

Attention-based Adaptive Selection of Operations for Image Restoration in the Presence of Unknown Combined Distortions

Masanori Suganuma^{1,2} Xing Liu¹ Takayuki Okatani^{1,2}

¹Graduate School of Information Sciences, Tohoku University ²RIKEN Center for AIP

{suganuma, ryu, okatani}@vision.is.tohoku.ac.jp

Abstract

Many studies have been conducted so far on image restoration, the problem of restoring a clean image from its distorted version. There are many different types of distortion which affect image quality. Previous studies have focused on single types of distortion, proposing methods for removing them. However, image quality degrades due to multiple factors in the real world. Thus, depending on applications, e.g., vision for autonomous cars or surveillance cameras, we need to be able to deal with multiple combined distortions with unknown mixture ratios. For this purpose, we propose a simple yet effective layer architecture of neural networks. It performs multiple operations in parallel, which are weighted by an attention mechanism to enable selection of proper operations depending on the input. The layer can be stacked to form a deep network, which is differentiable and thus can be trained in an end-to-end fashion by gradient descent. The experimental results show that the proposed method works better than previous methods by a good margin on tasks of restoring images with multiple combined distortions.

1. Introduction

The problem of image restoration, which is to restore a clean image from its degraded version, has a long history of research. Previously, researchers tackled the problem by modeling (clean) natural images, where they design image prior, such as edge statistics [7, 34] and sparse representation [1, 45], based on statistics or physics-based models of natural images. Recently, learning-based methods using convolutional neural networks (CNNs) [18, 16] have been shown to work better than previous methods that are based on the hand-crafted priors, and have raised the level of performance on various image restoration tasks, such as denoising [44, 50, 39, 51], deblurring [29, 38, 17], and super-resolution [6, 19, 53].

There are many types of image distortion, such as

Gaussian/salt-and-pepper/shot noises, defocus/motion blur, compression artifacts, haze, raindrops, etc. Then, there are two application scenarios for image restoration methods. One is the scenario where the user knows what image distortion he/she wants to remove; an example is a deblurring filter tool implemented in a photo editing software. The other is the scenario where the user does *not* know what distortion(s) the image undergoes but wants to improve its quality, e.g., applications to vision for autonomous cars and surveillance cameras.

In this paper, we consider the latter application scenario. Most of the existing studies are targeted at the former scenario, and they cannot be directly applied to the latter. Considering that real-world images often suffer from a combination of different types of distortion, we need image restoration methods that can deal with combined distortions with unknown mixture ratios and strengths.

There are few works dealing with this problem. A notable exception is the work of Yu et al. [48], which proposes a framework in which multiple light-weight CNNs are trained for different image distortions and are adaptively applied to input images by a mechanism learned by deep reinforcement learning. Although their method is shown to be effective, we think there is room for improvements. One is its limited accuracy; the accuracy improvement gained by their method is not so large, as compared with application of existing methods for a single type of distortion to images with combined distortions. Another is its inefficiency; it uses multiple distortion-specific CNNs in parallel, each of which also needs pretraining.

In this paper, we show that a simple attention mechanism can better handle aforementioned combined image distortions. We design a layer that performs many operations in parallel, such as convolution and pooling with different parameters. We equip the layer with an attention mechanism that produces weights on these operations, intending to make the attention mechanism to work as a switcher of these operations in the layer. Given an input feature map, the proposed layer first generates attention weights on the multiple operations. The outputs of the operations are mul-

plied with the attention weights and then concatenated, forming the output of this layer to be transferred to the next layer.

We call the layer *operation-wise attention* layer. This layer can be stacked to form a deep structure, which can be trained in an end-to-end manner by gradient descent; hence, any special technique is not necessary for training. We evaluate the effectiveness of our approach through several experiments.

The contribution of this study is summarized as follows:

- We show that a simple attention mechanism is effective for image restoration in the presence of multiple combined distortions; our method achieves the state-of-the-art performance for the task.
- Owing to its simplicity, the proposed network is more efficient than previous methods. Moreover, it is fully differentiable, and thus it can be trained in an end-to-end fashion by a stochastic gradient descent.
- We analyze how the attention mechanism behaves for different inputs with different distortion types and strengths. To do this, we visualize the attention weights generated conditioned on the input signal to the layer.

2. Related Work

This section briefly reviews previous studies on image restoration using deep neural networks as well as attention mechanisms used for computer vision tasks.

2.1. Deep Learning for Image Restoration

CNNs have proved their effectiveness on various image restoration tasks, such as denoising [44, 50, 39, 51, 37], deblurring [29, 38, 17], single image super-resolution [6, 19, 10, 52, 54, 53], JPEG artifacts removal [5, 9, 50], raindrop removal [35], deraining [46, 20, 49, 22], and image inpainting [33, 13, 47, 23]. Researchers have studied these problems basically from two directions. One is to design new network architectures and the other is to develop novel training methods or loss functions, such as the employment of adversarial training [8].

Mao *et al.* [28] proposed an architecture consisting of a series of symmetric convolutional and deconvolutional layers, with skip connections [36, 11] for the tasks of denoising and super-resolution. Tai *et al.* [39] proposed the memory network having 80 layers consisting of a lot of recursive units and gate units, and applied it to denoising, super-resolution, and JPEG artifacts reduction. Li *et al.* [22] proposed a novel network based on convolutional and recurrent networks for deraining, i.e., a task of removing rain-streaks from an image.

