Media Fundamental Technologies for Ultra Realistic Tera-media).

REFERENCES

- Eirikur Agustsson, Michael Tschannen, Fabian Mentzer, Radu Timofte, and Luc Van Gool. Generative adversarial networks for extreme learned image compression. *arXiv preprint arXiv:1804.02958*, 2018. URL <http://arxiv.org/abs/1804.02958>.
- Johannes Ballé, Valero Laparra, and Eero P. Simoncelli. End-to-end optimized image compression. In *the 5th Int. Conf. on Learning Representations*, 2017. URL <http://arxiv.org/abs/1611.01704>.
- Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. In *the 6th Int. Conf. on Learning Representations*, 2018. URL <http://arxiv.org/abs/1802.01436>.
- Fabrice Bellard. Bpg image format, 2014. URL <http://bellard.org/bpg/>.
- Chao Dong, Yubin Deng, Chen Change Loy, and Xiaoou Tang Tang. Compression artifacts reduction by a deep convolutional network. In *IEEE International Conference on Computer Vision (ICCV)*, 2015. URL <http://arxiv.org/abs/1504.06993>.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (eds.), *Advances in Neural Information Processing Systems 27*, pp. 2672–2680. Curran Associates, Inc., 2014. URL <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>.
- ISO/IEC 23008-2, ITU-T H.265. Information technology – high efficiency coding and media delivery in heterogeneous environments – part 2: High efficiency video coding. Standard, ISO/IEC, 2013.
- Nick Johnston, Damien Vincent, David Minnen, Michele Covell, Saurabh Singh, Troy Chinen, Sung Jin Hwang, Joel Shor, and George Toderici. Improved lossy image compression with priming and spatially adaptive bit rates for recurrent networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *the 3rd Int. Conf. on Learning Representations*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *the 2nd Int. Conf. on Learning Representations*, 2014. URL <http://arxiv.org/abs/1312.6114>.
- Eastman Kodak. Kodak lossless true color image suite (photocd pcd0992), 1993. URL <http://r0k.us/graphics/kodak/>.
- Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *the 33rd International Conference on Machine Learning*, 2016. URL <http://arxiv.org/abs/1601.06759>.
- Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alexei A. Efros. Context encoders: Feature learning by inpainting. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *the 4th Int. Conf. on Learning Representations*, 2016. URL <http://arxiv.org/abs/1511.06434>.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *the 31st Int. Conf. on Machine Learning*, 2014. URL <http://arxiv.org/abs/1401.4082>.
- Oren Rippel and Lubomir Bourdev. Real-time adaptive image compression. In *International Conference on Machine Learning*, 2017. URL <http://arxiv.org/abs/1705.05823>.

- Shibani Santurkar, David M. Budden, and Nir Shavit. Generative compression. In *The 33rd Picture Coding Symposium*, 2018. URL <http://arxiv.org/abs/1703.01467>.
- Pavel Svoboda, Michal Hradis, David Barina, and Pavel Zemcík. Compression artifacts removal using convolutional neural networks. In *WSCG 24th International Conference on Computer Graphics, Visualization and Computer Vision*, 2016. URL <http://arxiv.org/abs/1605.00366>.
- Lucas Theis, Wenzhe Shi, Andrew Cunningham, and Ferenc Huszár. Lossy image compression with compressive autoencoders. In *the 5th Int. Conf. on Learning Representations*, 2017. URL <http://arxiv.org/abs/1703.00395>.
- Bart Thomee, David A. Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016.
- 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. on Computer Vision and Pattern Recognition (CVPR)*, 2017. doi: 10.1109/CVPR.2017.577. URL <http://arxiv.org/abs/1608.05148>.
- Zhou Wang, Eero P. Simoncelli, and Alan C. Bovik. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers*, 2003. doi: 10.1109/ACSSC.2003.1292216.
- Chao Yang, Xin Lu, Zhe Lin, Eli Shechtman, Oliver Wang, and Hao Li. High-resolution image inpainting using multi-scale neural patch synthesis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- Raymond A. Yeh, Chen Chen, Teck-Yian Lim, Alexander G. Schwing, Mark Hasegawa-Johnson, and Minh N. Do. Semantic image inpainting with deep generative models. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6882–6890. IEEE Computer Society, 2017.
- Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017. doi: 10.1109/TIP.2017.2662206. URL <http://arxiv.org/abs/1608.03981>.
- Junbo Jake Zhao, Michaël Mathieu, and Yann LeCun. Energy-based generative adversarial network. In *the 5th Int. Conf. on Learning Representations*, 2017. URL <http://arxiv.org/abs/1609.03126>.




