

Image Restoration by Estimating Frequency Distribution of Local Patches

Jaeyoung Yoo^{1,2,*}Sang-ho Lee^{1,*}Nojun Kwak¹¹Seoul National University²NALBI Inc.

{yoojy31 | shlee223 | nojunk}@snu.ac.kr

Abstract

In this paper, we propose a method to solve the image restoration problem, which tries to restore the details of a corrupted image, especially due to the loss caused by JPEG compression. We have treated an image in the frequency domain to explicitly restore the frequency components lost during image compression. In doing so, the distribution in the frequency domain is learned using the cross entropy loss. Unlike recent approaches, we have reconstructed the details of an image without using the scheme of adversarial training. Rather, the image restoration problem is treated as a classification problem to determine the frequency coefficient for each frequency band in an image patch. In this paper, we show that the proposed method effectively restores a JPEG-compressed image with more detailed high frequency components, making the restored image more vivid.

1. Introduction

As multimedia and the Internet have become indispensable in our ordinary life, low quality compressed images are used more often because the quality of images and the consumption of data resource are highly correlated. In this environment, tasks such as compression artifact removal that removes artifact from a lossy-compressed image and restoration of a high quality image have recently become important areas of computer vision.

Among various image formats, JPEG is the most commonly used lossy image compression format. A lossy compression like this reduces the volume of data by permanently removing some of its information. Thus, the reverse process of restoration of an image is basically a problem of generating information that the input image does not have, and is therefore an ill-posed problem. In most cases, there are several possible output images corresponding to a given input image and the problem can be seen as a task of select-



Figure 1: From top to bottom, the first image is a JPEG compressed image with quality factor 10, the second is a restored result by our baseline method that uses a typical encoder-decoder model whose loss function is pixel-wise MSE. The third is a restored result of our method, which utilizes the classifier network, trained to predict frequency distribution of possible outputs. Our result has more detail than that of the baseline.

ing the most proper one from all the possible outputs. That is, the image restoration problem can be formulated as the problem of estimating the distribution conditioned on the input image.

Since convolutional neural networks (CNNs) have been actively applied in the field of computer vision, there have been many attempts to recover the lost information of an image due to compression, using CNN. These methods mainly approximate the mapping function from the input image to the output image using CNNs in a supervised manner. In

*J. Yoo and S. Lee equally contributed to this work.

This work was supported by NRF of Korea (2017M3C4A7077582) and NALBI Inc.

most cases, the mean squared error (MSE) or the mean absolute error (MAE) between the output and the target image is minimized. These approaches generally achieve good performances in commonly used metrics like PSNR (peak signal to noise ratio) and SSIM (structural similarity) [25], but the output images are blurry for human eyes with absence of high frequency details as can be seen in the second image of Figure 1. The reason can be attributed to the fact that the learning method tries to minimize a loss function like MSE and MAE, which is based on the pixel distance between the output and the target, forcing the model to converge to the mean or the median of all the possible solutions [2, 24, 26, 1].

In recent years, many studies have shown good results by applying generative adversarial networks (GAN) [7] to the problems of lossy compression artifact removal [9, 5]. However, since training a GAN actually requires to find a Nash equilibrium between the generator and the discriminator, the learning is unstable and difficult, thus still it is very hard to reproduce the good results reported in the original papers [6].

In this work, instead of using a generative model, we treat the image restoration problem as a classification task. The frequency distribution of the target image is directly estimated from the input image using the cross-entropy loss function. By using this information together with the existing encoder-decoder neural network model, the output image can be brought closer to a natural image. Our method, as shown in the third row of Figure 1, can generate more sharp output images with realistic details. The contribution of this paper is threefold:

1. The image restoration problem is reformulated as a task of recovering original frequency components, thus the viewpoint of an image is changed from the pixel domain to the frequency domain to explicitly recover the lost high frequency information.
2. Unlike previous works that tackle the image restoration problem by solving a regression problem or by using a generative model, we estimated the distribution of the lost information of an image by treating this task as a classification problem in the frequency domain.
3. Especially, the proposed method is applied to the tasks of JPEG compression artifact removal and the results show that our work restores high frequency components well and produces visually satisfactory outputs.

2. Related works

Since deep neural networks (DNN) have attracted researchers' interests, many studies have been conducted to remove the lossy compression artifacts with DNN, most of which focused on enhancing the quality of JPEG images.

