Learning Context-Based Non-local Entropy Modeling for Image Compression

Mu Li, Kai Zhang, Wangmeng Zuo, *Senior Member, IEEE*, Radu Timofte, *Member, IEEE*, and David Zhang, *Fellow, IEEE*

Abstract—The entropy of the codes usually serves as the rate loss in the recent learned lossy image compression methods. Precise estimation of the probabilistic distribution of the codes plays a vital role in the performance. However, existing deep learning based entropy modeling methods generally assume the latent codes are statistically independent or depend on some side information or local context, which fails to take the global similarity within the context into account and thus hinder the accurate entropy estimation. To address this issue, we propose a non-local operation for context modeling by employing the global similarity within the context. Specifically, we first introduce the proxy similarity functions and spatial masks to handle the missing reference problem in context modeling. Then, we combine the local and the global context via a non-local attention block and employ it in masked convolutional networks for entropy modeling. The entropy model is further adopted as the rate loss in a joint rate-distortion optimization to guide the training of the analysis transform and the synthesis transform network in transforming coding framework. Considering that the width of the transforms is essential in training low distortion models, we finally produce an U-Net block in the transforms to increase the width with manageable memory consumption and time complexity. Experiments on Kodak and Tecnick datasets demonstrate the superiority of the proposed context-based non-local attention block in entropy modeling and the U-Net block in low distortion compression against the existing image compression standards and recent deep image compression models.

Index Terms—Non-local, entropy modeling, U-Net block, image compression.

I. INTRODUCTION

Image compression is a crucial problem in computer science that has been studied for decades. Many image compression standards like JPEG [1] have already been used in our daily life. Nevertheless, the last decade has witnessed a population of artificial intelligence and social media, which brought new challenges for sharing and storing huge amounts of high-definition media. Better image compression methods are still demanding.

Deep neural networks (DNNs) have proven to be effective and get extraordinary results in numerous vision tasks such

Mu Li is with the School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen, Guangdong, 518172, P.R. China (e-mail: limuhit@gmail.com).

Kai Zhang and Radu Timofte are with the Computer Vision Laboratory, ETH Zurich, Zurich, Switzerland.

Wangmeng Zuo is with the School of Computer Science and Technology, Harbin Institute of Technology, Harbin, 150001, China (e-mail: cswmzuo@gmail.com).

David Zhang is with the School of Science and Engineering, The Chinese University of Hong Kong (Shenzhen), Shenzhen Research Institute of Big Data, and Shenzhen Institute of Artificial Intelligence and Robotics for Society, Shenzhen, China, e-mail: (davidzhang@cuhk.edu.cn).

as image restoration [2–5], image quality assessment [6, 7], and image classification [8, 9], which throws light on learning better lossy image compression methods. Deep neural networks are natural transforms. As a result, nearly all of the recent attempts on image compression with DNNs follow the transforming coding framework and are optimized for a rate-distortion trade-off. A transforming coding framework usually consists of three components, i.e., transform, quantization and entropy coding. To be specific, transform maps the input images to code representations; quantization transforms the representation into discrete counterparts; entropy coding compresses the quantized codes into the bitstream in a lossless manner.

Previous transforms [10] are linear, invertible and fixed for all bit rates. Distortion only arises from the quantization. As for the recent deep lossy image compression methods, transforms are modeled by learning non-linear and powerful DNNs. With the quantization operation and the structure of the network, the learned transforms are generally non-invertible, which tends to encourage discarding the perceptual negligible information for better visual quality at low bit rates. However, the performance is usually limited at high bit rates where the distortion is required to be as small as possible. Empirically, the width of the transforms, i.e., the number of feature maps at each layer, is effective in reducing the information losing brought by the transforms. Many deep image compression methods [11–13] suggest adopting wider transforms at high bit rates. For a deep network, the growing of width will inevitably increase the computational complexity and GPU memory consumption. We suggest adopting an U-Net like block in the transforms which could help reduce the time complexity and memory usage in the transforms. With the paired down-sampling, up-sampling operations and skip connection, the U-Net like architecture can not only speed up the transforms but also combine the information in different scales and facilitate the information propagation.

The learned image compression methods are usually modeled as a rate-distortion optimization problem where entropy modeling of the codes plays an important role. According to Shannon's source coding theorem [14], given a sequence of codes $\boldsymbol{y} = \{y_0, \ldots, y_M\}$, the optimal code length of $\boldsymbol{y}$ should be $\lceil -\log_2 P(\boldsymbol{y}) \rceil$ with the codebook constructed by binary symbols. Thus, building accurate discrete probability distribution functions for the codes is essential in determining the compression rate. Most entropy coding schemes directly assume the codes in $\boldsymbol{y}$ are independent and follow the same marginal distribution, resulting in a code length of

$\lceil -\sum_{i=0}^{M} \log_2 P(y_i) \rceil$. Generally, a better approximation could be given with the chain rule when considering the context.

$$P(\boldsymbol{y}) \approx \prod_{i=0}^{M} P(y_i | y_0, \dots, y_{i-1}), \quad (1)$$

where the context of y_i is defined as all the codes coded before it in $\boldsymbol{y}$. One typical application of the context in entropy modeling is the context-based adaptive binary arithmetic coding (CABAC) [15] used in H.264/AVC, which gets a clear performance improvement over previous image compression standards by considering only two nearest codes in the context. However, the traditional context-based entropy modeling based on histograms fails to cover more codes. To employ more codes in the context and the shared computation of fully convolutional network, a set of masked convolutional networks including PixelCNN [13], MCN [16], and CCN [17] are proposed for entropy modeling in recent lossy image compression methods. Nevertheless, all the entropy models based on CNNs have a limited receptive filed and only consider the local context within the receptive field in entropy modeling. Global information and content similarity in the context are ignored.

To explore global information, we introduce a context-based non-local operation for entropy modeling. In image denoising, non-local means [18] predicts the target pixel with a weighted sum on the whole image, where the weights are produced by the content similarity between the target pixel and the others. Inspired by the non-local means, we regress a target code, $y_r(p, q)$, with a weighted sum over all the codes in the r-th channel within the context. Here, $y_r(p, q)$ is a code from the 3D code block $\boldsymbol{y}$ at the position (p, q) in the r-th channel. However, due to the definition of context, the target code itself is unknown in the calculation of the weights. Without the reference, it is unable to evaluate the similarity and produce the weights. We call this challenge missing reference problem. Considering that $y_r(p, q)$ and the codes from the vector $\boldsymbol{\xi}_r(p, q) = \{y_0(p, q), \dots, y_{r-1}(p, q)\}$ are produced by the same image patch, there should be a close relationship between them. We alternatively evaluate the similarity between the target code $y_r(p, q)$ and $y_r(u, v)$ as a weighted L2 distance between $\boldsymbol{\xi}_r(p, q)$ and $\boldsymbol{\xi}_r(u, v)$. An attention block is introduced to combine the local representations produced by CCNs and the global representations from the non-local operation for context-based entropy modeling.

The context-based non-local attention block is further adopted in CCNs to estimate the entropy of the lossy image compression method. In detail, the probability of each code is modeled on a mixture of Gaussian (MoG) distribution whose parameters are predicted by the non-local CCNs. We jointly optimize transforms, *i.e.*, analysis transform and synthesis transform built on U-Net blocks and the context-based non-local entropy model with respect to the trade-off of rate-distortion performance in an end-to-end manner. Experiments on the Kodak and Tecnick datasets show that the proposed method can outperform state-of-art lossy image compression standards.

II. RELATED WORK

A. DNN-Based Lossy Image Compression Methods

Recent DNN-based lossy image compression methods are optimized for the rate-distortion trade-off, where the rate loss is usually modeled as the entropy of the codes. Previous methods [11, 19–21] only adopt DNNs for transforms and suppose that all the codes are i.i.d. and follow the same probability distribution function (PDF) for easy entropy models. With a lack of accurate estimation of the entropy, the performance of such methods is limited. Arguably speaking, the code representations generated from the image via a highly nonlinear analysis transform still exhibit strong statistical redundancies [12]. Therefore, several attempts have been made by estimating the conditional PDFs of the codes based on extra information, *i.e.*, context [13, 16, 17, 22, 23] and hyperprior [12, 13]. Specially, for context-based entropy modeling, DNN-based entropy models such as RNN [24], LSTM [25] are shown to be powerful in modeling the context of sequential data in natural language processing. PixelRNN [26] and PixelCNN [27] are further proposed for modeling long-range dependency among pixels and employ much larger context for processing image data. Taking both of the context and effeciency into account, a set of masked convolutional networks, such as PixelCNN [13], MCN [16], and CCN [17], are proposed for modeling the entropy of the codes in lossy image compression. In the following, we will give an overview of some representative DNN-based lossy image compression methods.

Toderici *et al*. [28] proposed an RNN to compress 32×32 images in a progressive manner. They later extended the job to full-size images and introduced a BinaryRNN for context-based entropy modeling [22]. Johnston *et al*. [29] further introduced content-adaptive bit allocation, warm-start training tricks, and perceptual losses to boost the compression performance in terms of MS-SSIM.