A recent trend is to use generative adversarial networks (GANs), where two networks are trained in an adversarial fashion; a generator is trained to perform image restoration, and a discriminator is trained to distinguish whether its input is a clean image or a restored one. Kupyn *et al.* [17] employed GANs for blind motion deblurring. Qian *et al.* [35] introduced an attention mechanism into the framework of GANs and achieved the state-of-the-art performance on the task of raindrop removal. Pathak *et al.* [33] and Iizuka *et al.* [13] employed GANs for image inpainting. Other researchers proposed new loss functions, such as the perceptual loss [14, 19, 23]. We point out that, except for the work of [48] mentioned in Sec. 1, there is no study dealing with combined distortions.

2.2. Attention Mechanisms for Vision Tasks

Attention has been playing an important role in the solutions to many computer vision problems [42, 21, 12, 2]. While several types of attention mechanisms are used for language tasks [40] and vision-language tasks [2, 30], in the case of feed-forward CNNs applied to computer vision tasks, researchers have employed the same type of attention mechanism, which generates attention weights from the input (or features extracted from it) and then apply them on some feature maps also generated from the input. There are many applications of this attention mechanism. Hu *et al.* [12] proposed the squeeze-and-excitation block to weight outputs of a convolution layer in its output channel dimension, showing its effectiveness on image classification. Zhang *et al.* [53] incorporated an attention mechanism into residual networks, showing that channel-wise attention contributes to accuracy improvement for image super-resolution. There are many other studies employing the same type of attention mechanism; [31, 43, 25] to name a few.

Our method follows basically the same approach, but differs from previous studies in that we consider attention over multiple different operations, not over image regions or channels, aiming at selecting the operations according to the input. To the authors' knowledge, there is no study employing exactly the same approach. Although it is categorized as neural architecture search (NAS) methods, the study of Liu *et al.* [24] is similar to ours in that they attempt to select multiple operations by computing weights on them. However, in their method, the weights are fixed parameters to be learned along with network weights in training. The weights are binarized after optimization, and do not vary depending on the input. The study of Veit *et al.* [41] is also similar to ours in that operation(s) in a network can vary depending on the input. However, their method only chooses whether to perform convolution or not in residual blocks [11] using a gate function; it does not select multiple different operations.

3. Operation-wise Attention Network

In this section, we describe the architecture of an entire network that employs the proposed operation-wise attention layers; see Fig.1 for its overview. It consists of three parts: a feature extraction block, a stack of operation-wise attention layers, and an output layer. We first describe the operation-wise attention layer (Sec.3.1) and then explain the feature extraction block and the output layer (Sec.3.2).

3.1. Operation-wise Attention Layer

3.1.1 Overview

The operation-wise attention layer consists of an operation layer and an attention layer; see Fig.2. The operation layer contains multiple parallel operations, such as convolution and average pooling with different parameters. The attention layer takes the feature map generated by the previous layer as inputs and computes attention weights on the parallel outputs of the operation layer. The operation outputs are multiplied with their attention weights and then concatenated to form the output of this layer. We intend this attention mechanism to work as a selector of the operations depending on the input.

3.1.2 Operation-wise Attention

We denote the output of the l -th operation-wise attention layer by $\mathbf{x}_l \in \mathbb{R}^{H \times W \times C}$ ($l = 1, \dots$), where H , W , and C are its height, width, and the number of channels, respectively. The input to the first layer in the stack of operation-wise attention layers, denoted by \mathbf{x}_0 , is the output of the feature extraction block connecting to the stack. Let \mathcal{O} be a set of operations contained in the operation layer; we use the same set for any layer l . Given \mathbf{x}_{l-1} , we calculate the attended value \bar{a}_l^o on an operation $o(\cdot)$ in \mathcal{O} as

$$\bar{a}_l^o = \frac{\exp(\mathcal{F}_l(\mathbf{x}_{l-1}))}{\sum \exp(\mathcal{F}_l(\mathbf{x}_{l-1}))} = \frac{\exp(a_l^o)}{\sum_{o=1}^{|\mathcal{O}|} \exp(a_l^o)}, \quad (1)$$

where \mathcal{F}_l is a mapping realized by the attention layer, which is given by

$$\mathcal{F}_l(\mathbf{x}) = \mathbf{W}_2 \sigma(\mathbf{W}_1 \mathbf{z}), \quad (2)$$

where $\mathbf{W}_1 \in \mathbb{R}^{T \times C}$ and $\mathbf{W}_2 \in \mathbb{R}^{|\mathcal{O}| \times T}$ are learnable weight matrices; $\sigma(\cdot)$ denotes a ReLU function; and $\mathbf{z} \in \mathbb{R}^C$ is a vector containing the channel-wise averages of the input \mathbf{x} as

$$z_c = \frac{1}{H \times W} \sum_{i=1}^H \sum_{j=1}^W x_{i,j,c}. \quad (3)$$

Thus, we use the channel-wise average \mathbf{z} to generate attention weights $\mathbf{a}_l = [\bar{a}_l^1, \dots, \bar{a}_l^{|\mathcal{O}|}]$ instead of using full feature maps, which is computationally expensive.

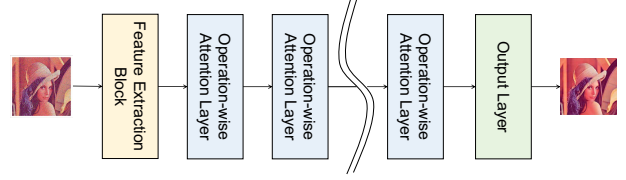


Figure 1. Overview of the operation-wise attention network. It consists of a feature extraction block, a stack of operation-wise attention layers, and an output layer.