Many have attempted to solve this problem by forwarding a lossy compressed image to a DNN to obtain a restored image directly from the output. The work in [2] is an early study of applying the DNN to the artifact removal, in which a relatively light neural network was used with pairs of a lossy compressed image and the corresponding lossless image. Better results were obtained in [24] by adding a loss function that emphasizes the edges using the Sobel filter.

Some attempts have been made to remove artifacts from lossy compression using discrete cosine transform (DCT), which is highly utilized in image compression algorithms such as JPEG and MJPEG. In lossy compression, it is used to remove high frequency components with low energy using the fact that frequency components are well separated by frequency bands. In the case of [14, 8, 26], the compression artifact removal problem was approached from both pixel domain and DCT domain using a neural network.

Most works in this line of research cast the image restoration problem in the framework of regression and tried to minimize the loss function defined as a distance in the pixel domain, which has the disadvantage that the resultant images are blurred because the neural network takes pixel-wise average [12, 23]. Some studies have attempted to solve this problem by using a classification framework directly or indirectly. Zhang *et al.* [28] showed good performances in colorization by directly classifying color pixels from grayscale pixels. Iizuka *et al.* [10] also dealt with colorization and mitigated the disadvantages of minimizing distance loss by mixing high level features learned in classification as prior knowledge.

GAN [7] is another way to solve this problem of artifact removal and image quality enhancement. The work [9] applied a GAN to remove lossy compression artifacts, which obtained better results using the GAN loss, in combination with the perceptual loss [11] and JPEG-related loss. Galteri *et al.* [5] also tackled the problem of artifact removal using the GAN. Ledig *et al.* [12] suggested SRGAN, a way to create a super resolution image using a GAN, and showed that the GAN generates more realistic images compared to the conventional methods. However, training a GAN is very difficult and unstable. Mescheder *et al.* [17] pinpoints the difficulties of convergence of GAN with the recent training algorithms and proposed a new methods.

Unlike the previous studies which apply GAN to solve the problems caused by dealing with the image restoration as a regression problem in the pixel domain, we utilize not only spatial representations but also frequency representations for the problem. More specifically, a classification method is used for estimating the frequency distribution, which is further utilized in a regression framework whose loss function is based on pixel-wise distance. Also, we propose a new architecture and training scheme to do this efficiently.

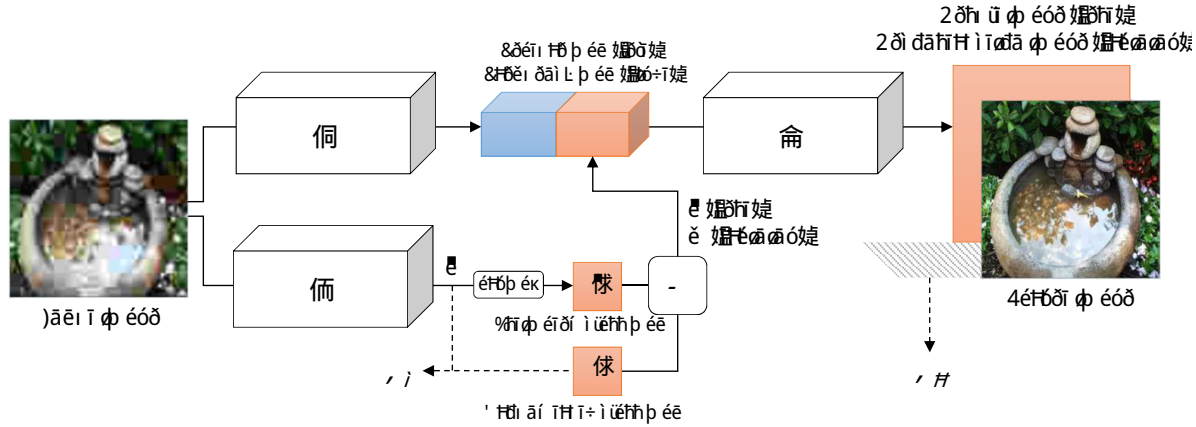


Figure 2: **The overview of the proposed image restoration process.** The proposed network is composed of an encoder (E), a decoder (D) and a classifier (C). E outputs a feature map (blue) which are considered to have low frequency information, while C outputs estimated frequency components for each image patch (orange). Taking the concatenation of the two as an input, D outputs a reconstructed image. The solid line represents the forward path and the dashed line represents the loss calculation path.

3. Estimating Frequency Distribution for Image Restoration

3.1. Problem formulation

The neural network approaches to the existing artifact removal problem mainly try to minimize the pixel-wise distance between the ground truth I^G and the restored output I^R for the JPEG-compressed input I^J in the sense of MSE [2, 24] or MAE [13]. The output images obtained with this distance-based loss are good in the sense of MSE-based metrics such as PSNR and SSIM, but they are blurry for human eyes because they are learned by taking the average of various possible solutions for I^G .