Ballé *et al*. [19] proposed to learn shallow transforms with a GDN activation function in an end-to-end manner and model the entropy with a shared linear piece-wise PDF. In their following work [12], each code is supposed to follow a zero-mean Gaussian distribution with the deviation estimated from a side information network depending on the hierarchical hyperprior. Minnen *et al*. [13] combined the hierarchical hyperprior with context-based auto-regressive prior to boosting the compression performance.

Theis *et al*. [20] proposed a strait-forward relaxation of the quantization operations and exploited the Gaussian scale mixture for entropy modeling. Rippel *et al*. [11] proposed a pyramid-based network structure for analysis and synthesis transforms and an adaptive code length regularization for real-time image compression. And a generative adversarial loss [30] was introduced to generate visually better decoded images at low bit rates. Later, Agustsson *et al*. [21] proposed a soft-to-hard quantization scheme with a parametric softmax function.

Li *et al*. [23] learned an importance map as the side information for content variant rate controlling and exploited a simple CNN for context-based entropy modeling. Mentzer *et al*. [16] further proposed a mask convolutional network
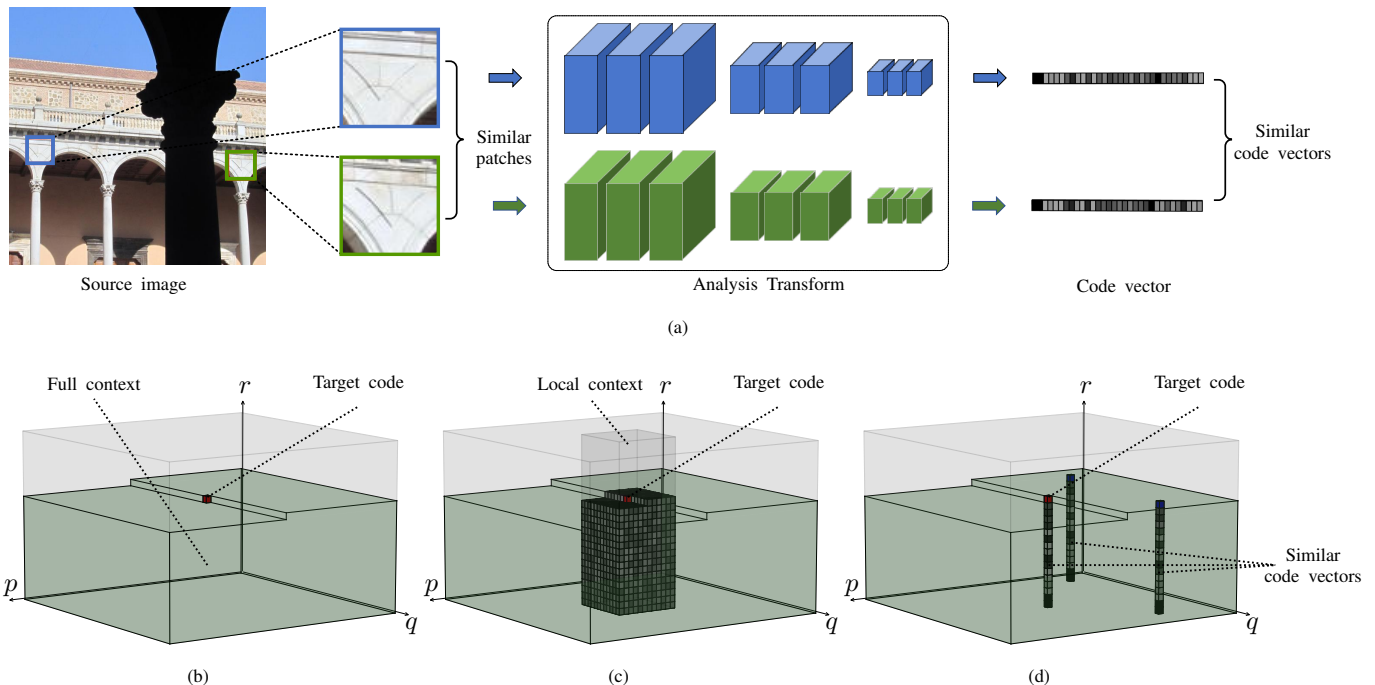
Fig. 1. Illustration of context-based non-local entropy modeling. (a) indicates the non-local similarity among the codes generated by the analysis transform. (b) shows the global context of a target code in 3D code block $\boldsymbol{y}$ with a raster scanning order. (c) gives the local context used in CNN-based entropy modeling. (d) illustrates the non-local similarity in the context.

for context-based entropy modeling and alternatively optimized the entropy model, transforms, and the importance map. Limited by the raster coding order, the context-based models should be decoded in serial coding order and thus are computational inefficiency. The heavy computational burden makes them not practical in many real applications. To speed up the entropy modeling and employ large context, context-based convolutional networks (CCNs) [17] introduce a specially defined coding order and group context for entropy modeling in both lossy and lossless image compression. With a given code dividing scheme, the CCNs can perform parallel computation within a group in decoding and largely speed up the decoding speed.

Previous deep context-based entropy models only focus on local context in the receptive field of CNNs and usually ignore the non-local similarity in the global context. To tackle this issue, we take the non-local similarity within the global context into account and combine it with the local representation for entropy modeling.

B. Non-local Methods for Image Processing

Non-local methods were first proposed for employing the global similarity among the pixels in image denoising. Later, DNN-based image processing methods embed it into DNNs to employ the global information and boost the performance in different tasks. Buades *et al.* [18] explored the self-similarity among pixels and proposed the non-local means based on a content weighted non-local averaging of all pixels in the image for image denoising. Wang *et al.* [31] formulated the non-local operation as a uniform block and adopt it in DNNs to combine local and non-local information for object detection.

Liu *et al.* [32] proposed a non-local recurrent network which incorporates non-local operations into a recurrent network for image restoration. Zhang *et al.* [33] exploited the non-local operation to building attention masks to capture long-range dependency between pixels and pay more attention to the challenging parts for high-quality image restoration. In this paper, we tackle the missing reference problem in context modeling with the proposed proxy similarity metric and introduce the non-local operation for the context-based entropy modeling. And a non-local attention block is introduced to combine local and global context for more accurate entropy estimation.

III. CONTEXT-BASED ENTROPY MODELING

Modeling the entropy of the code from its context is an auto-regression problem where the probability distribution functions (PDFs) of a code is regressed from the context of it. Let $\boldsymbol{y} \in \mathbb{R}^{M \times H \times W}$ denote a 3D code block generated by the analysis transform, where M, H, and W separately are the channel, height, and width of the code block. $y_r(p, q)$ is a code at the position (p, q) in r-th channel. With a given coding order, the context of a target $y_r(p, q)$ among a code block $\boldsymbol{y}$ is defined as all the codes scanned before it, *i.e.*, $\mathrm{CTX}(y_r(p, q), \boldsymbol{y})$. Fig. 1 (b) gives an example of the context of the target code in red with a raster scanning order.

Thanks to the CNN's efficiency in processing 2D and 3D media data, previous context-based entropy modeling usually adopts masked CNNs to predict the PDFs of the codes. As shown in Fig. 1 (c), limited by the structure, CNNs only focus on the local context surrounding the target code in the receptive field. The content similarity and global context are generally ignored. To exploit the global information and the content similarity between the target code and its context, we

introduce the non-local operation for context-based entropy modeling.

Non-local operation is usually defined as a weighted sum over the whole image, where the distance between a target pixel and other pixels is employed to produce the weights. Similar pixels play a more important role in estimating the target. As for the code block, we calculate the non-local representation of a target code $y_r(p, q)$ by summarizing all the codes in r-th channel of the context with content-related weights. However, it is hard to determine the weights when the target code itself is unavailable for computation in the context. Given a vector $\boldsymbol{\xi}_r(p, q) = \{y_0(p, q), \ldots, y_{r-1}(p, q)\}$, $\boldsymbol{\xi}_r(p, q)$ and $y_r(p, q)$ are produced by the same image patch. It is reasonable to evaluate the similarity between $y_r(p, q)$ and $y_r(u, v)$ with the distance between the two vectors, $\boldsymbol{\xi}_r(p, q)$ and $\boldsymbol{\xi}_r(u, v)$. As shown in Fig. 1(d), by finding similar code vectors, we predict the non-local representation of the target code in red with the similar codes in blue. Fig. 1 (a) shows that similar patches produce similar code vectors. As proved in many image restoration methods that employ similar data patches such as BM3D [34], there are many similar patches in a single image. Correspondingly, there should be many similar codes and code vectors in the code block produced by the image. Consequently, the introduced non-local operation should contribute to more accurate estimation of the entropy.

In this paper, we introduce a non-local attention block that combines the information from the local and global context for entropy modeling. Context-based convolutional networks (CCNs) [17] are adopted to modeling the local context for their efficiency and effectiveness. In the following, we will first give a short description of the CCNs and then introduce the context-based non-local operation.

A. CCNs in Modeling Local Context

Given the code block $\boldsymbol{y} \in \mathbb{R}^{M \times H \times W}$, the output of the t-th CCN layer is a 4D tensor $\boldsymbol{v}^{(t)} \in \mathbb{R}^{M \times H \times W \times N_t}$ with N_t feature blocks. Each feature block has the same size as the code block. The feature $v_{i,r}^{(t)}(p, q)$ in r-th channel and i-th feature block at spatial location (p, q) is a representation of $y_r(p, q)$ and only convey information from the $\text{CTX}(y_i(p, q), \boldsymbol{y})$.