We found in our preliminary experiments that it makes training more stable to generate attention weights in the first layer of every few layers rather than to generate and use attention weights within each individual layer. (By layer, we mean operation-wise attention layer here.) To be specific, we compute the attention weights to be used in a group of k consecutive layers at the first layer of the group; see Fig.2. Letting $l = nk + 1$ for a non-negative integer n , we compute attention weights $\mathbf{a}_{nk+1}, \dots, \mathbf{a}_{nk+k}$ at the l -th layer, where the computation of Eq.(2) is performed using different \mathbf{W}_1 and \mathbf{W}_2 for each of $\mathbf{a}_{nk+1}, \dots, \mathbf{a}_{nk+k}$ but the same \mathbf{x} and \mathbf{z} of the l -th layer. We will refer to this attention computation as *group attention*.

We multiply the outputs of the multiple operations with the attention weights computed as above. Let f_o be the o -th operation and $\mathbf{h}_l^o (\equiv f_o(\mathbf{x}_{l-1})) \in \mathbb{R}^{H \times W \times C}$ be its output for $o = 1, \dots, |\mathcal{O}|$. We multiply \mathbf{h}_l^o 's with the attention weights \bar{a}_l^o 's, and then concatenate them in the channel dimension, obtaining $\mathbf{s}_l \in \mathbb{R}^{H \times W \times C|\mathcal{O}|}$:

$$\mathbf{s}_l = \text{Concat}[\bar{a}_l^1 \mathbf{h}_l^1, \dots, \bar{a}_l^{|\mathcal{O}|} \mathbf{h}_l^{|\mathcal{O}|}]. \quad (4)$$

The output of the l -th operation-wise attention layer is calculated by

$$\mathbf{x}_l = \mathcal{F}_c(\mathbf{s}_l) + \mathbf{x}_{l-1}, \quad (5)$$

where \mathcal{F}_c denotes a 1×1 convolution operation with C filters. This operation makes activation of different channels interact with each other and adjusts the number of channels. We employ a skip connection between the input and the output of each operation-wise attention layer, as shown in Fig. 2.

3.1.3 Operation Layer

Considering the design of recent successful CNN models, we select 8 popular operations for the operation layer: separable convolutions [4] with filter sizes $1 \times 1, 3 \times 3, 5 \times 5, 7 \times 7$, dilated separable convolutions with filter sizes $3 \times 3, 5 \times 5, 7 \times 7$ all with dilation rate = 2, and average pooling with a 3×3 receptive field. All convolution operations use $C = 16$ filters with stride = 1, which is

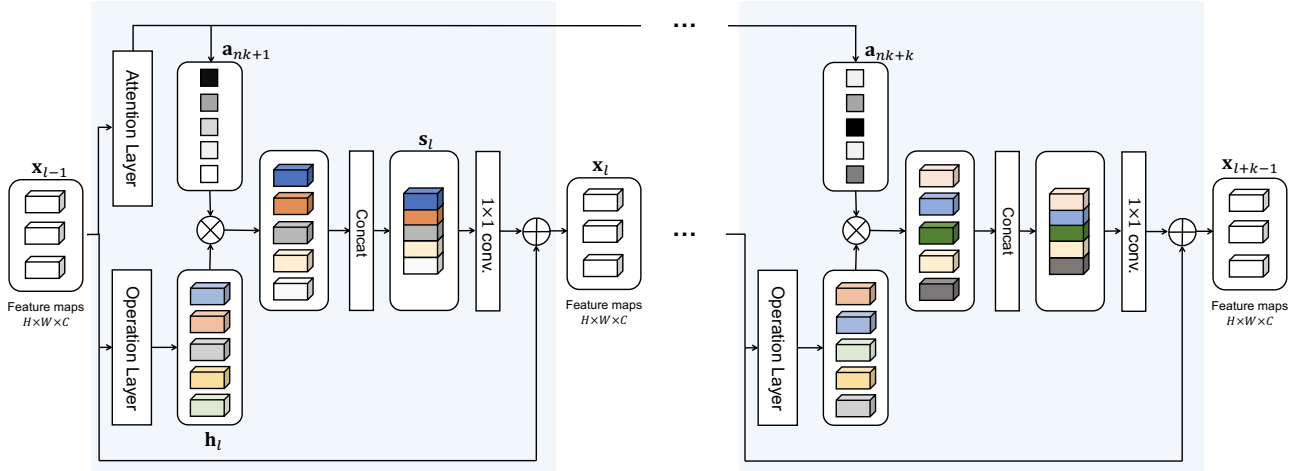


Figure 2. Architecture of the operation-wise attention layer. It consists of an attention layer, an operation layer, a concatenation operation, and 1×1 convolution. Attention weights over operations of each layer are generated at the first layer in a group of consecutive k layers. Note that different attention weights are generated for each layer.

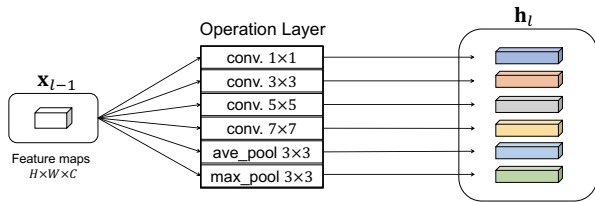


Figure 3. An example of the operation layer in the operation-wise attention layer.

followed by a ReLU. Also, we zero-pad the input feature maps computed in each operation not to change the sizes of its input and output. As shown in Fig.3, the operations are performed in parallel, and they are concatenated in the channel dimension as mentioned above.

3.2. Feature Extraction Block and Output Layer

As mentioned earlier, our network consists of three parts, the feature extraction block, the stack of operation-wise attention layers, and the output layer. For the feature extraction block, we use a stack of standard residual blocks, specifically, K residual blocks ($K = 4$ in our experiments), in which each residual block has two convolution layers with 16 filters of size 3×3 followed by a ReLU. This block extracts features from a (distorted) input image and passes them to the operation-wise attention layer stack. For the output layer, we use a single convolution layer with kernel size 3×3 . The number of filters (i.e., output channels) is one if the input/output is a gray-scale image and three if it is a color image.