Figure 6: The structure of the hybrid network for higher bit-rate environments. The same notations as in the figure 4 are used. The representation \mathbf{y} is divided into two parts and quantized. One of the resulting parts, $\hat{\mathbf{y}}_1$, is encoded using the proposed model, whereas the other, $\hat{\mathbf{y}}_2$, is encoded using a simpler model in which only the standard deviations are estimated using side information. The detailed structure of the proposed model is illustrated in figure 4. All concatenation and split operations are performed in a channel-wise manner.

6 APPENDIX

6.1 HYBRID NETWORK FOR HIGHER BIT-RATE COMPRESSIONS

We combined the lightweight entropy model with the context-adaptive entropy model to reduce the implementation costs for high-bpp configurations. The lightweight model exploits the scale (standard deviation) estimation, assuming that the PMF approximations of the quantized representations follow zero-mean Gaussian distributions convolved with a standard uniform distribution.

Figure 6 illustrates the network structure of this hybrid network. The representation \mathbf{y} is split channel-wise into two parts, \mathbf{y}_1 and \mathbf{y}_2 , which have M_1 and M_2 channels, respectively, and is then quantized. Here, $\hat{\mathbf{y}}_1$ is entropy coded using the proposed entropy model, whereas $\hat{\mathbf{y}}_2$ is coded with the lightweight entropy model. The standard deviations of $\hat{\mathbf{y}}_2$ are estimated using h_a and h_s . Unlike the context-adaptive entropy model, which uses the results of h_a (\hat{c}') as the input source to the estimator f , the lightweight entropy model retrieves the estimated standard deviations from h_a directly. Note that h_a takes the concatenated $\hat{\mathbf{y}}_1$ and $\hat{\mathbf{y}}_2$ as input, and h_s generates \hat{c}' as well as σ_2 , at the same time.

The total loss function also consists of the rate and distortion terms, although the rate is divided into three parts, each of which is for $\hat{\mathbf{y}}_1$, $\hat{\mathbf{y}}_2$, and $\hat{\mathbf{z}}$, respectively. The distortion term is the same as before, but note that $\hat{\mathbf{y}}$ is the channel-wise concatenated representation of $\hat{\mathbf{y}}_1$ and $\hat{\mathbf{y}}_2$:

$$\mathcal{L} = R + \lambda D \quad (7)$$

$$\text{with } R = \mathbb{E}_{\mathbf{x} \sim p_{\mathbf{x}}} \mathbb{E}_{\tilde{\mathbf{y}}_1, \tilde{\mathbf{y}}_2, \tilde{\mathbf{z}} \sim q} \left[-\log p_{\tilde{\mathbf{y}}_1 | \tilde{\mathbf{z}}}(\tilde{\mathbf{y}}_1 | \tilde{\mathbf{z}}) - \log p_{\tilde{\mathbf{y}}_2 | \tilde{\mathbf{z}}}(\tilde{\mathbf{y}}_2 | \tilde{\mathbf{z}}) - \log p_{\tilde{\mathbf{z}}}(\tilde{\mathbf{z}}) \right],$$

$$D = \mathbb{E}_{\mathbf{x} \sim p_{\mathbf{x}}} \left[-\log p_{\mathbf{x} | \hat{\mathbf{y}}}(\mathbf{x} | \hat{\mathbf{y}}) \right]$$