This approach can be an efficient method under the assumption that the true distribution $p(I^G|I^J)$ of the lossless image I^G corresponding to the lossy-compressed image I^J is unimodal. However, $p(I^G|I^J)$ is ambiguous, because the mapping from the I^J to the I^R is one-to-many function which involves in quantization in each of the frequency channel.

Our goal is to create non-blurry I^R , which has sharp edges with vivid details. The proposed network tackles the problem in the frequency domain as the problem of estimating DCT coefficients q of the I^G . Furthermore, instead of using the conventional MSE loss to directly estimate a single point q in the space of DCT coefficients, a network is trained to estimate the distribution $p(q|I^J)$ of q by minimizing the KL-divergence;

$$\theta = \operatorname{argmin}_{\theta} D_{KL}(p(q|I^J) \parallel \hat{p}_{\theta}(q|I^J)). \quad (1)$$

Here, \hat{p}_{θ} is the estimated distribution by the network for the

input I^J with the parameter vector θ .

The problem formulation of image restoration using KL-divergence has an advantage over that using MSE as follows: Consider two images I_1^G and I_2^G result in the same JPEG image I^J . If a network is trained to directly estimate the target using two training samples (I^J, I_1^G) and (I^J, I_2^G) , the network never learns because only a single point I^R is outputted by the network. In this case, I^R will be different from both targets I_1^G and I_2^G . On the other hand, if a distribution is trained with the same samples, the network can learn the bimodal distribution which has two peaks at I_1^G and I_2^G . If we take the maximal point of the distribution, one of the peaks will be selected and a perfect restoration is possible.

In what follows, we treat $p(q|I^J)$ as the ground truth discrete distribution of the DCT coefficient classes in each frequency channel, and restore the image by using the information on the estimated distribution $\hat{p}(q|I^J)$ as the input to the traditional encoder-decoder neural network architecture.

3.2. Overview of the proposed method

Figure 2 shows the structure of the proposed method for image restoration. As can be seen in the figure, our compressed image restoration framework consists of three networks: a classifier, an encoder and a decoder. Through classification, classifier C outputs a discrete distribution \hat{p} which contains probability of per-patch frequency coefficient class for each frequency channel. Then the class with maximum probability is written in \hat{y} . From now on, \hat{p} and \hat{y} will be referred to as the *class distribution map* and *estimated class map*, respectively. Mathematically, it becomes $\hat{p} = C(I^J)$ and $\hat{y} = \operatorname{argmax}_k \hat{p}(k)$. By a simple mapping

M , the estimated class map \hat{y} is further mapped to a map of real frequency coefficients \hat{q} , i.e. $\hat{q} = M(\hat{y})$.

The encoder E takes I^J to generate a feature map and the decoder D produces an output image I^R , which receives the detailed frequency information from the coefficients \hat{q} estimated by the classifier as well as the output feature maps from the encoder:

$$I^R = D(E(I^J), \hat{q}). \quad (2)$$

Instead of using the adversarial training scheme, the proposed method learns the network through typical supervised learning in the frequency domain. The cross entropy loss is used to train the classifier and the MSE loss is used to train the encoder and the decoder, thus the learning process is simple.

For practical reason, we define q as DCT coefficient from Laplacian image I^L of I^G , not from I^G . this preprocessing can highlight the details of an image. In our method, the wider the range of the DCT coefficient, the larger becomes the quantization error. Using I^L that has a smaller DCT coefficient range than I^G , we can reduce the quantization error and also can focus on detailed texture.

3.3. Classifier

As shown in Figure 3, the overall classifier network consists of a feature extractor F and a multi-stage part composed of T stage blocks.

The feature extractor part F creates a feature map $f \in \mathbb{R}^{n_w \times n_h \times n_f}$ that contains spatial information of an input image $I^J \in \mathbb{R}^{w \times h \times c}$. The parameters n_w and n_h are computed by $n_w = \frac{w}{w_b}$ and $n_h = \frac{h}{h_b}$, respectively, where w_b and h_b are the width and height of an image patch, w and h are the width and height of an input image, c is the number of input channels which is typically 1 (gray) or 3 (RGB).