4. Experiments

We conducted several experiments to evaluate the proposed method.

4.1. Experimental Configuration

We used a network with 40 operation-wise attention layers in all the experiments. We set the dimension of the weight matrices \mathbf{W}_1 , \mathbf{W}_2 in each layer to $T = 32$ and use 16 convolution filters in all the convolutional layers. For the proposed group attention, we treat four consecutive operation-wise attention layers as a group (i.e., $k = 4$).

For the training loss, we use ℓ_1 loss between the restored images and their ground truths as

$$L = \frac{1}{N} \sum_{n=1}^N \|\text{OWAN}(\mathbf{y}_n) - \mathbf{x}_n\|_1, \quad (6)$$

where \mathbf{x} is a clean image, \mathbf{y} is its corrupted version, N is the number of training samples, and OWAN indicates the proposed operation-wise attention network.

In our experiments, we employed the Adam optimizer [15] with parameters $\alpha = 0.001$, $\beta_1 = 0.9$, and $\beta_2 = 0.99$, along with the cosine annealing technique of [27] for adjusting the learning rate. We trained our model for 100 epochs with mini-batch size 32. We conducted all the experiments using PyTorch [32].

4.2. Performance on the Standard Dataset

4.2.1 DIV2K Dataset

In the first experiment on combined distortions, we follow the experimental procedure described in [48]. They use the

Table 1. **Results on DIV2K.** Comparison of DnCNN, RL-Restore, and our operation-wise attention network using DIV2K test sets. RL-Restore* displays the PSNR and SSIM values reported in [48].

Test set	Mild (unseen)		Moderate		Severe (unseen)	
Metric	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
DnCNN [50]	27.51	0.7315	26.50	0.6650	25.26	0.5974
RL-Restore* [48]	28.04	0.6498	26.45	0.5544	25.20	0.4629
RL-Restore [48]	28.04	0.7313	26.45	0.6557	25.20	0.5915
Ours	28.33	0.7455	27.07	0.6787	25.88	0.6167

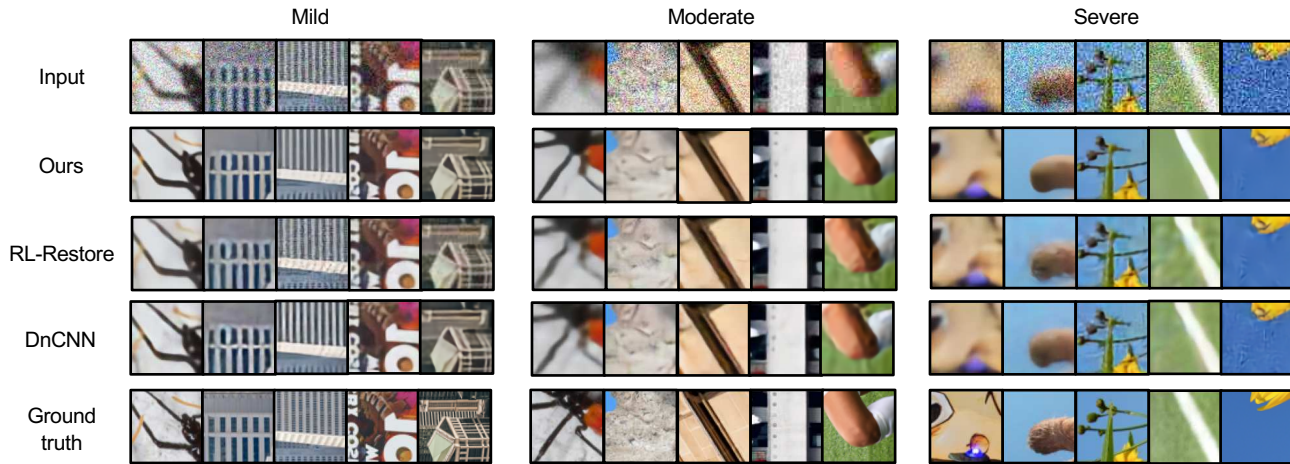


Figure 4. Examples of restored images by our method, RL-Restore [48], and DnCNN [50].

DIV2K dataset containing 800 high-quality, large-scale images. The 800 images are divided into two parts: (1) the first 750 images for training and (2) the remaining 50 images for testing. Then 63×63 pixel patches are cropped from these images, yielding a training set and a testing set consisting of 249, 344 and 3, 584 patches, respectively.

They then apply multiple types of distortion to these patches. Specifically, a sequence of Gaussian blur, Gaussian noise and JPEG compression is added to the training and testing images with different degradation levels. The standard deviations of Gaussian blur and Gaussian noise are randomly chosen from the range of $[0, 5]$ and $[0, 50]$, respectively. The quality of JPEG compression is randomly chosen from the range of $[10, 100]$. The resulting images are divided into three categories based on the applied degradation levels; mild, moderate, and severe (examples of the images are shown at the first row of Fig.4). The training are performed using only images of the moderate class, and testing are conducted on all three classes.

4.2.2 Comparison with State-of-the-art Methods

We evaluate the performance of our method on the DIV2K dataset and compared with previous methods. In [48], the authors compared the performances of DnCNN [50] and

their proposed method, RL-Restore, using PSNR and SSIM metrics. However, we found that the SSIM values computed in our experiments for these methods tend to be higher than those reported in [48], supposedly because of some difference in algorithm for calculating SSIM values. There are practically no difference in PSNR values. For fair comparisons, we report here PSNR and SSIM values we calculated using the same code for all the methods. To run RL-Restore and DnCNN, we used the author’s code for each¹.