Here, the noisy representations of $\tilde{\mathbf{y}}_1$, $\tilde{\mathbf{y}}_2$, and $\tilde{\mathbf{z}}$ follow a standard uniform distribution, the mean values of which are \mathbf{y}_1 , \mathbf{y}_2 , and \mathbf{z} , respectively. In addition, \mathbf{y}_1 and \mathbf{y}_2 are channel-wise split representations from \mathbf{y} , the results of the transform g_a , and have M_1 and M_2 channels, respectively:

$$\begin{aligned}
q(\tilde{\mathbf{y}}_1, \tilde{\mathbf{y}}_2, \tilde{\mathbf{z}} \mid \mathbf{x}, \phi_g, \phi_h) &= \prod_i \mathcal{U}(\tilde{y}_{1i} \mid y_{1i} - \frac{1}{2}, y_{1i} + \frac{1}{2}) \cdot \\
&\quad \prod_j \mathcal{U}(\tilde{y}_{2j} \mid y_{2j} - \frac{1}{2}, y_{2j} + \frac{1}{2}) \cdot \\
&\quad \prod_k \mathcal{U}(\tilde{z}_k \mid z_k - \frac{1}{2}, z_k + \frac{1}{2}) \tag{8} \\
&\text{with } \mathbf{y}_1, \mathbf{y}_2 = S(g_a(\mathbf{x}; \phi_g)), \hat{\mathbf{y}} = Q(\mathbf{y}_1) \oplus Q(\mathbf{y}_2), \mathbf{z} = h_a(\hat{\mathbf{y}}; \phi_h).
\end{aligned}$$

The rate term for $\hat{\mathbf{y}}_1$ is the same model as that of equation (4). Note that σ_2 does not contribute here, but does contribute to the model for $\hat{\mathbf{y}}_2$:

$$\begin{aligned}
p_{\tilde{\mathbf{y}}_1|\hat{\mathbf{z}}}(\tilde{\mathbf{y}}_1 \mid \hat{\mathbf{z}}, \theta_h) &= \prod_i \left(\mathcal{N}(\mu_{1i}, \sigma_{1i}^2) * \mathcal{U}(-\frac{1}{2}, \frac{1}{2}) \right) (\tilde{y}_{1i}) \tag{9} \\
&\text{with } \mu_{1i}, \sigma_{1i} = f(\mathbf{c}'_i, \mathbf{c}''_i), \\
&\quad \mathbf{c}'_i = E'(\mathbf{c}', i), \\
&\quad \mathbf{c}''_i = E''(\langle \hat{\mathbf{y}}_1, i \rangle), \\
&\quad \mathbf{c}', \sigma_2 = S(h_s(\hat{\mathbf{z}}; \theta_h))
\end{aligned}$$

The rate term for $\hat{\mathbf{y}}_2$ is almost the same as Ballé et al. (2018), except that noisy representations are only used as the inputs to the entropy models for training, and not for the conditions of the models.

$$p_{\tilde{\mathbf{y}}_2|\hat{\mathbf{z}}}(\tilde{\mathbf{y}}_2 \mid \hat{\mathbf{z}}, \theta_h) = \prod_j \left(\mathcal{N}(0, \sigma_{2j}^2) * \mathcal{U}(-\frac{1}{2}, \frac{1}{2}) \right) (\tilde{y}_{2j}) \tag{10}$$

The model of \mathbf{z} is the same as in equation (5). For implementation, we used this hybrid structure for the top-four configurations in bit-rate descending order. We set N , M_1 , and M_2 to 400, 192, and 408, respectively, for the top-two configurations, and to 320, 192, and 228, respectively, for the next two configurations.

In addition, we measured average execution time per image, spent for encoding and decoding Kodak PhotoCD image dataset (Kodak, 1993), to clarify benefit of the hybrid model. The test was conducted under CPU environments, Intel i9-7900X. Note that we ignored time for actual entropy coding because all models with the same values of N and M spend the same amount of time for entropy coding. As shown in figure 7, the hybrid models clearly reduced execution time of the models. Setting N and M to 320 and 420, respectively, we obtained 46.83% of speed gain. With the higher number of parameters, 400 of N and 600 of M , we obtained 57.28% of speed gain.