Each stage block S^t , $t = 1, \dots, T$, receives the feature map f from the feature extractor. The first stage block S^1 generates a discrete class probability map $\hat{p}_{ch}^{(1)} \in \mathbb{R}^{n_w \times n_h \times n_{cl}}$ of the frequency coefficient corresponding to each frequency channel ch of a local patch, where n_{cl} is the number of frequency coefficient classes per channel¹. Stage blocks S^t , $t = 2, \dots, T$, concatenate the output $\hat{p}^{(t-1)}$ of the previous block S^{t-1} and the feature map f from the feature extractor as an input:

$$\begin{aligned} (t = 1) & : \hat{p}_{ch}^{(1)} = S_{ch}^1(f) \\ (t > 1) & : \hat{p}_{ch}^{(t)} = S_{ch}^t(f, \hat{p}^{(t-1)}). \end{aligned} \quad (3)$$

Here, S_{ch}^t denotes the softmax output of the stage block S^t for the frequency channel $ch \in \{1, \dots, n_{ch}\}$ and the cor-

¹If we use a 4×4 patch as a base block ($w_b = h_b = 4$), there are 16 frequency channels. In our implementation, we set $n_{cl} = 7$, thus the output dimension of each stage is $112 (= 16 \times 7)$ per an image patch.

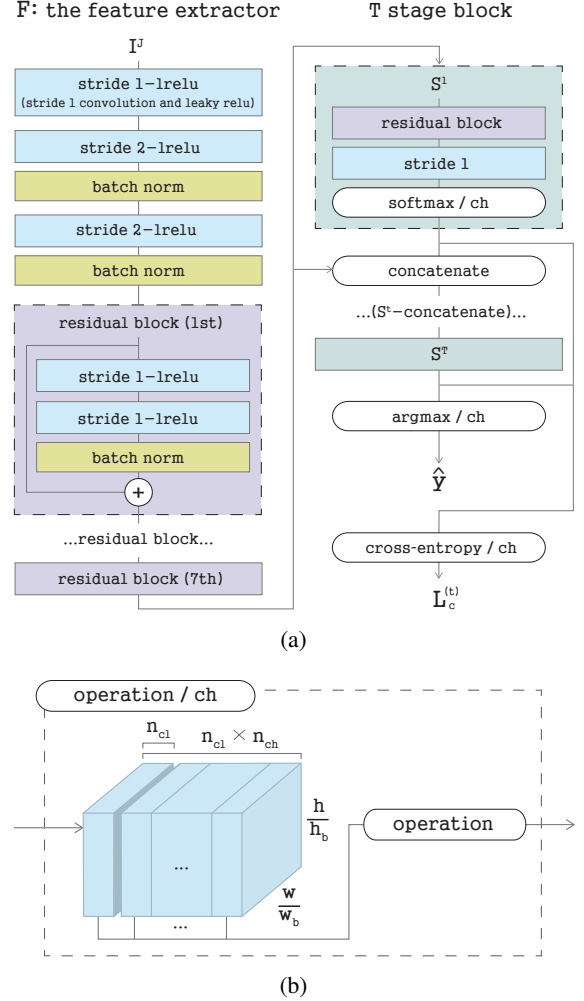


Figure 3: **The architecture of the classifier.** (a) is an illustration of the feature extractor F and the multi stage block S_{ch}^t . In F there are two stride 2 blocks because our w_b and h_b are set to 4. (b) represents per-channel operation of softmax, argmax, cross-entropy drawn as white box in (a). These operations are performed separately channel by channel.

responding $\hat{p}_{ch}^{(t)}$ can be interpreted as a probabilistic estimation of the target class label y_{ch} .

We design the multi-stage part to consider correlations between frequency channels and spatial blocks more deeply. Table 1 is the accuracy comparison of various levels of the stage blocks in the multi-stage classifier. As shown in the table, the stage 2 block shows a better result than the stage 1, but after the stage 2, the improvement is relatively low, so we decide to use only two stages ($T = 2$).

The classification result $\hat{y}_{ch} \in \mathbb{R}^{n_w \times n_h}$ is obtained by

	stage 1	stage 2	stage 3
accuracy	0.3094	0.3169	0.3178

Table 1: Classification accuracy of each stage evaluated on a three-stage classifier, trained on 96×96 luminance images.

taking the index of the maximum element of the $\hat{p}_{ch}^{(T)}$ along the class axis as follows:

$$\hat{y}_{ch} = \underset{k}{\operatorname{argmax}}(\hat{p}_{ch}^{(T)}(k)), \quad k \in \{1, \dots, n_{cl}\}. \quad (4)$$

The class \hat{y}_{ch} estimated by the classifier is converted back to the final DCT coefficient value $\hat{q}_{ch} = M_{ch}(\hat{y}_{ch})$ to represent the real information. Here, $M_{ch}(\hat{y}_{ch})$ is our class-to-coefficient mapping.