Table 1 shows the PSNR and SSIM values of the three methods on different degradation levels of DIV2K test sets. RL-Restore* indicates the numbers reported in [48]. It is seen that our method outperforms the two methods in all of the degradation levels in both the PSNR and SSIM metrics.

It should be noted that RL-Restore requires a set of pre-trained CNNs for multiple expected distortion types and levels in advance; specifically, they train 12 CNN models, where each model is trained on images with a single distortion type and level, e.g., Gaussian blur, Gaussian noise, and JPEG compression with a certain degradation level. Our method does not require such pretrained models; it performs a standard training on a single model in an end-to-end man-

¹ <https://github.com/yuke93/RL-Restore> for RL-Restore and https://github.com/cszn/DnCNN/tree/master/TrainingCodes/dnncn_pytorch for DnCNN.

Table 2. **Results on PASCAL VOC.** Comparison of RL-Restore [48] and our method. A pretrained SSD300 [26] is applied to distorted images (“w/o restoration”) and their restored versions.

VOC	Method	mAP	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv
2007	w/o restoration	39.1	42.4	45.7	29.5	26.0	36.3	63.9	48.1	33.6	27.9	40.6	32.5	43.9	52.1	41.6	44.2	18.3	34.6	40.8	55.1	26.0
	RL-Restore [48]	66.0	74.9	78.2	60.1	43.2	49.9	80.8	72.9	68.1	51.5	69.9	63.4	76.2	78.6	82.6	68.3	41.6	50.0	69.1	75.7	65.8
	Ours	69.3	78.7	78.3	64.8	41.2	47.8	80.8	74.7	69.4	62.1	68.7	74.7	76.5	79.2	86.0	69.0	40.8	61.3	80.4	76.1	76.6
2012	w/o restoration	37.0	39.8	50.0	36.7	25.0	32.9	57.4	36.4	42.2	28.3	39.3	30.0	41.7	44.7	34.7	43.4	23.0	33.2	25.8	49.2	28.2
	RL-Restore [48]	62.0	73.4	73.8	64.0	41.2	50.5	76.7	59.3	69.8	45.8	58.8	59.9	72.2	72.3	76.6	69.3	41.3	51.0	53.0	71.5	59.6
	Ours	67.3	77.6	76.3	69.2	42.4	52.0	78.2	64.5	76.6	56.6	61.9	70.8	74.6	76.0	81.4	71.0	43.2	57.7	69.3	72.1	75.7

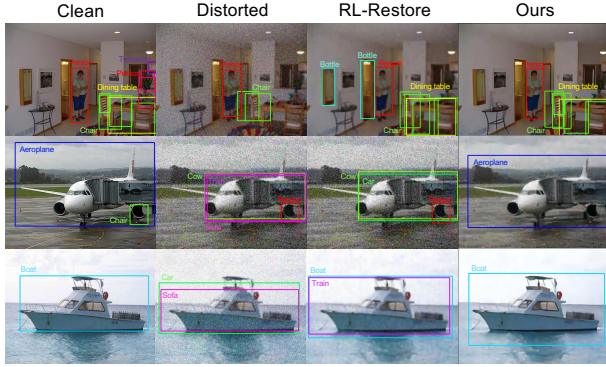


Figure 5. Examples of results of object detection on PASCAL VOC. The box colors indicate class categories.

ner. This may be advantageous in application to real-world distorted images, since it is hard to identify in advance what types of distortion they undergo.

We show examples of restored images obtained by our method, RL-Restore, and DnCNN along with their input images and ground truths in Fig.4. They agree with the above observation on the PSNR/SSIM values shown in Table 1.

4.3. Evaluation on Object Detection

To further compare the performance of our method and RL-Restore, we use the task of object detection. To be specific, we synthesized distorted images, then restored them and finally applied a pretrained object detector (we used SSD300 [26] here) to the restored images. We used the images from the PASCAL VOC detection dataset. We added Gaussian blur, Gaussian Noise, and JPEG compression to the images of the validation set. For each image, we randomly choose the degradation levels (i.e., standard deviations) for Gaussian blur and Gaussian noises from uniform distributions in the range [0, 10] and [30, 60], respectively, and the quality of JPEG compression from a uniform distribution with [10, 50].

Table 2 shows the mAP values obtained on (the distorted version of) the validation sets of PASCAL VOC2007 and VOC2012. As we can see, our method improves detection accuracy by a large margin (around 30% mAP) compared to the distorted images. Our method also provides better mAP results than RL-Restore for almost all categories. Figure 5

shows a few examples of detection results. It can be seen that our method removes the combined noises effectively, which we think contributes to the improvement of detection accuracy.

4.4. Ablation Study

The key ingredient of the proposed method is the attention mechanism in the operation-wise attention layer. We conduct an ablation test to evaluate how it contributes to the performance.

4.4.1 Datasets

In this test, we constructed and used a different dataset of images with combined distortions. We use the Raindrop dataset [35] for the base image set. Its training set contains 861 pairs of a clean image of a scene and its distorted version with raindrops, and the test set contains 58 images. We first cropped 128×128 pixel patches containing raindrops from each image and then added Gaussian noise, JPEG compression, and motion blur to them. We randomly changed the number of distortion types to be added to each patch, and also randomly chose the level (i.e., the standard deviation) of the Gaussian noise and the quality of JPEG compression from the range of [10, 30] and [15, 35], respectively. For motion blur, we followed the method of random trajectories generation of [3]. To generate several levels of motion blur, we randomly sampled the max length of trajectories from [10, 80]; see [3, 17] for the details. As a result, the training and test set contain 50,000 and 5,000 patches, respectively. We additionally created four sets of images with a single distortion type, each of which consists of 1,000 patches with the same size. We used the same procedure of randomly choosing the degradation levels for each distortion type. Note that we do not use the four single-distortion-type datasets for training.