To speed up the efficiency and break up the serial decoding process, the CCNs divide the codes into $K = M + H + W - 2$ non-overlap groups and parallel process each group, where $\text{GP}_k(\boldsymbol{y}) = \{y_r(p, q)|r + p + q = k\}$ denotes the k-th group. Then, the group context is defined as $\text{PTX}(y_r(p, q), \boldsymbol{y}) = \{y_{r'}(p', q')|r' + p' + q' < k\}$. With the group context, codes within each group share the same context and thus can be processed in parallel. Without a clear performance drop in the entropy modeling, the special group context could dramatically accelerate the decoding efficiency.

The CCNs are built on mask convolution layers which are defined as

$$v_{i,r}^{(t)}(p, q) = \sum_{j=1}^{N_t} \sum_{s=1}^{M} \left(u_{j,s}^{(t)} * \left(m_{r,s}^{(t)} \odot w_{i,j,r,s}^{(t)} \right) \right)(p, q) + b_i^{(t)},$$
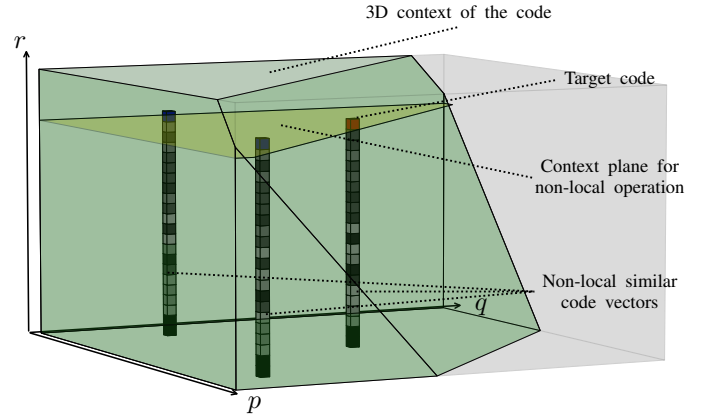
(2)

Fig. 2. Illustration of the non-local similarity of the codes within the group context of CCNs for entropy modeling. The red block is the target code. And the green region is the group context for the CCNs. Content similar codes in blue are located by the proxy similarity metric g_d and then employed to predict the target code. The yellow plane indicates the available codes in the non-local operation. Zoom in for better visual quality and details.

where $\{i, j\}$ and $\{r, s\}$ are indexes for the feature block and channel dimensions, respectively. $\boldsymbol{u}^{(t)}$ and $\boldsymbol{v}^{(t)}$ are the input and output of the t-th convolution layer. $\boldsymbol{v}^{(t-1)}$ is activated by a nonlinear element-wise function to produce $\boldsymbol{u}^{(t)}$. $w_{i,j,r,s}^{(t)}$ is the weight to connect feature maps $u_{j,s}^{(t)}$ and $v_{i,r}^{(t)}$, and $b_i^{(t)}$ is the bias term. Given $w_{i,j,r,s}^{(t)} = \{w_{i,j,r,s}^{(t)}(u, v)| - k_s \leq u, v \leq k_s\}$ where k_s is kernel size, the corresponding mask for the input layer is defined as

$$m_{r,s}^{(0)}(u, v) = \begin{cases} 1, & \text{if } s + u + v < r \\ 0, & \text{otherwise.} \end{cases}$$

(3)

For the t-th hidden layer, the mask is modified to include the codes in the same group

$$m_{r,s}^{(t)}(u, v) = \begin{cases} 1, & \text{if } s + u + v \leq r \\ 0, & \text{otherwise.} \end{cases}$$

(4)

B. Context-based Non-local Operation

Different channels in the code block are generated by different convolutional filters. Only the codes in the same channel are produced by the same transform and thus should follow the same distribution. We only consider the codes in the same channel in the non-local operation. Besides, the context should be considered. For the target code $y_r(p, q)$, masks $\boldsymbol{m}^l$ are introduced to exclude the codes outside of the context. Following the group context $\text{PTX}(y_r(p, q), \boldsymbol{y})$ in CCNs, we set the mask $m^l(p, q, u, v)$ to be 1 if $y_r(u, v) \in \text{PTX}(y_r(p, q), \boldsymbol{y})$, i.e., $v + v < p + q$. Otherwise, $m^l(p, q, u, v) = 0$. Fig. 2 shows the codes used in the non-local operation with a yellow plane where the masks are 1s.

Evaluating the similarity between the target and the others is necessary for non-local operation. Due to the target code is not in the context, it is unable to calculate the distance between the target and others directly. Instead, we propose a proxy similarity function by considering the available codes in the same position but the other channels. The similarity metric

$g_d(y_r(p,q), y_r(u,v))$ between the code $y_r(p,q)$ and $y_r(u,v)$ is defined as,

$$g_d(y_r(p,q), y_r(u,v)) = \sum_{j=0}^{r-1} w_{r,j}^d \| y_j(p,q) - y_j(u,v) \|^2. \quad (5)$$

Here, $w_{r,j}^d$ is the weight to balance the contribution of the proxy codes in different channels. $w_{r,j}^d$ is set to be $1/(r+1)$ in initialization and dynamically optimized with the training of the whole framework. g_d is a proxy based on the available codes. When considering more codes in different channels, the result will be more accurate and close to the real similarity.

The context-based non-local operation is defined as,

$$g_{nlc}(y_r(p,q)) = \sum_{u,v} w_r^s(p,q,u,v) y_r(u,v), \quad (6)$$

where the weight $w_r^s(p,q,u,v)$ is defined as,

$$w_r^s(p,q,u,v) = \frac{m^l(p,q,u,v) e^{-g_d(y_r(p,q), y_r(u,v))}}{\sum_{u',v'} m^l(p,q,u',v') e^{-g_d(y_r(p,q), y_r(u',v'))}}. \quad (7)$$

An attention block is introduced to combine the non-local representation of the target code with the local representation produced by CCNs. Given the constraint of context and the proxy similarity function, the non-local representation of a target code may be predicted by the dissimilar codes and thus hinders the performance. To indicate whether the target code is estimated by content similar codes and reduce the contribution of non-local representation in those bad estimations, a confidence indicator is introduced as the weighed similarity among the proxy codes.

$$g_c(y_r(p,q)) = \sum_{u,v} w_r^s(p,q,u,v) g_d(y_r(p,q), y_r(u,v)). \quad (8)$$

We combine the confidence indicators and local representation to generate the attention weight. And the Hadamard product of the attention weights and non-local representation is then concatenated with the local representation to produce output of the attention block. Fig. 3(b) gives the structure of the block.

IV. CONTEXT-BASED NON-LOCAL ENTROPY MODELING FOR LOSSY IMAGE COMPRESSION

A lossy image compression method usually involves transforms based on DNNs and is optimized for a joint rate-distortion objective function. Fig. 3 gives the framework of our lossy image compression method which consists of the analysis transform g_a, quantizer g_q, synthesis transform g_s and the context-based non-local entropy model g_e. The analysis transform maps a color image x to the code representation z, which is further discretized by g_q to produce the code block y. The synthesis transform g_s takes y as input to produce a color image $\hat{x}$ as the reconstruction of x. In this section, we will first describe the transforms, quantization function and the objective function and then introduce a post entropy model to simplify the entropy coding.

A. Network Structure for Transforms

Deep neural networks usually are non-invertible. As a result, transformations based on DNNs will lose some information in mapping the input image to the code representation. Such information loss could be ignored for models at low bit rates when the distortion between the input and the decoded image is large enough. However, with the decrease of the distortion, the information loss brought by DNNs will start to hinder the performance. To reduce the information loss, some learned image compression methods [11–13] suggest adopting a wider network at high bit rates. Consequently, the time complexity and memory consumption also increase rapidly. A network structure taking both the computational efficiency and the width into account is needed.

U-Net [35, 36] is a light structure originally proposed for medical image segmentation. With paired down-sampling, up-sampling operations and skip connections, U-Net is fast and can utilize representations in different scales. Besides, the skip connection could facilitate information propagation and ease the training of the network [37]. As shown in Fig. 3(a), we adopt the U-Net structure in building the basic block of our transforms, *i.e.*, UnetBlock, which can keep the width and reduce the time complexity and memory consumption of the transforms at the same time. Each UnetBlock consists of 3 down-sampling and 3 up-sampling convolution layers, where the down-sampling/up-sampling multipliers are determined by a_0, a_1 and a_2 separately. And skip connections are introduced to combine the features at each scale. The down-sampling convolution layer is conducted by the stride convolution, while the up-sampling convolution layer is modeled as a normal convolution to increase the channels followed by the depth-to-space reshaping [22, 38].

The analysis transformation g_a is composed of 3 down-sampling layers with each followed by a UnetBlock. Each convolution layer is followed by a PReLU activation function. And an additional convolutional layer with sigmoid nonlinearity is adopted after the last UnetBlock to produce $z \in (0,1)$ with M channels. The synthesis transformation g_s is a mirror of the analysis transform. And the upsampling operation is the same as used in UnetBlock. The last convolution layer makes use of three filters to produce the RGB decoded image. ϕ and ψ represent the parameters to be optimized in g_a and g_s separately. Please refer to Fig. 3 for more details about the structure of the transforms.