We use the cross entropy loss to train the classifier. The classification loss at each stage $L_c^{(t)}$, $t = 1, \dots, T$, is first calculated as the average of the cross entropy losses in all the spatial blocks and frequency channels using the class distribution map \hat{p}^t and the ground truth class map y . Then, these are averaged to define the final classification loss:

$$L_c = \frac{1}{T} \sum_{t=1}^T L_c^{(t)}. \quad (5)$$

The cross entropy loss is closely related to the KL-divergence and plays a very important role in our method in that it matches the predicted distribution \hat{p} with the ground truth frequency coefficient class y .

3.4. Encoder and Decoder

Like the classifier, our encoder-decoder architecture consists of several residual blocks and convolution layers. The structure of the encoder is exactly same as the feature extractor (F) of the classifier, and that of the decoder is exactly symmetric with the encoder except that there is no activation function after the output convolution layer. In the decoder, upsampling (convolution - pixel shuffle [22] - leaky ReLU [15]), instead of downsampling, is conducted. The encoder E in Fig. 2 produces a feature map of an input image, while the decoder D takes the output of the encoder E and the result of the classifier C together as an input to produce an output image.

The output of the classifier \hat{q} has frequency information that can concatenate to the feature maps induced by the input image I^J . In order for \hat{q} to be used appropriately to produce the output image, the decoder needs to learn the mapping from the frequency domain to the pixel domain. If we use estimated \hat{q} in the training of D with the general pixel-wise distance loss, the decoder has a tendency to ignore \hat{q} and mostly rely on the feature map induced by the input image I^J . This is due from the fact that \hat{q} does not



Figure 4: The effect of using ground truth q instead of classifier output \hat{q} for training the encoder-decoder. Both images are using the same classifier, but the encoder-decoder is trained by different loss. The image on the right is the output of the network trained by equation (6) using q , while the left one is the result of the network trained by \hat{q} instead of q in equation (6). In test time, both images are generated using \hat{q} . The right one has more vivid detail and lively edges.

have a perfect information on the target image I^G because of the classification error in C and this small but incorrect error in the frequency domain can cause large errors in each pixel value for the entire patch.

However, if we use the exact frequency information of the target image, the decoder can learn the mapping from the frequency to the pixel domain correctly. Therefore, in the training phase, instead of using the imperfect frequency information \hat{q} , we use q , the actual frequency information obtained from the target image I^G , in the reconstruction loss, as follows:

$$L_r = \frac{1}{wh} \sum_{x=1}^w \sum_{y=1}^h \|D(E(I^J), q)_{x,y} - I_{x,y}^G\|^2, \quad (6)$$

where the subscript x, y indicates the pixel location. The reconstruction loss L_r is back-propagated to the encoder and decoder and does not affect the classifier.

As shown in the Figure 4, the image reconstructed by the network trained with ground truth q has more vivid and natural detail. We have calculated the mean absolute values of convolution filters in the first layer of the decoder for both cases of using q and \hat{q} in the training. Then, the elements corresponding to the feature vector f from the encoder and the ones corresponding to the DCT coefficients q or \hat{q} from the classifier are summed to yield w_f and w_q , respectively. In both networks, w_f is bigger than w_q , but the difference of this value $w_f - w_q$ of using \hat{q} is about 1.74 time bigger than that using q . Therefore, we can assume that the decoder trained using \hat{q} ignores the features from the classifier much.

Since q contains DCT coefficients of the target image I^G , it is used only for training. In the test phase, the output image I^R is generated by using the estimated \hat{q} .




Figure 5: Comparison of restored result according to loss function used in learning. The first and second rows are image extracted on RGB channel. The third low is result on Y channel. The QF in jpeg images means the quality factor. In every samples, while the others’ results look blurred, ours are most clear with high frequency details and edges being the closest to the ground truths.

4. Experiments

4.1. Experiment setting

For all the experiments in this section, our baseline (ED) uses the encoder(E) and the decoder(D) networks without the classifier, trained using only the corresponding reconstruction loss. The proposed method, CED-GT, is trained and tested with the classifier (C) and the reconstruction loss. In this case, the ground truth DCT coefficient map q is used for the training of E and D . For comparison, we also trained E and D using the estimated DCT coefficient map \hat{q} and name this model as CED-EST. For the classifier, the number of stage modules (T) is set to two.