4.4.2 Baseline Methods

We consider two baseline methods for comparison, which are obtained by invalidating some parts of the attention mechanism in our method. One is to remove the attention layer from the operation-wise attention layer. In each layer, the parallel outputs from its operation layer are simply concatenated in the channel dimension and inputted into the

Table 3. **Effects of the proposed attention mechanism.** Comparison with two baseline methods (*w/o attention* and *fixed attention*) on test sets of the new dataset we created. Each model is trained on a single set of images with combined distortions and evaluated on different sets of images with combined or single-type distortion.

Test set	Mix		Raindrop		Blur		Noise		JPEG	
Metric	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
w/o attention	23.24	0.7342	26.93	0.8393	21.74	0.7546	29.88	0.8771	29.09	0.8565
fixed attention	23.37	0.7345	26.83	0.8433	21.99	0.7590	30.58	0.8907	29.20	0.8609
Ours	24.71	0.7933	28.24	0.8764	23.66	0.8211	31.93	0.9102	30.07	0.8779

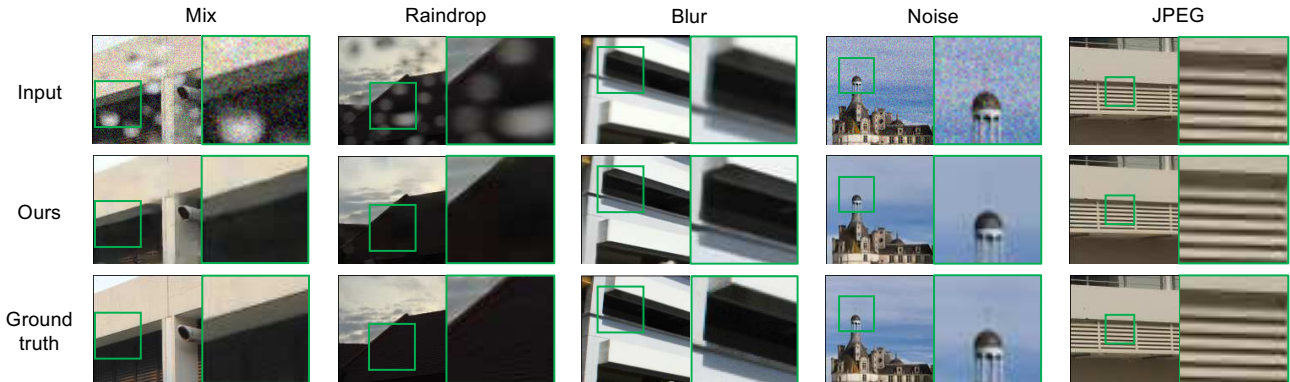


Figure 6. Examples of restored images by our method from distorted images with a different type of distortion. The model is trained on a single set of images with combined distortions.

1×1 convolution layer. Others are kept the same as in the original network. In other words, we set all the attention weights to 1 in the proposed network. We refer to this configuration as “w/o attention”.

The other is to remove the attention layer but employ constant attention weights on the operations. The attention weights no longer depend on the input signal. We determine them together with the convolutional kernel weights by gradient descent. More rigorously, we alternately update the kernel weights and attention weights as in [24], a method proposed as differentiable neural architecture search. These attention weights are shared in the group of k layers; we set $k = 4$ in the experiments as in the proposed method. We refer to this model as “fixed attention” in what follows. These two models along with the proposed model are trained on the dataset with combined distortions and then evaluated on each of the aforementioned datasets.

4.4.3 Results

The quantitative results are shown in Table 3. It is observed that the proposed method works much better than the two baselines on all distortion types in terms of PSNR and SSIM values. This confirms the effectiveness of the attention mechanism in the proposed method. We show several examples of images restored by the proposed method in Fig.6. It can be seen that not only images with combined

distortions but those with a single type of distortion (i.e., raindrops, blur, noise, JPEG compression artifacts) are recovered fairly accurately.

4.5. Analysis of Operation-wise Attention Weights

To analyze how the proposed operation-wise attention mechanism works, we visualize statistics of the attention weights of all the layers for the images with a single type of distortion. Figure 7 shows the mean and variance of the attention weights over the four single-distortion datasets, i.e., raindrops, blur, noise, and JPEG artifacts. Each row and column indicates one of the forty operation-wise attention layers and one of eight operations employed in each layer, respectively.

We can make the following observations from the map of mean attention weights (Fig. 7, left): i) 1×1 convolution tends to be more attended than the others throughout the layers; and ii) convolution with larger-size filters, such as 5×5 and 7×7 , tend to be less attended. It can be seen from the variance map (Fig. 7, right) that iii) attention weights have higher variances for middle layers than for other layers, indicating the tendency that the middle layers more often change the selection of operations than other layers.