B. Quantization Function

The quantization function g_q maps the output of analysis transform z to L discrete quantization centres in each channel and generates the quantized code block y. We parametrize the quantization interval for the r-th channel by $\{\sigma_{r,0}, \ldots, \sigma_{r,L-1}\}$. Then, the quantization centres for r-th channel is represented as,

$$\omega_{r,i} = \sum_{j=0}^{i} e^{\sigma_{r,j}}, \text{for } i = 0, \ldots, L-1. \quad (9)$$
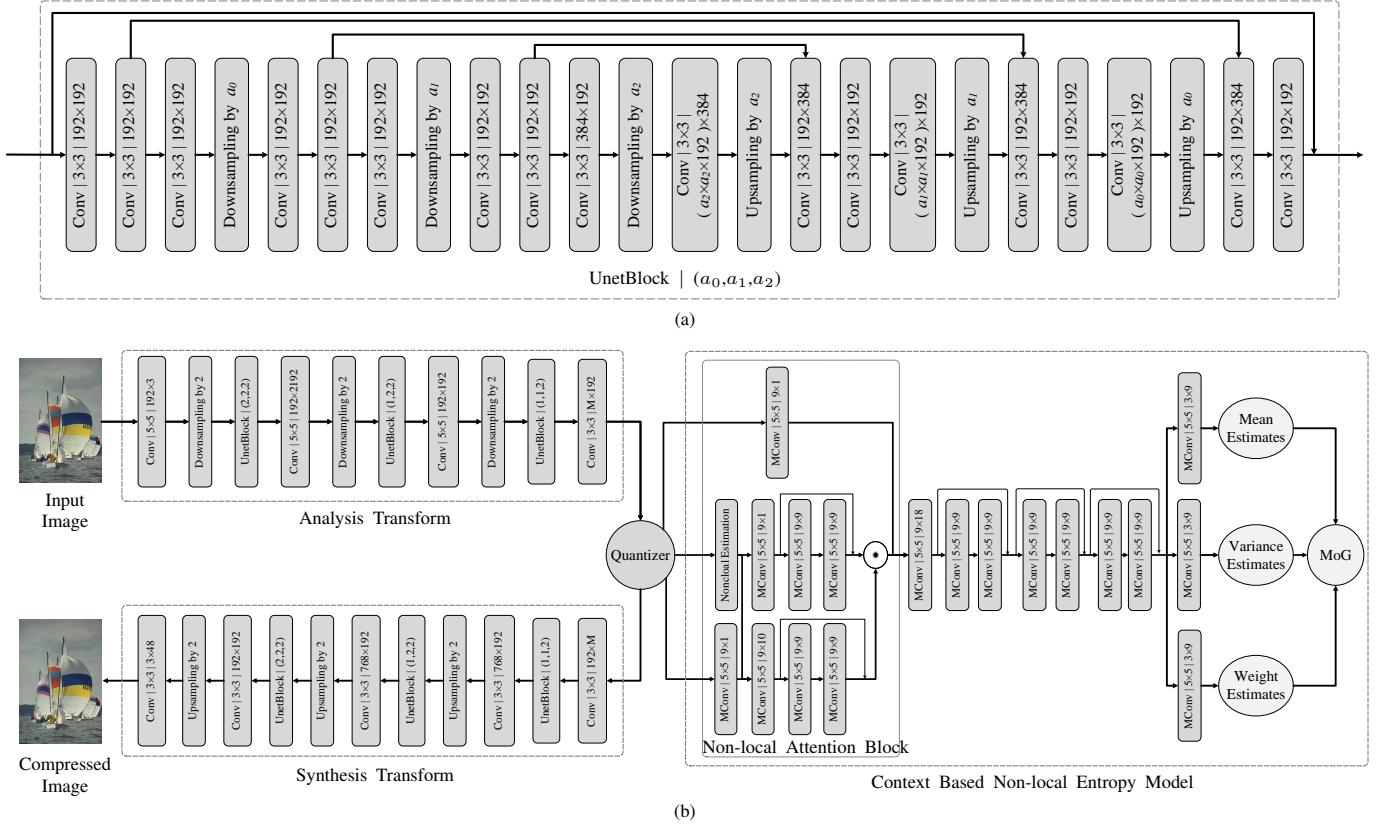
Fig. 3. The architecture of the proposed lossy image compression method, including an analysis transform g_a, an adaptive trainable quantizer g_q, a context-based non-local entropy model g_e, and a synthesis transform g_s. Conv: regular convolution with filter support (kernel_height × kernel_width) and number of channels (output×input). UnetBlock: a_0, a_1 and a_2 are the down-sampling multipliers in the U-Net structure. MConv: mask convolution used in CCNs with filter size (kernel_height× kernel_width) and number of feature blocks (output×input). Note that the number of channels is fixed in MConv, and is the same as the input of the entropy model, *i.e.*, $\boldsymbol{y}$.

In quantization, the code representation $z_r(p,q)$ is assigned to the nearest quantization center in r-th channel.

$$g_q(z_r(p,q)) = \arg\min_{\omega_{r,l}} \|z_r(p,q) - \omega_{r,l}\|_2^2. \quad (10)$$

Here, g_q has zero gradients almost everywhere, which blocks the training via back-propagation. Inspired by the binarized neural networks [39–41] and the DNN-based compression framework [20], an identify mapping $\hat{g}_q(z_r(p,q)) = z_r(p,q)$ is adopted as a continuous proxy for the quantization function in back-propagation.

We initialize $\boldsymbol{\sigma}$ as a uniform quantization in $(0,1)$, which are then optimized and updated according to the distribution of $\boldsymbol{z}$ by minimizing the mean squared quantization error,

$$\mathcal{L}_q(\boldsymbol{\sigma}) = \frac{1}{MHW} \sum_{r,p,q} \|y_r(p,q) - z_r(p,q))\|_2^2, \quad (11)$$

C. Modeling the Objective Function

The whole framework is optimized for a joint rate-distortion objective function, where the distortion loss is directly modeled as the difference between the decoded image $\hat{\boldsymbol{x}}$ and the input image $\boldsymbol{x}$. Two separate metrics, *i.e.*, standard mean square error (MSE) and the perceptual metric MS-SSIM [42], are adopted as the distortion loss. The MSE distortion loss

$\mathcal{L}_D^{\text{MSE}}$ and MS-SSIM distortion loss $\mathcal{L}_D^{\text{MS-SSIM}}$ are defined as follows

$$\mathcal{L}_D^{\text{MSE}}(\boldsymbol{x};\boldsymbol{\phi},\boldsymbol{\psi}) = \frac{1}{3H_IW_I}\|\hat{\boldsymbol{x}} - \boldsymbol{x}\|_2^2, \quad (12)$$

and

$$\mathcal{L}_D^{\text{MS-SSIM}}(\boldsymbol{x};\boldsymbol{\phi},\boldsymbol{\psi}) = 100 - 100\,\text{MS-SSIM}(\hat{\boldsymbol{x}},\boldsymbol{x}), \quad (13)$$

where $\hat{\boldsymbol{x}} = g_s\Big(g_q(g_a(\boldsymbol{x};\boldsymbol{\phi});\boldsymbol{\sigma});\boldsymbol{\psi}\Big)$. H_I and W_I are separately the height and width of the image $\boldsymbol{x}$. We denote our method optimized for $\mathcal{L}_D^{\text{MSE}}$ as Ours(MSE) and $\mathcal{L}_D^{\text{MS-SSIM}}$ as Ours(MS-SSIM).

Each code in $\boldsymbol{y}$ is assumed to follow an MoG distribution depending on its context. We introduce a context-based non-local entropy model g_e to produce the mixture weight, mean, and variance of these MoG distributions. As shown in Fig. 3, the entropy model consists of a non-local attention block, several CCN residual blocks, and three final CCN layers to produce the mean, variance, and weight estimates separately. $\boldsymbol{\theta}$ is the parameters for the proposed entropy model.

Given $y_r(p,q) = \omega_{r,l}$, the discretized probability of a code $y_r(p,q)$ is defined as,

$$P(y_r(p,q);\theta) = \int_{\frac{\omega_{l-1}+\omega_l}{2}}^{\frac{\omega_l+\omega_{l+1}}{2}} \sum_{i=0}^{C-1} \pi_i \mathcal{N}(\alpha;\mu_i,\sigma_i^2)d\alpha, \quad (14)$$

where C is the number of Gaussian distribution in the mixture; π_i, μ_i and δ_i^2 are the mixture weight, mean, and variance of the i-th component, respectively. Specially, $\omega_{r,-1} = -\infty$ and $\omega_{r,L} = \infty$ for the boundary cases. Then, the entropy of the codes are adopted as the rate loss.

$$\mathcal{L}_R(\boldsymbol{x}; \boldsymbol{\phi}, \boldsymbol{\theta}) = -\sum \log P\big(g_q\big(g_a(x_r(p,q); \boldsymbol{\phi}); \boldsymbol{\sigma}\big), \boldsymbol{\theta}\big). \quad (15)$$

Finally, we propose a rate-distortion objective function for the parameters $\{\boldsymbol{\phi}, \boldsymbol{\psi}, \boldsymbol{\theta}\}$ over the training set $\mathcal{X}$ as

$$\mathcal{L}(\boldsymbol{\phi}, \boldsymbol{\psi}, \boldsymbol{\theta}) = \sum_{\boldsymbol{x} \in \mathcal{X}} \mathcal{L}_D(\boldsymbol{x}; \boldsymbol{\phi}, \boldsymbol{\psi}) + \lambda \mathcal{L}_R(\boldsymbol{x}; \boldsymbol{\phi}, \boldsymbol{\theta}), \quad (16)$$

where λ is the trade-off parameter to balance the rate and distortion in the objective function.