For our training data, the test set of ILSVRC 2015 [19] which contains 10^5 images was chosen. For data augmentation, these images were horizontally flipped and total 2×10^5 images were used for training. For test data, and the validation sets of BSDS500 [16] and LIVE1 [20] were used. For fair comparison with previous works, we used MATLAB JPEG encoder as others did. Depending on the

experiment, the quality factor of 10 or 20 was used.

At the training phase, the training images were resized to 96×96 or 128×128 depending on the dimension of the network. In the test phase, full-sized images were directly restored using the fully convolutional property of the network.

To generate a label for the training of the classifier network, we first extracted the luminance channel Y from the YCbCr format of the target image and then took the Laplacian to highlight the image detail. Then, the image was divided into 4×4 patches. By performing DCT on each patch, 16 channels of DCT coefficients were obtained, and then each of the 16 coefficients was labeled to one of the 7 classes. To prohibit class imbalance problem, the DCT coefficient spaces are evenly separated such that each class bin has the same number of training samples.

4.2. Qualitative result

Figure 5 is the results of the proposed method. The proposed method, CED-GT, is compared with the baseline ED




Figure 6: The distribution of some DCT coefficients of laplacian images, obtained through 2D-DCT of 8×8 image patches in LIVE1 dataset. The 1-1 means first row, first column in frequency table. we compare the ground truth, JPEG (quality factor 10) and restored image (using baseline ED and our methods CED-GT). The distributions of our method are more similar to the ground truths, yet baseline method ED is similar to the JPEG.

and SA-DCT [4]. The grayscale model of AR-CNN [2] by the author, which is publicly available is also used for comparison. Compared to the source image JPEG-10, the output of ED is closer to the ground truth image, but the details still look blurred. The proposed CED-GT shows the best result, hardly distinguishable from the ground truth with bare eyes. Our method removes a block boundary effects, as well as a ringing artifacts. For the high frequency detailed surfaces, such as feather, flowers, the results of ours have more visually plausible texture compared with those of others. Especially, the blurry edges that the ED method generates turn to very vivid ones in ours. Our method does not make such a texture randomly. For the low frequency region such as flat surface, our method does not make such a rough texture. From this, we can conjecture that the network knows where to generate the detail or not. Furthermore, the network does not just sprinkle the pattern, rather it generates visually natural edges.

4.3. Quantitative result

Frequency distribution To show our method successfully restores the frequency detail, the distribution of DCT coefficients are compared between methods. To focus on detail, we extracted DCT coefficients on Laplacian images. Figure 6 shows the distribution of frequency coefficients for various frequency channels. The distribution was obtained by applying 2D-DCT on 8×8 image patches of LIVE1 dataset [20]. Except the DC component (1-1 channel), the DCT coefficients of target images are widely distributed, whereas the distribution of JPEG image is very narrow and concentrated near zero. This tendency is strengthened with increasing frequency. This is because the JPEG compression absolutely removes its high frequency components in most of the time. As shown in Figure 6, the distribution of the images that were restored by the proposed CED-GT is more similar to that of the target image, compared to the ED baseline. On the other hand, the distribution of the ED

is similar to that of JPEG images, especially at high frequencies.

BEF and quantitative metric For quantitative comparison, we adopt the well-known traditional metrics PSNR, PSNR-B [27], and SSIM [25] scores. However, according to the previous state of the art works using GAN [12, 9, 5], this measurement may not be perfectly related to human vision. It is clear that the images generated by their methods are visually better for human than the plain MSE-based images, but their PSNR scores are generally lower than those of the MSE-based methods. The image restoration problem can be seen as estimating a solution that is plausible to human eyes, not perfectly restoring an unknown target pixel-by-pixel. Therefore, if the human visual evaluation is satisfied, it is acceptable to have a lower value for MSE-related measure.

For the above reason, the PSNR score does not fit for the assessment of modern state-of-the-art image restoration methods. Inspired by [27], we also evaluate the generated images in the BEF metric, which is defined by the difference in MSEs between block boundary pixels ($D_B(I^R)$) and non-block boundary pixels ($D_B^C(I^R)$):factor

$$\mathbf{BEF}(I^R) = \eta(D_B(I^R) - D_B^C(I^R)), \quad (7)$$

where η is a constant that depends only on the image and block sizes. If the boundary artifact effect is zero, there is clearly no reason that $D_B(I^R)$ is different with $D_B^C(I^R)$. So the lower BEF indicates lower block boundary effect. If the BEF score is normalized by MSE of an entire image (**BEF/MSE**), it can be a good measure of how much boundary effect is removed. However, as we do not know the each image's BEF of other works, we can only derive the lower bound of mean BEF/MSE via Jensen's inequality, using mean PSNR and mean PSNR-B that are provided.