Figure 8 shows the absolute differences between the mean attention weights over each of four single-distortion datasets and the mean attention weights over all the four datasets. It is observed that the attention weights differ de-

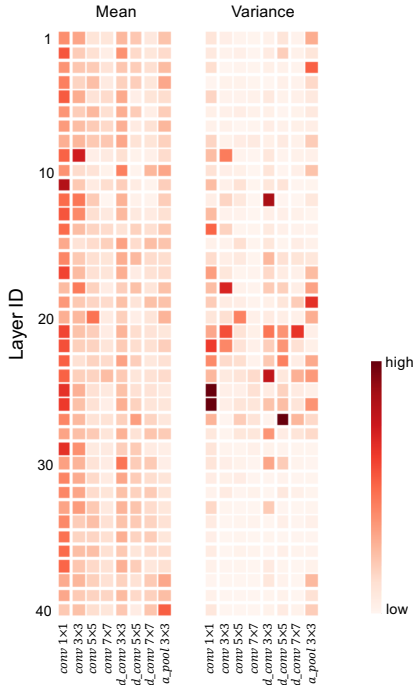


Figure 7. Mean and variance of attention weights over the input images with a single type of distortion, i.e., raindrop, blur, noise, JPEG artifacts.

pending on the distortion type of the input images. This indicates that the proposed attention mechanism does select operations depending on the distortion type(s) of the input image. It is also seen that the map for raindrop and that for blur are considerably different from each other and the other two, whereas those for noise and JPEG are mostly similar, although there are some differences in the lower layers. This implies the (dis)similarity among the tasks dealing with these four types of distortion.

4.6. Performance on novel strengths of distortion

To test our method on a wider range of distortions, we evaluate its performance on novel strengths of distortion that have not been learned. Specifically, we created a training set by randomly sampling the standard deviations of Gaussian noise and the quality of JPEG compression from $[0, 20]$ and $[60, 100]$, respectively and sampling the max length of trajectories for motion blur from $[10, 40]$. We then apply trained models to a test set with a different range of distortion parameters; the test set is created by sampling the standard deviations of Gaussian noise, the quality of JPEG compression, and the max trajectory length from $[20, 40]$, $[15, 60]$, and $[40, 80]$, respectively. We used the same image set used in the ablation study for the base image set. The results are shown in Table 4, which shows that our method outperforms the baseline method by a large margin.

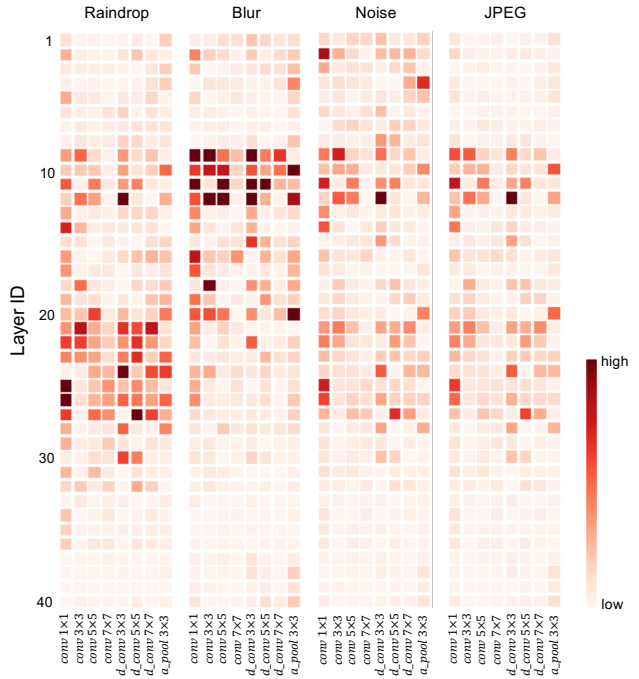


Figure 8. Visualization of mean attention weights for four distortion types. Each element indicates the difference between the mean attention weights over each distortion-type images and those over all the images (the left map of Fig.7).

Table 4. Performance on novel strengths of distortion.

Metric	PSNR	SSIM
DnCNN [50]	18.94	0.3594
Ours	23.07	0.6795

5. Conclusion

In this paper, we have presented a simple network architecture, named operation-wise attention mechanism, for the task of restoring images having combined distortions with unknown mixture ratios and strengths. It enables to attend on multiple operations performed in a layer depending on input signals. We have proposed a layer with this attention mechanism, which can be stacked to build a deep network. The network is differentiable and can be trained in an end-to-end fashion using gradient descent. The experimental results show that the proposed method works better than the previous methods on the tasks of image restoration with combined distortions, proving the effectiveness of the proposed method.

Acknowledgement

This work was partly supported by JSPS KAKENHI Grant Number JP15H05919 and JST CREST Grant Number JPMJCR14D1.