6.2 TEST RESULTS OF THE MODELS TRAINED USING DIFFERENT TYPES OF REPRESENTATIONS

In this section, we provide test results of the two models, the proposed model trained using discrete representations as inputs to the synthesis transforms, g_s and h_s , and the same model but trained using noisy representations following the training process of Ballé et al. (2018)’s approach. In detail, in training phase of the proposed model, we used quantized representations $\hat{\mathbf{y}}$ and $\hat{\mathbf{z}}$ as inputs to the transforms g_s and h_s , respectively, to ensure the same conditions of training and testing phases. On the other hand, for training the compared model, representations $\tilde{\mathbf{y}}$ and $\tilde{\mathbf{z}}$ are used as inputs to the transforms. An additional change of the proposed model is using $\hat{\mathbf{y}}$, instead of \mathbf{y} , as inputs to h_a , but note that this has nothing to do with the mismatches between training and testing. We used them to match inputs to h_a to targets of model estimation via f . As shown in figure 8, the proposed model, trained using discrete representations, was 5.94% better than the model trained using noisy representations, in terms of the BD-rate of PSNR. Compared with Ballé et al. (2018)’s approach, the performance gains of the two models, trained using discrete representations and noisy representations, were 11.97% and 7.20%, respectively.

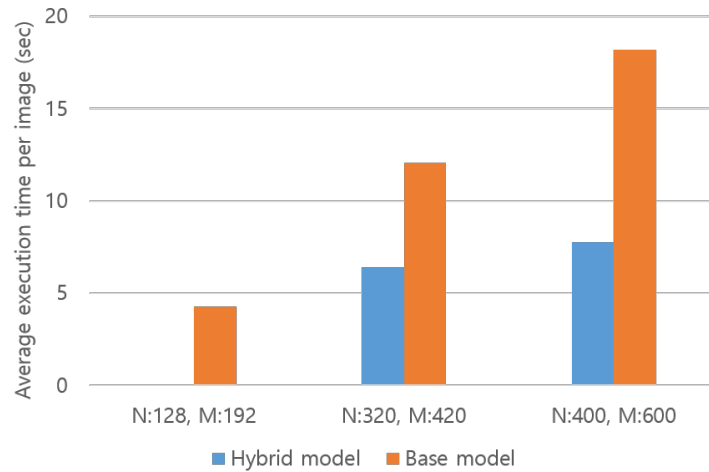


Figure 7: Comparison results of execution time between the base model and hybrid model




Figure 8: Evaluation results of the models trained using noisy/discrete representations as input to synthesis transforms

6.3 SAMPLES OF THE EXPERIMENTS

In this section, we provide a few more supplemental test results. Figures 9, 10, and 11 show the results of the MSE optimized version of our method, whereas figures 12 and 13 show the results of the MS-SSIM optimized version. All figures include a ground truth image and the reconstructed images for our method, BPG, and Ballé et al. (2018)’s approach, in the clockwise direction.



Figure 9: Sample test results. Top left, ground truth; top right, our method (MSE optimized; bpp, 0.2040; PSNR, 32.2063); bottom left, BPG (bpp, 0.2078; PSNR, 32.0406); bottom right, Ballé et al. (2018)'s approach (MSE optimized; bpp, 0.2101; PSNR, 31.7054)



Figure 10: Sample test results. Top left, ground truth; top right, our method (MSE optimized; bpp, 0.1236; PSNR, 32.4284); bottom left, BPG (bpp, 0.1285; PSNR, 32.0444); bottom right, Ballé et al. (2018)'s approach (MSE optimized, bpp, 0.1229; PSNR, 31.0596)



Figure 11: Sample test results. Top left, ground truth; top right, our method (MSE optimized; bpp, 0.1501; PSNR, 34.7103); bottom left, BPG (bpp, 0.1477; PSNR, 33.9623); bottom right, Ballé et al. (2018)'s approach (MSE optimized; bpp, 0.1520; PSNR, 34.0465)




Figure 12: Sample test results. Top left, ground truth; top right, our method (MS-SSIM optimized; bpp, 0.2507; MS-SSIM, 0.9740); bottom left, BPG (bpp, 0.2441; MS-SSIM, 0.9555); bottom right, Ballé et al. (2018)'s approach (MS-SSIM optimized; bpp, 0.2101; MS-SSIM, 0.9705)



Figure 13: Sample test results. Top left, ground truth; top right, our method (MS-SSIM optimized; bpp, 0.2269; MS-SSIM, 0.9810); bottom left, BPG (bpp, 0.2316; MS-SSIM, 0.9658); bottom right, Ballé et al. (2018)'s approach (MS-SSIM optimized; bpp, 0.2291; MS-SSIM, 0.9786)