Table 2 shows a quantitative comparison of various algorithms on LIVE1 and BSDS500 datasets. All the quantita-

QF	Method	LIVE1				BSDS500			
		BEF/MSE>	PSNR	PSNR-B	SSIM	BEF/MSE>	PSNR	PSNR-B	SSIM
10	JPEG	0.754	27.77	25.33	0.791	0.824	27.58	24.97	0.769
	*AR-CNN[2]	0.094	29.13	28.74	0.823	0.086	28.74	28.38	0.796
	Galteri-MSE[5]	0.067	29.41	29.13	0.832	0.089	28.93	28.56	0.805
	*Galteri-MSE[5]	0.084	29.45	29.10	0.834	0.102	29.03	28.61	0.807
	*Galteri-GAN[5]	0.148	27.29	26.69	0.773	0.178	27.01	26.30	0.746
	ED	0.074	29.40	29.09	0.833	0.094	28.96	28.57	0.806
	CED-EST	0.076	29.40	29.08	0.832	0.094	28.95	28.56	0.805
	CED-GT	0.007	26.54	26.51	0.767	0.007	26.00	25.97	0.731
20	JPEG	0.778	30.07	27.57	0.868	0.884	29.72	26.97	0.852
	*AR-CNN[2]	0.178	31.40	30.69	0.890	0.180	30.80	30.08	0.868
	Galteri-MSE[5]	0.122	31.70	31.20	0.896	0.180	31.09	30.37	0.876
	*Galteri-MSE[5]	0.125	31.77	31.26	0.896	0.180	31.20	30.48	0.832
	*Galteri-GAN[5]	0.059	28.35	28.10	0.817	0.740	28.07	27.76	0.794
	ED	0.132	31.68	31.14	0.895	0.189	31.08	30.33	0.875
	CED-EST	0.127	31.65	31.13	0.895	0.180	31.04	30.32	0.875
	CED-GT	0.002	29.33	29.32	0.854	0.009	28.62	28.58	0.825

Table 2: The lower bound of mean BEF/MSE and other quantitative result on LIVE1 and BSDS500. All the experiments are done with luminance images. The BEF/MSE of our CED-GT methods are quite lower than others, which means block artifact effect is very low. For the Galteri *et al.* [5], the asteric (*) marked results are brought directly from their paper, and the other ones are reproduced by ourselves. The results of AR-CNN [2] are also brought from [5]

tive experiments were done using the luminance (Y) channel of the original YCbCr images. The quality factors of JPEG compression were 10 and 20. The proposed methods were compared with the methods presented by Galteri *et al.* [5]. The MSE version in [5] was implemented and trained by ourselves but we could not reproduce the GAN version. In the table, the results with * mark are directly from the paper [5]. Considering the MSE score we reproduce (Galteri-MSE) is not that much different from their report (*Galteri-MSE), we believe this comparison is plausible.

As shown in Table 2 the PSNR score of the proposed CED-GT are lower than that of the other methods, but still its output images are well restored and visually better than other methods. The BEF/MSE score of CED-GT is much lower ($\times 5 \sim \times 10$) than that of the other methods. Note that the GAN-based method also has high BEF values which means that block boundary effect is not treated well.

4.4. High level tasks: detection and segmentation

If the image is successfully restored, so that the image is abundant in high frequency details, we can assume that the image can yield better performance on high level vision tasks, such as object detection and semantic segmentation. We tested Faster RCNN [18] and FCN-8s [21] as a benchmark algorithms for object detection and semantic segmentation, respectively, and compared their scores on images generated by various methods. Note that ARCNN is tested only on grayscale images.

Table 3 shows the result of detection and segmentation on VOC 2007 [3]. We subsampled the VOC 2007 to collect sample that labeled to both detection and segmentation task. All the experiments on Table 3 is performed on this sub-

set. For both detection and segmentation tasks, the proposed CED-GT obtained the highest scores of mAP and mIoU.

Method	Detection		Segmentation	
	mAP (RGB)	mAP (Y)	mIoU (RGB)	mIoU (Y)
JPEG (QF 10)	0.359	0.292	0.414	0.311
SA-DCT[4]	0.485	0.426	0.456	0.363
ARCNN[2]	-	0.429	-	0.375
Galteri-MSE[5]	0.519	0.438	0.462	0.378
ED	0.525	0.437	0.471	0.380
CED-EST	0.526	0.440	0.474	0.384
CED-GT	0.550	0.440	0.475	0.389
Original	0.705	0.637	0.631	0.556

Table 3: Object detection and semantic segmentation performance measured on the subset of Pascal VOC 2007 dataset.