D. Post Processing for Entropy Coding

In entropy coding, a discrete probability table that contains the probability $P(y_r(p,q) = \omega_{r,l})$ for $l = 0, \ldots, L-1$ is needed. With the current entropy model built on MoG distributions, we should first produce the parameters for MoG distributions and then cut the PDF into L intervals and get the discrete possibility table with an integral on each interval. Considering the amount of code to be processed is very large, it will bring further computational burden and slow down the entropy coding. We simplify this process by training a post entropy model which takes the codes $\boldsymbol{y}$ as input and directly output the discrete probability table that the code belongs to each quantization center. For the network structure, we only make a small modification on the last layer. Instead of taking three separate CCN layers to produce the mean, variance, and weight estimations used in MoGs, we make use of a single CCN layer followed by a softmax nonlinearity to produce the possibility table $\boldsymbol{u}$ where $u_{i,r}(p,q)$ represents the possibility that $y_r(p,q) = \omega_{r,i}$. The post entropy network is optimized by minimizing the expected code length

$$\mathcal{L}_{post}(\boldsymbol{\theta}) = -\mathbb{E}_{\boldsymbol{y}}\Big[\sum_{r,p,q} \sum_i \mathbb{1}(y_r(p,q) = \omega_{r,i}) \log_2(u_{i,r}(p,q)) \Big], \quad (17)$$

where $\mathbb{1}(\cdot)$ is an indicator function. We first extract $\boldsymbol{y}$ for $\boldsymbol{x} \in \mathcal{X}$ with the trained analysis transform g_a and then optimize the post entropy model on the extracted code blocks. Finally, we implement our own arithmetic coding with the context-based non-local entropy model to compress $\boldsymbol{y}$ to bitstreams, and report performance using actual bit rates.

V. EXPERIMENTS

In this section, we test the proposed context-based non-local entropy model and Unet-block in the lossy image compression. $10,000$ high-quality images are collected from the photo-sharing website Flickr and down-sampled to further reduce possibly compression artifacts. We crop $640,000$ color patches of size $3 \times 256 \times 256$ as the training sets. For the post entropy model, we first extract the code blocks with the analysis transform from the full-size images and then crop code blocks with the size of $M \times 60 \times 60$ for training. We test our models on two benchmark datasets - Kodak and Tecnick [43], and
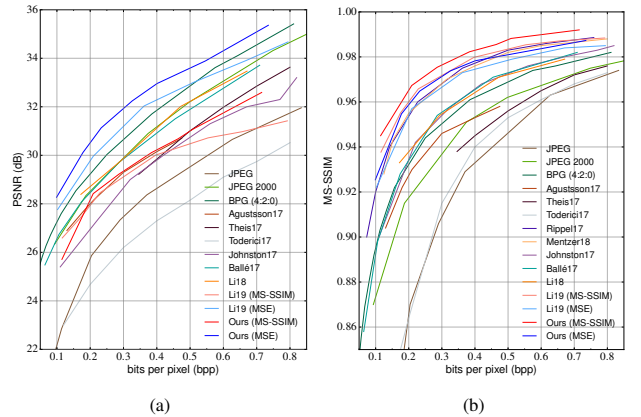

Fig. 4. Rate-distortion curves of different compression methods on the Kodak dataset. (a) PSNR. (b) MS-SSIM.

then compare them to the recent deep image compression algorithms and state-of-the-art image compression standards. The pre-trained models for testing are now available at https://github.com/limuhit/Nonlocal-CCN.

A. Experimental Setup

The quantization centers L is set to be 8 and the number of Gaussian components in MoG C is set to be 3. We follow the warmup strategy in [17] and jointly optimize the transforms and entropy model by minimizing the rate-distortion objective function in Eq. 16 with an Adam solver [44]. Starting from a learning rate of 10^{-5}, smaller learning rates, *i.e.*, 10^{-6} and 10^{-7}, are adopted until the loss does not decrease for 5 successive epochs. The post entropy model is also optimized with Adam solver in the same way. We train 14 models for seven different bit rates and two distinct distortion metrics, i.e., MSE and MS-SSIM. For testing, the compression rate is evaluated by bits per pixel (bpp), which is the total amount of bits used to compress the image divided by the whole number of pixels in the image. Two quantitative metrics, *i.e.*, Multi-Scale Structural Similarity (MS-SSIM) and the Peak Signal-to-Noise Ratio (PSNR), are considered in evaluating the image distortion.

B. Quantitative Evaluation

Using MS-SSIM and PSNR as distortion metrics, we compare our methods with existing image compression standards such as JPEG [1], JPEG2000 [45] and BPG [46] and recent DNN-based compression models in terms of rate-distortion curves. DNN-based compression models include Agustsson17 [21], Theis17 [20], Toderici17 [22], Rippel17 [11], Mentzer18 [16], Johnston17 [29], Ballé17 [19], Li18 [23] and Li19 [17]. Both JPEG (with 4:2:0 chroma subsampling) and JPEG2000 are based on the optimized implementations in MATLAB2017. For BPG, we adopt the latest version from its official website with the default setting. When it comes to DNN-based compression models, the implementations are generally not available. For those models, we carefully digitalize the rate-distortion curves and report the results from their respective papers.
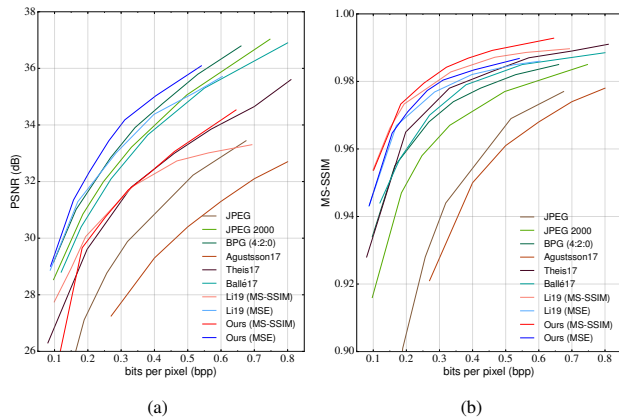
Fig. 5. Rate-distortion curves of different compression methods on the Tecnick dataset. (a) PSNR. (b) MS-SSIM.

TABLE I
RUNNING TIME IN SECONDS, GPU MEMORY IN GBS AND DISTORTION EVALUATED WITH PSNR OF THREE NETWORK STRUCTURES.

Network Structure	Running Time (seconds)	GPU Memory (GB)	Distortion (dB)
DenseBlock [17]	0.025	0.86	35.54
ResidualBlock	0.136	2.65	36.34
UnetBlock	0.048	1.52	36.29

TABLE II
ENTROPY CODING FOR CODES OPTIMIZED BY MSE AND MS-SSIM. THE RESULTS ARE EVALUATED BY BITS PER CODE. $y^{a,i}$ ($y^{b,i}$) REPRESENTS THE CODE OF THE i-TH MODEL OPTIMIZED FOR THE MSE(MS-SSIM).

Entropy coding for codes optimized by MSE							
Code set	$y^{a,0}$	$y^{a,1}$	$y^{a,2}$	$y^{a,3}$	$y^{a,4}$	$y^{a,5}$	$y^{a,6}$
CCN [17]	1.65	1.56	1.45	1.45	1.43	1.46	1.58
Non-local	1.58	1.42	1.24	1.30	1.28	1.34	1.47

Entropy coding for codes optimized by MS-SSIM							
Code set	$y^{b,0}$	$y^{b,1}$	$y^{b,2}$	$y^{b,3}$	$y^{b,4}$	$y^{b,5}$	$y^{b,6}$
CCN [17]	1.86	1.78	1.66	1.60	1.64	1.39	1.55
Non-local	1.84	1.67	1.53	1.54	1.48	1.26	1.43

Fig. 4 shows the rate-distortion curves on the Kodak dataset. The results for PSNR of Rippel17 [11] and Mentzer18 [16] are missing as they do not report the PSNR curve in the paper. We find that both Ours(MSE) and Ours(MS-SSIM) outperform all competing methods in terms of PSNR and MS-SSIM separately. Especially, compared to Li19 [17] with local context-based entropy modeling, our models with non-local attention block are much better, which support the effectiveness of the introduced non-local operation in context modeling. And at the high bit rate region, our models outperform Li19 by a lot, which also indicates the effectiveness of the proposed UnetBlock in reducing the information loss in transforms. Fig. 5 shows the rate-distortion curves on the Tecnick dataset, where similar trends as Koadak dataset for both PSNR and MS-SSIM can be observed. We fail to compare with Rippel17 [11], Theis17 [20], Agustsson17 [21] and Mentzer18 [16] on Tecnick dataset due to the unavailable results in the papers.