References

- [1] M. Aharon, M. Elad, and A. Bruckstein. k-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on Signal Processing*, 54(11):4311–4322, 2006. 1
- [2] P. Anderson, X. He, C. Buehler, D. Teney, M. Johnson, S. Gould, and L. Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018. 2
- [3] G. Boracchi and A. Foi. Modeling the performance of image restoration from motion blur. *IEEE Transactions on Image Processing*, 21(8):3502–3517, 2012. 6
- [4] F. Chollet. Xception: Deep learning with depthwise separable convolutions. In *CVPR*, 2017. 3
- [5] C. Dong, Y. Deng, C. Change Loy, and X. Tang. Compression artifacts reduction by a deep convolutional network. In *ICCV*, 2015. 2
- [6] C. Dong, C. C. Loy, K. He, and X. Tang. Learning a deep convolutional network for image super-resolution. In *ECCV*, 2014. 1, 2
- [7] R. Fattal. Image upsampling via imposed edge statistics. In *SIGGRAPH*, 2007. 1
- [8] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [9] J. Guo and H. Chao. Building dual-domain representations for compression artifacts reduction. In *ECCV*, 2016. 2
- [10] M. Haris, G. Shakhnarovich, and N. Ukita. Deep backprojection networks for super-resolution. In *CVPR*, 2018. 2
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2
- [12] J. Hu, L. Shen, and G. Sun. Squeeze-and-excitation networks. In *CVPR*, 2018. 2
- [13] S. Iizuka, E. Simo-Serra, and H. Ishikawa. Globally and locally consistent image completion. In *SIGGRAPH*, 2017. 2
- [14] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. 2
- [15] D. P. Kingma and J. L. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 4
- [16] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 1
- [17] O. Kupyn, V. Budzan, M. Mykhailych, D. Mishkin, and J. Matas. Deblurgan: Blind motion deblurring using conditional adversarial networks. In *CVPR*, 2018. 1, 2, 6
- [18] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 1
- [19] 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. In *CVPR*, 2017. 1, 2
- [20] G. Li, X. He, W. Zhang, H. Chang, L. Dong, and L. Lin. Non-locally enhanced encoder-decoder network for single image de-raining. In *ACMMM*, 2018. 2
- [21] K. Li, Z. Wu, K.-C. Peng, J. Ernst, and Y. Fu. Tell me where to look: Guided attention inference network. In *CVPR*, 2018. 2
- [22] X. Li, J. Wu, Z. Lin, H. Liu, and H. Zha. Recurrent squeeze-and-excitation context aggregation net for single image de-raining. In *ECCV*, 2018. 2
- [23] G. Liu, F. A. Reda, K. J. Shih, T.-C. Wang, A. Tao, and B. Catanzaro. Image inpainting for irregular holes using partial convolutions. In *ECCV*, 2018. 2
- [24] H. Liu, K. Simonyan, and Y. Yang. Darts: Differentiable architecture search. *arXiv:1806.09055*, 2018. 2, 7
- [25] N. Liu, J. Han, and M.-H. Yang. Picanet: Learning pixel-wise contextual attention for saliency detection. In *CVPR*, 2018. 2
- [26] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. Ssd: Single shot multibox detector. In *ECCV*, 2016. 6
- [27] I. Loshchilov and F. Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *ICLR*, 2017. 4
- [28] X. Mao, C. Shen, and Y. Yang. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. In *NIPS*, 2016. 2
- [29] S. Nah, T. H. Kim, and K. M. Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *CVPR*, 2017. 1, 2
- [30] D.-K. Nguyen and T. Okatani. Improved fusion of visual and language representations by dense symmetric co-attention for visual question answering. In *CVPR*, 2018. 2
- [31] J. Park, S. Woo, J.-Y. Lee, and I. S. Kweon. Bam: bottleneck attention module. In *BMVC*, 2018. 2
- [32] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. In *NIPS Workshop*, 2017. 4
- [33] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. 2
- [34] D. Perrone and P. Favaro. Total variation blind deconvolution: The devil is in the details. In *CVPR*, 2014. 1
- [35] R. Qian, R. T. Tan, W. Yang, J. Su, and J. Liu. Attentive generative adversarial network for raindrop removal from a single image. In *CVPR*, 2018. 2, 6
- [36] R. K. Srivastava, K. Greff, and J. Schmidhuber. Training very deep networks. In *NIPS*, 2015. 2
- [37] M. Suganuma, M. Ozay, and T. Okatani. Exploiting the potential of standard convolutional autoencoders for image restoration by evolutionary search. In *ICML*, 2018. 2
- [38] J. Sun, W. Cao, Z. Xu, and J. Ponce. Learning a convolutional neural network for non-uniform motion blur removal. In *CVPR*, 2015. 1, 2
- [39] Y. Tai, J. Yang, X. Liu, and C. Xu. Memnet: A persistent memory network for image restoration. In *CVPR*, 2017. 1, 2
- [40] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *NIPS*, 2017. 2
- [41] A. Veit and S. Belongie. Convolutional networks with adaptive inference graphs. In *ECCV*, 2018. 2

- [42] F. Wang, M. Jiang, C. Qian, S. Yang, C. Li, H. Zhang, X. Wang, and X. Tang. Residual attention network for image classification. In *CVPR*, 2017. 2
- [43] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon. Cbam: Convolutional block attention module. In *ECCV*, 2018. 2
- [44] J. Xie, L. Xu, and E. Chen. Image denoising and inpainting with deep neural networks. In *NIPS*, 2012. 1, 2
- [45] J. Yang, J. Wright, T. S. Huang, and Y. Ma. Image super-resolution via sparse representation. *IEEE Transactions on Image Processing*, 19(11):2861–2873, 2010. 1
- [46] W. Yang, R. T. Tan, J. Feng, J. Liu, Z. Guo, and S. Yan. Deep joint rain detection and removal from a single image. In *CVPR*, 2017. 2
- [47] R. A. Yeh, C. Chen, T. Y. Lim, A. G. Schwing, M. Hasegawa-Johnson, and M. N. Do. Semantic image inpainting with deep generative models. In *CVPR*, 2017. 2
- [48] K. Yu, C. Dong, L. Lin, and L. C. Change. Crafting a toolchain for image restoration by deep reinforcement learning. In *CVPR*, 2018. 1, 2, 4, 5, 6
- [49] H. Zhang and V. M. Patel. Density-aware single image de-raining using a multi-stream dense network. In *CVPR*, 2018. 2
- [50] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017. 1, 2, 5, 8
- [51] K. Zhang, W. Zuo, and L. Zhang. Ffdnet: Toward a fast and flexible solution for cnn based image denoising. *IEEE Transactions on Image Processing*, 27(9):4608–4622, 2018. 1, 2
- [52] K. Zhang, W. Zuo, and L. Zhang. Learning a single convolutional super-resolution network for multiple degradations. In *CVPR*, 2018. 2
- [53] Y. Zhang, K. Li, K. Li, L. Wang, B. Zhong, and Y. Fu. Image super-resolution using very deep residual channel attention networks. In *ECCV*, 2018. 1, 2
- [54] Y. Zhang, Y. Tian, Y. Kong, B. Zhong, and Y. Fu. Residual dense network for image super-resolution. In *CVPR*, 2018. 2