5. Conclusion

In this work, we proposed a new image restoration method that is based on the estimation of the DCT coefficient distribution and showed that it can solve the JPEG artifact removal task well. The proposed architecture uses the typical encoder-decoder model in generating restored image by the help of the classifier output which is an estimated distribution of DCT coefficients of an image patch. The resultant images generated by the proposed method have good visual quality with many sharp edges. Especially, our method is very good at removing the blocking artifacts and restoring high frequency texture information.

References

- [1] L. Cavigelli, P. Hager, and L. Benini. Cas-cnn: A deep convolutional neural network for image compression artifact suppression. In *Neural Networks (IJCNN), 2017 International Joint Conference on*, pages 752–759. IEEE, 2017.
- [2] C. Dong, Y. Deng, C. C. Loy, and X. Tang. Compression artifacts reduction by a deep convolutional network. In *Computer Vision (ICCV), 2015 IEEE International Conference on*, pages 576–584. IEEE, 2015.
- [3] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, June 2010.
- [4] A. Foi, V. Katkovnik, and K. Egiazarian. Pointwise shape-adaptive dct for high-quality denoising and deblocking of grayscale and color images. *IEEE Transactions on Image Processing*, 16(5):1395–1411, 2007.
- [5] L. Galteri, L. Seidenari, M. Bertini, and A. Del Bimbo. Deep generative adversarial compression artifact removal. *arXiv preprint arXiv:1704.02518*, 2017.
- [6] I. Goodfellow. Nips 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160*, 2016.
- [7] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [8] J. Guo and H. Chao. Building dual-domain representations for compression artifacts reduction. In *European Conference on Computer Vision*, pages 628–644. Springer, 2016.
- [9] J. Guo and H. Chao. One-to-many network for visually pleasing compression artifacts reduction. *arXiv preprint arXiv:1611.04994*, 2016.
- [10] S. Iizuka, E. Simo-Serra, and H. Ishikawa. Let there be color!: joint end-to-end learning of global and local image priors for automatic image colorization with simultaneous classification. *ACM Transactions on Graphics (TOG)*, 35(4):110, 2016.
- [11] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision*, pages 694–711. Springer, 2016.
- [12] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint arXiv:1609.04802*, 2016.
- [13] B. Lim, S. Son, H. Kim, S. Nah, and K. Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 136–144, 2017.
- [14] X. Liu, X. Wu, J. Zhou, and D. Zhao. Data-driven sparsity-based restoration of jpeg-compressed images in dual transform-pixel domain. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5171–5178, 2015.
- [15] A. L. Maas, A. Y. Hannun, and A. Y. Ng. Rectifier nonlinearities improve neural network acoustic models. In *Proc. ICML*, volume 30, 2013.
- [16] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. 8th Int'l Conf. Computer Vision*, volume 2, pages 416–423, July 2001.
- [17] L. Mescheder, S. Nowozin, and A. Geiger. The numerics of gans. *arXiv preprint arXiv:1705.10461*, 2017.
- [18] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- [19] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [20] H. R. Sheikh, Z. Wang, L. Cormack, and A. C. Bovik. Live image quality assessment database release 2 (2005), 2016.
- [21] E. Shelhamer, J. Long, and T. Darrell. Fully convolutional networks for semantic segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(4):640–651, 2017.
- [22] W. Shi, J. Caballero, F. Huszár, 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1874–1883, 2016.
- [23] C. K. Sønderby, J. Caballero, L. Theis, W. Shi, and F. Huszár. Amortised map inference for image super-resolution. *arXiv preprint arXiv:1610.04490*, 2016.
- [24] P. Svoboda, M. Hradis, D. Barina, and P. Zemicik. Compression artifacts removal using convolutional neural networks. *arXiv preprint arXiv:1605.00366*, 2016.
- [25] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [26] Z. Wang, D. Liu, S. Chang, Q. Ling, Y. Yang, and T. S. Huang. D3: Deep dual-domain based fast restoration of jpeg-compressed images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2764–2772, 2016.
- [27] C. Yim and A. C. Bovik. Quality assessment of deblocked images. *IEEE Transactions on Image Processing*, 20(1):88–98, 2011.
- [28] R. Zhang, P. Isola, and A. A. Efros. Colorful image colorization. In *European Conference on Computer Vision*, pages 649–666. Springer, 2016.