C. Visual Quality Evaluation

We further compare the decoded images by our method against Li19, JPEG2K, and BPG in visual quality and show the sample decoded images and the uncompressed images on the Kodak dataset in Fig. 6 and Tecnick dataset in Fig. 7, respectively. For images at low bit rate, the methods optimized with MS-SSIM are visually much better due to that MS-SSIM takes structural similarity in different scale into account and are more consistent with the human visual system. When it comes to the images at a high bit rate, the methods optimized with MSE are better at keeping small scale details. Thus, we compare Ours(MS-SSIM) with Li19(MS-SSIM) at low bit rates as shown in Fig. 7 and compare Ours(MSE) with Li19(MSE) at higher bit rates in Fig 6. In Fig. 7, JPEG2K and BPG exhibit artifacts (such as blocking, ringing, blurring, and aliasing) that are common to all handcrafted transform coding methods. Li19(MS-SSIM) is effective at suppressing most of the artifacts bus still surfers from blurring in some parts of the image. In contrast, our method optimized for MS-SSIM is more able to keep the details and has less visible distortions. In Fig. 4, the whole visual quality is nearly the same. But when zooming into details, similar to BPG, our

methods optimized for MSE shows to have better small scale edges and textures. Li19(MSE) blurring the small edges and losing important information, such as the text.

D. Ablation Experiments

We conduct thorough ablation experiments to analyze the impact of individual components, *i.e.*, the UnetBlock, and context-based non-local entropy modeling, to final compression performance. For fair comparisons, we use the same set of parameters and training set for all the competing models.

1) UnetBlock for analysis and synthesis transforms: The width of the transforms, *i.e.*, the number of feature maps in the output of each layer, is supposed to have a significant influence on the performance of low distortion image compression models. A narrow transform will inevitably lose some information and thus introduce extra distortion in reconstructing the input image. To support the effectiveness of increasing the width in low distortion cases, we introduce a baseline model by combining the entropy model from Li19 [17] and transforms based on UnetBlocks. The width for the baseline model and Li19 are 196 and 64.

We compare the baseline models separately optimized for two distortion metrics, *i.e.*MS-SSIM and MSE, to Li19 [17] in Fig. 8. At a low bit rate region (< 0.4bpp), the baseline model optimized for MSE, *i.e.*, Baseline(MSE), and Li19(MSE) have nearly the same performance in both of the MS-SSIM and MSE. But when the bit rate grows, the Baseline(MSE) has a significant improvement, which supports the width of the transforms is important at high bit rate regions where the distortion is quite small. Similar trends can be observed for our model optimized with MS-SSIM.

With the same width and depth, Unet is usually fast in speed and needs less computational resources. To evaluate the performance of the UnetBlock in transforms, we further compare the UnetBlock to residual block [47] on a computer
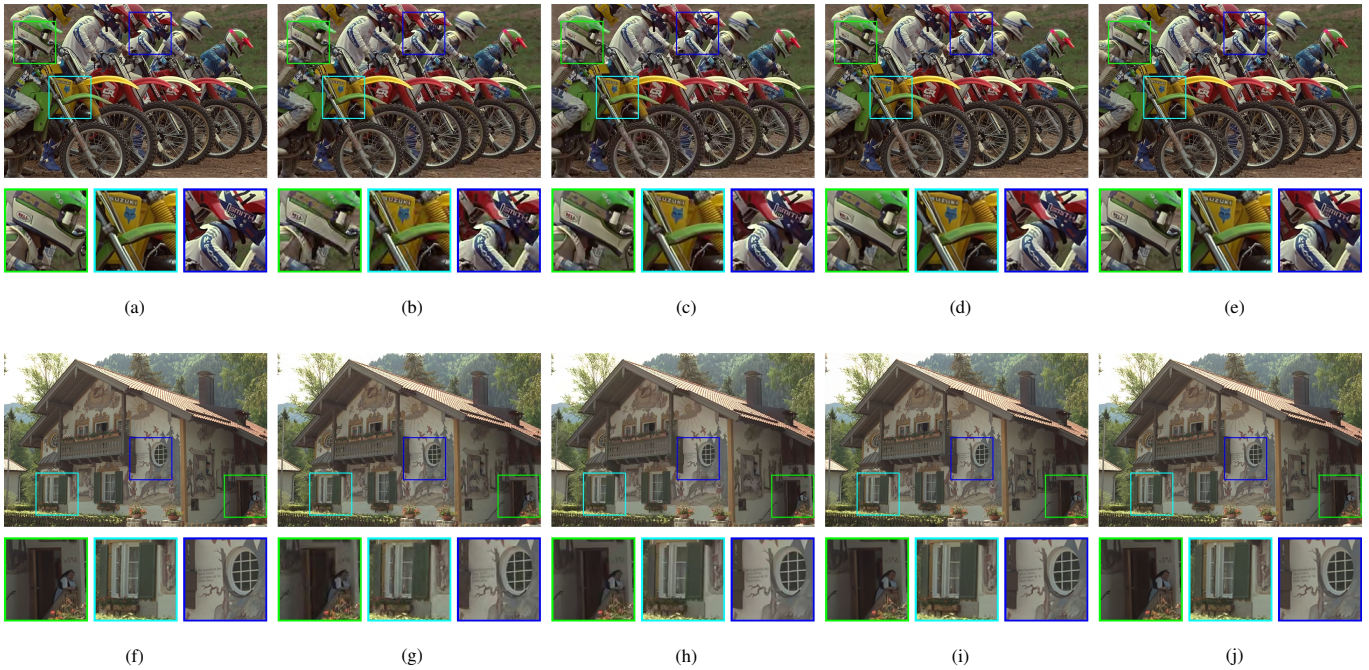
Fig. 6. Compressed images by different compression methods on the Kodak dataset. The quantitative measures are in the format of "bpp / PSNR / MS-SSIM". (a) Uncompressed "Motorcycle" image. (b) JPEG2K. 0.673 / 27.75 / 0.962. (c) BPG. 0.712 / 29.65 / 0.973. (d) Li19 [17] optimized for MSE. 0.694 / 28.81 / 0.986. (e) Ours(MSE). 0.670 / 30.53 / 0.984. (f) Uncompressed "House" image. (g) JPEG2K. 0.879 / 31.52 / 0.972. (h) BPG. 0.877 / 32.64 / 0.979. (i) Li19 optimized for MSE. 0.871 / 29.99 / 0.988. (j) Ours(MSE). 0.865 / 33.04 / 0.989.
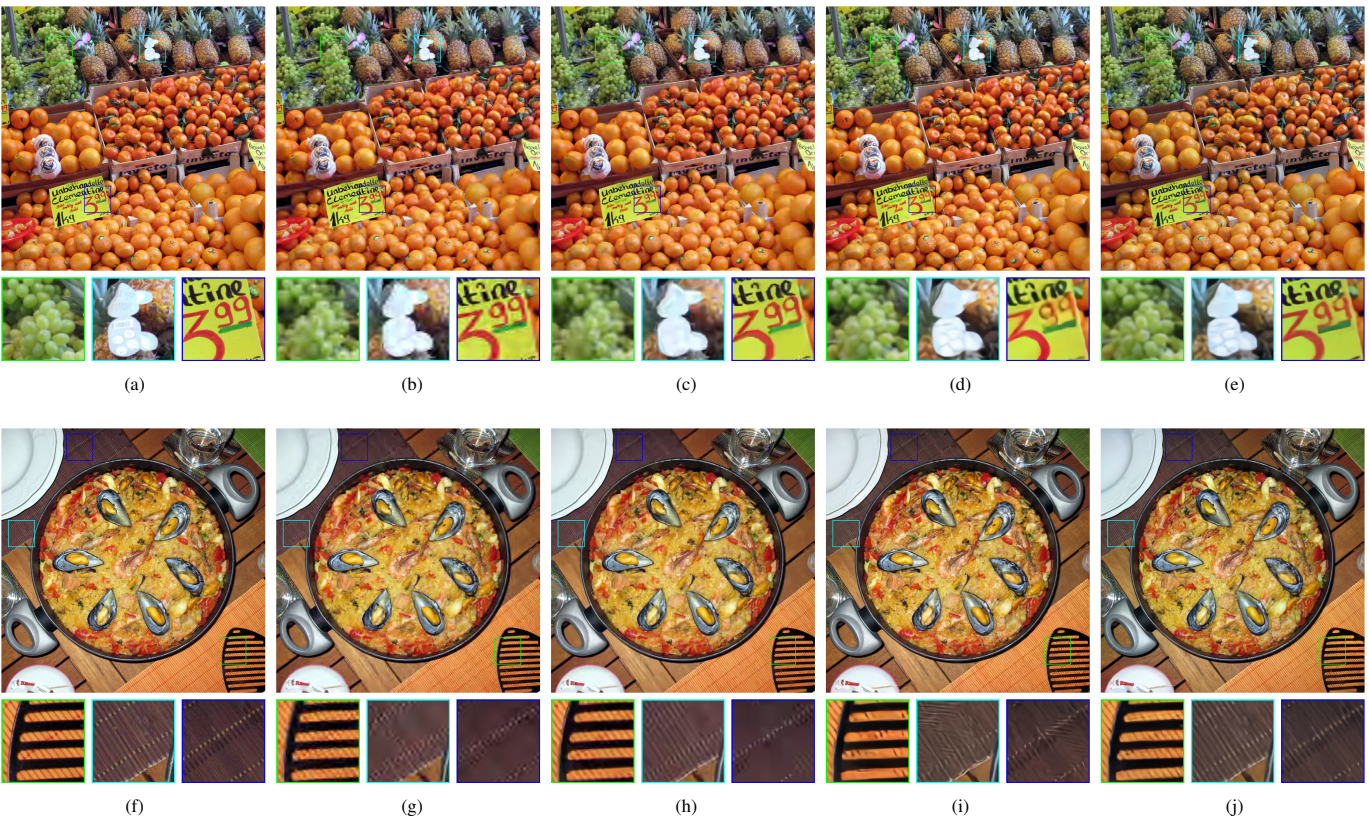

Fig. 7. Compressed images by different compression methods on the Tecnick dataset. The quantitative measures are in the format of "bpp / PSNR / MS-SSIM". (a) Uncompressed "Fruit" image. (b) JPEG2K. 0.137 / 24.95 / 0.890. (c) BPG. 0.136 / 25.58 / 0.901. (d) Li19 [17] optimized for MS-SSIM. 0.139 / 25.63 / 0.938. (e) Ours(MS-SSIM). 0.134 / 24.46 / 0.939. (f) Uncompressed "Seafood" image. (g) JPEG2K. 0.137 / 25.52 / 0.882. (h) BPG. 0.135 / 26.42 / 0.905. (i) Li19 optimized for MS-SSIM. 0.138 / 25.62 / 0.925. (j) Ours(MS-SSIM). 0.133 / 25.03 / 0.932.
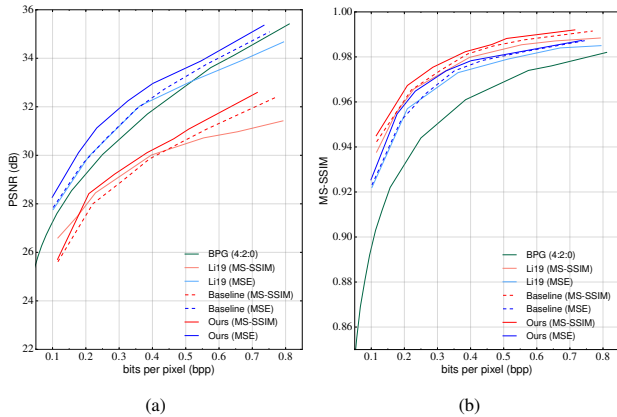
Fig. 8. Rate-distortion curves of different variants of the proposed methods on the Kodak dataset. (a) PSNR. (b) MS-SSIM.

with an Intel(R) Xeon(R) Processor E5-2620 v4, 64 GB of RAM and a NVIDIA TITAN Xp GPU. We adopt several residual blocks to replace a single UnetBlock to keep the depth of the competing transforms. The DenseBlocks in Li19 [17] concatenate all the output of previous sub-blocks as the input for the next sub-block. The computational consumption increases dramatically with the growth of the width of the input to the block. Limited by the GPU memory, we fail to compare to transforms based on DenseBlock with a width of 192. Thus, we only adopt the transforms with a width of 64 for DenseBlock in comparison. Without the limitation of the entropy model ($\lambda = \infty$), the three competing models are optimized for the MSE distortion loss. Table I gives the results of the three structures on running time in seconds, GPU memory in GB and distortion performance evaluated by PSNR. The transforms based on the ResidualBlock can not process the whole image from the Kodak dataset due to the GPU memory. We only report the performance on 256×256 patches sampled from the Kodak dataset. The UnetBlock performs on par with ResidualBlock on distortion performance and overwhelm the narrow network, *i.e.*, DenseBlock. When it comes to efficiency, the transforms based on UnetBlock is much faster and needs less GPU memory than the transforms based on ResidualBlock. Thus, the UnetBlock is a good trade-off between the speed and performance.

2) Context-based non-local entropy modeling: The context-base non-local entropy modeling exploits both of the global similarity and the local representations of the context for entropy modeling. We compare it to the entropy model only makes use of local representations [17]. The same baseline models are used for comparison. Compared to Ours(MSE) and Ours(MS-SSIM), the baselines adopt the CCN-based entropy models which focus on local representations. As shown in Fig. 8, both Ours(MSE) and Ours(MS-SSIM) outperform the counterpart baseline models by a large margin, which strongly supports the effectiveness of the proposed non-local attention block in context modeling.

To remove the effectiveness of different initialization and training process, we test the performance of the context-based non-local entropy models and the CCN-based entropy models on the extracted codes. $\boldsymbol{y}^{a,0}, \dots, \boldsymbol{y}^{a,6}$ denote seven

code sets produced by Ours(MSE) at seven different bit rates. $\boldsymbol{y}^{b,0}, \dots, \boldsymbol{y}^{b,6}$ are produced by Ours(MS-SSIM). The performance of the two entropy models is evaluated by bits per code, which is the rate between the total bits used to coding the code block and the number of codes in the code block. Table II gives the performance of the two entropy models, the context-based non-local entropy model is significantly better than the CCN-based entropy model on the 14 code sets, which supports the contribution of the introduced non-local attention block in entropy modeling.

We further visualize the estimated probability for each code plane by mapping $P(y_r(p, q))$ to an integer in the range of $[0, 255]$ and shows the code plane as a gray image in Fig. 9. Fig. 9 (a) gives the uncompressed image; (b)-(i) are the visualized probability of the CCN-based entropy model on 8 code planes; (j)-(q) are the corresponding results of the context-based non-local entropy model. In bottom code planes, as shown in Fig. 9 (b), (c), (j) and (k), CCN-based entropy model has similar performance as the proposed non-local entropy model. This can be illustrated by the proxy similarity function. The similarity of codes inner one plane is estimated by the proxy codes in the bottom code planes. With a few numbers of bottom code planes, the similarity is not accurate and thus lead to poor non-local estimation. The attention block tends to focus on the local representations instead of the non-local estimation. Thus, the competing entropy models share a similar performance. With the increase of planes in the bottom, the proxy similarity begins more accurate and the non-local estimations star to contribute to the performance. The context-based non-local entropy modeling shows to have overwhelming performance in the top code planes as described in Fig. 9 (f), (g), (h), (i), (n), (o), (p) and (q).

VI. Conclusion

In this paper, we solved the missing reference problem in the context modeling with a proxy similarity metric and introduced the non-local operation for contex-based entropy modeling. With a non-local attention block to combine the local and global representations, our non-local entropy models can provide more precise estimation for the entropy of the codes. To reduce the information loss and boost the performance for low distortion compression, we further designed an effective and efficient network structure, *i.e.*, UnetBlock, and adopted it in building the transforms of the proposed lossy image compression framework. We tested the context-based non-local entropy model and the UnetBlock in lossy image compression. Both of our models optimized for MSE and MS-SSIM got significant improvements over state-of-art image compression standards and recent DNN-based models.

References

[1] G. K. Wallace, "The JPEG still picture compression standard," *IEEE Trans. Consumer Electron.*, vol. 38, no. 1, pp. 18–34, 1992.
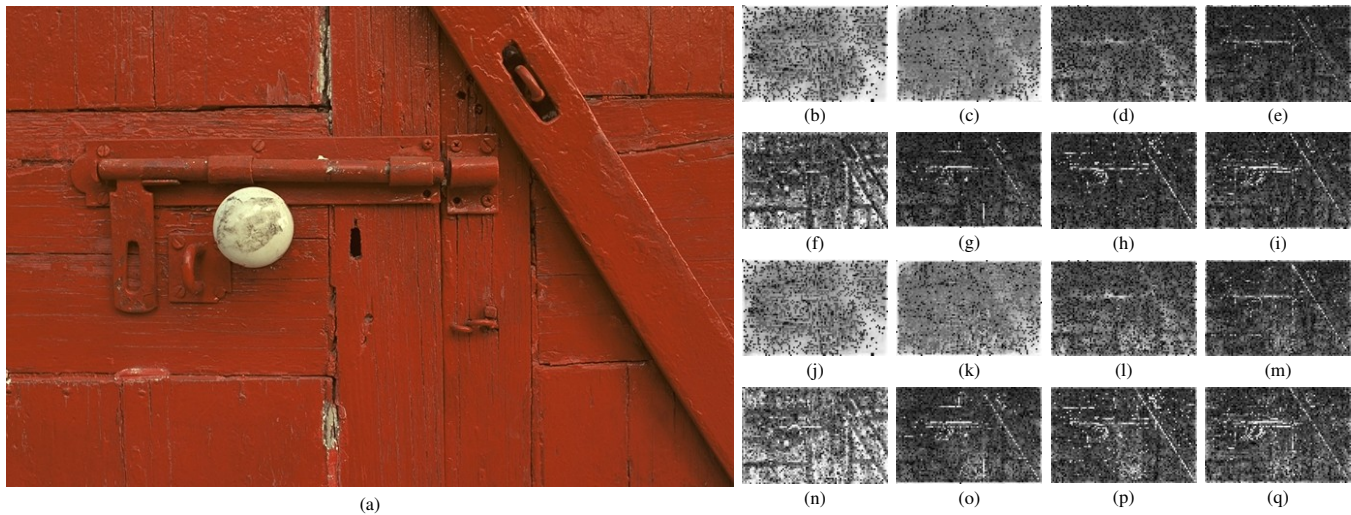
Fig. 9. Visualization of the estimated probability for the context-based non-local entropy model and the CCN-based entropy model. Each gray figure represents one code plane. The bright and dark colors indicate large and small probability of the corresponding codes. The quantitative measures are in the format of channel number of the code plane / bits per code. For simplification, "NLC" represents the context-based non-local entropy model and "LC" represents the CCN-based entropy model. (a) Uncompressed "Door" image. (b) LC. 1 / 1.05. (c) LC. 7 / 1.16. (d) LC. 16 / 1.85. (e) LC. 27 / 2.40. (f) LC. 28 / 1.67. (g) LC. 29 / 2.37. (h) LC. 30 / 2.37. (i) LC. 31 / 2.30. (j) NLC. 1 / 1.05. (k) NLC. 7 / 1.14. (l) NLC. 16 / 1.75. (m) NLC. 27 / 2.11. (n) NLC. 28 / 1.21. (o) NLC. 29 / 2.03. (p) NLC. 30 / 1.97. (q) NLC. 31 / 1.94.

[2] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising," *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3142–3155, 2017.

[3] C. Dong, C. C. Loy, K. He, and X. Tang, "Learning a deep convolutional network for image super-resolution," in *Eur. Conf. Comput. Vis.*, 2014, pp. 184–199.

[4] ——, "Image super-resolution using deep convolutional networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 2, pp. 295–307, 2016.

[5] A. Majumdar, "Blind denoising autoencoder," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 30, no. 1, pp. 312–317, 2018.

[6] K. Ma, W. Liu, K. Zhang, Z. Duanmu, Z. Wang, and W. Zuo, "End-to-end blind image quality assessment using deep neural networks," *IEEE Trans. Image Process.*, vol. 27, no. 3, pp. 1202–1213, 2018.

[7] W. Hou, X. Gao, D. Tao, and X. Li, "Blind image quality assessment via deep learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 26, no. 6, pp. 1275–1286, 2015.

[8] P. Tang, X. Wang, B. Shi, X. Bai, W. Liu, and Z. Tu, "Deep fishernet for image classification," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 30, no. 7, pp. 2244–2250, 2018.

[9] W. Luo, J. Li, J. Yang, W. Xu, and J. Zhang, "Convolutional sparse autoencoders for image classification," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 7, pp. 3289–3294, 2017.

[10] N. Ahmed, T. Natarajan, and K. R. Rao, "Discrete cosine transform," *IEEE Trans. Comput.*, vol. C-23, no. 1, pp. 90–93, 1974.

[11] O. Rippel and L. Bourdev, "Real-time adaptive image compression," in *Int. Conf. Mach. Learning*, 2017, pp. 2922–2930.

[12] J. Ballé, D. Minnen, S. Singh, S. J. Hwang, and N. John- ston, "Variational image compression with a scale hyper- prior," in *Int. Conf. Learning Representations*, 2018.

[13] D. Minnen, J. Ballé, and G. D. Toderici, "Joint autoregressive and hierarchical priors for learned image compression," in *Neural Inf. Process. Syst*, 2018, pp. 10 794–10 803.

[14] C. E. Shannon, "A mathematical theory of communication," *Bell System Tech. J.*, vol. 27, no. 3, pp. 379–423, 1948.

[15] D. Marpe, H. Schwarz, and T. Wiegand, "Context-based adaptive binary arithmetic coding in the H. 264/AVC video compression standard," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 13, no. 7, pp. 620–636, 2003.

[16] F. Mentzer, E. Agustsson, M. Tschannen, R. Timofte, and L. Van Gool, "Conditional probability models for deep image compression," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 4394–4402.

[17] M. Li, K. Ma, J. You, D. Zhang, and W. Zuo, "Efficient and effective context-based convolutional entropy modeling for image compression," *IEEE Transactions on Image Processing*, vol. 29, pp. 5900–5911, 2020.

[18] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *IEEE Conf. Comput. Vis. Pattern Recog.*, vol. 2. IEEE, 2005, pp. 60–65.

[19] J. Ballé, V. Laparra, and E. P. Simoncelli, "End-to-end optimized image compression," in *Int. Conf. Learning Representations*, 2017.

[20] L. Theis, W. Shi, A. Cunningham, and F. Huszár, "Lossy image compression with compressive autoencoders," in *Int. Conf. Learning Representations*, 2017.

[21] E. Agustsson, F. Mentzer, M. Tschannen, L. Cavigelli, R. Timofte, L. Benini, and L. V. Gool, "Soft-to-hard vector quantization for end-to-end learning compressible representations," in *Neural Inf. Process. Syst*, 2017, pp. 1141–1151.

[22] G. Toderici, D. Vincent, N. Johnston, S. J. Hwang, D. Minnen, J. Shor, and M. Covell, "Full resolution image compression with recurrent neural networks," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 5435–5443.

[23] 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.*, 2018, pp. 3214–3223.

[24] T. Mikolov, M. Karafiát, L. Burget, J. Černockỳ, and S. Khudanpur, "Recurrent neural network based language model," in *Int. Speech Commun. Assoc.*, 2010.

[25] M. Sundermeyer, R. Schlüter, and H. Ney, "LSTM neural networks for language modeling," in *Int. Speech Commun. Assoc.*, 2012.

[26] A. v. d. Oord, N. Kalchbrenner, and K. Kavukcuoglu, "Pixel recurrent neural networks," in *Int. Conf. Machine Learning*, 2016, pp. 1747–1756.

[27] A. v. d. Oord, N. Kalchbrenner, O. Vinyals, L. Espeholt, A. Graves, and K. Kavukcuoglu, "Conditional image generation with PixelCNN decoders," in *Neural Inf. Process. Syst*, 2016, pp. 4797–4805.

[28] G. Toderici, S. M. O'Malley, S. J. Hwang, D. Vincent, D. Minnen, S. Baluja, M. Covell, and R. Sukthankar, "Variable rate image compression with recurrent neural networks," *arXiv:1511.06085*, 2015.

[29] N. Johnston, D. Vincent, D. Minnen, M. Covell, S. Singh, T. Chinen, S. J. Hwang, J. Shor, and G. Toderici, "Improved lossy image compression with priming and spatially adaptive bit rates for recurrent networks," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 4385–4393.

[30] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Neural Inf. Process. Syst*, 2014, pp. 2672–2680.

[31] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 7794–7803.

[32] D. Liu, B. Wen, Y. Fan, C. C. Loy, and T. S. Huang, "Non-local recurrent network for image restoration," in *Neural Inf. Process. Syst*, 2018, pp. 1673–1682.

[33] Y. Zhang, K. Li, K. Li, B. Zhong, and Y. Fu, "Residual non-local attention networks for image restoration," in *Int. Conf. Learning Representations*, 2019.

[34] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising with block-matching and 3d filtering," in *Image Processing: Algorithms and Systems, Neural Networks, and Machine Learning*, vol. 6064. International Society for Optics and Photonics, 2006, p. 606414.

[35] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Int. Conf. Medical Image Computing Computer-assisted Intervention*. Springer, 2015, pp. 234–241.

[36] X. Li, H. Chen, X. Qi, Q. Dou, C.-W. Fu, and P.-A. Heng, "H-DenseUNet: hybrid densely connected UNet for liver and tumor segmentation from CT volumes," *IEEE Trans. Medical Imaging*, vol. 37, no. 12, pp. 2663–2674, 2018.

[37] Z. Zhang, Q. Liu, and Y. Wang, "Road extraction by deep residual U-Net," *IEEE Geosci. Remote Sensing Lett.*, vol. 15, no. 5, pp. 749–753, 2018.

[38] W. Shi, J. Caballero, F. Huszar, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, "Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016.

[39] M. Courbariaux, I. Hubara, D. Soudry, R. El-Yaniv, and Y. Bengio, "Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1," *arXiv:1602.02830*, 2016.

[40] S. Zhou, Y. Wu, Z. Ni, X. Zhou, H. Wen, and Y. Zou, "DoReFa-Net: Training low bitwidth convolutional neural networks with low bitwidth gradients," *arXiv:1606.06160*, 2016.

[41] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi, "Xnor-net: Imagenet classification using binary convolutional neural networks," in *Eur. Conf. Comput. Vis.*, 2016, pp. 525–542.

[42] Z. Wang, E. P. Simoncelli, and A. C. Bovik, "Multiscale structural similarity for image quality assessment," in *37th Asilomar Conf. Signals, Syst. and Comput.*, 2003, pp. 1398–1402.

[43] N. Asuni and A. Giachetti, "TESTIMAGES: A large data archive for display and algorithm testing," *Journal of Graphics Tools*, vol. 17, no. 4, pp. 113–125, 2013.

[44] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Int. Conf. Learning Representations*, 2015.

[45] A. Skodras, C. Christopoulos, and T. Ebrahimi, "The JPEG 2000 still image compression standard," *IEEE Signal Process. Mag.*, vol. 18, no. 5, pp. 36–58, 2001.

[46] F. Bellard, "BPG image format," 2019. [Online]. Available: https://bellard.org/bpg/

[47] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 770–778.