

# Fast and Flexible Image Blind Denoising via Competition of Experts

Shunta Maeda  
Navier Inc.

shunta@navier.co.jp

## Abstract

*Fast and flexible processing are two essential requirements for a number of practical applications of image denoising. Current state-of-the-art methods, however, still require either high computational cost or limited scopes of the target. We introduce an efficient ensemble network trained via a competition of expert networks, as an application for image blind denoising. We realize automatic division of unlabeled noisy datasets into clusters respectively optimized to enhance denoising performance. The architecture is scalable, can be extended to deal with diverse noise sources/levels without increasing the computation time. Taking advantage of this method, we save up to approximately 90% of computational cost without sacrifice of the denoising performance compared to single network models with identical architectures. We also compare the proposed method with several existing algorithms and observe significant outperformance over the standard denoising algorithm BM3D in terms of computational efficiency.*

## 1. Introduction

While a lot of image denoising algorithms have been developed over the past decades, blind removal of real noises still remains challenging. Blind denoising aims to recover a clean image from a noisy one without referring to any *a priori* information about the noise. This restriction can be widely seen as a realistic requirement of practical use cases. Fast processing is also important because the denoising often constitutes a crucial part of pre-processing pipelines in many vision tasks [31, 9]. However, current state-of-the-art methods still require either high computational cost or limited scopes of the target, unfitting for many real-world applications.

Image denoising methods are broadly classified into two categories: image prior based and discriminative learning based methods. The image prior based methods, which explicitly define models of noises, have achieved remarkable results [4, 8, 2, 13]. Nevertheless, their recovering perfor-

mances are limited because noise models are usually introduced manually. In contrast, discriminative learning based methods learn the underlying mapping between clean images and noisy ones by exploiting large modeling capacity of deep convolutional neural networks (CNNs) [33, 34, 15]. Although learning based methods demonstrate advanced performances if the evaluation data resembles the training data used, they are often outperformed by image prior based methods when tested on actual noisy images [27, 1]. This drawback of learning based methods is mainly due to a lack of high-quality denoising dataset. Some notable recent works have successfully produced such high-quality denoising dataset, for example, by systematic procedures estimating ground-truth images of noisy ones [1] or to create real noisy images utilizing generative adversarial networks [6].

Even if ideal pairs of clean images and noisy ones are available, learning based methods are usually inferior to image prior based methods in terms of adaptability for data outside the training domain. A possible solution overcoming this limitation is preparing a vast amount of clean-noisy image pairs across a variety of training domains. However, a large scale network is required to handle such dataset, resulting in a high computational cost. Taking into account the above, one natural approach would be this: training a specialized expert for each representative domain in given dataset and selecting a few suitable experts for the evaluation. Such an approach is widely known as mixture-of-experts (MoE) [18, 19]. An underlying difficulty of this approach is that real noisy images are usually unlabeled. Moreover, even though they are labeled, MoE trained with given labels is not guaranteed to minimize the task-oriented loss. Thus an unsupervised clustering method minimizing the loss of the denoising with MoE is desired.

In this work, we propose a stable method to train specialized experts via competition. The training procedure is designed for low-level vision tasks, and the network architecture is composed for fast processing applications. We apply the method to image blind denoising and realize automatic division of unlabeled noisy datasets into appropriate domains. To evaluate the effectiveness of the proposed method, we compare it with single network models using

identical architecture. Our extensive experiments reveal how the performances vary according to the number of experts and the size of the expert network. In addition, we examine perceptual quality along with various noise levels. We also compare the proposed method with several existing algorithms and observe significant outperformance over the standard denoising algorithm BM3D in terms of computational efficiency.

We highlight the notable contributions of this work:

- By utilizing our competition training method, unlabeled noisy dataset is automatically divided into appropriate clusters which are optimized respectively to enhance denoising performance.
- The proposed method is up to  $10\times$  faster than single network baseline without sacrifice of the denoising performance (Fig. 6).
- The proposed method is  $5\times$  faster than BM3D, which is a state-of-the-art patch-based algorithm, when comparing at similar denoising performance (Tab. 1).
- Our formulation is applicable to a variety of other low-level vision tasks with different model networks and loss functions.

## 2. Related Work

### 2.1. Image Prior Based Denoising

Image priors are utilized in most of traditional denoising methods. Among them, patch-based denoising methods, such as non-local means (NL-means) [4], BM3D [8], and WNNM [13] have achieved remarkable results. These methods explicitly define noise models based on *a priori* knowledge, and thus they are applicable to denoising problems with unknown noise levels. BM3D is often referred to as a benchmark because it is still one of the best denoising algorithms in terms of quality and computation time [1]. At first in BM3D algorithm, image fragments similar to a target patch are grouped into three-dimensional arrays. After that, the arrays are processed by collaborative filtering and then transformed back into the original two-dimensional form. While BM3D is one of the fastest patch-based algorithms, patch-based algorithms are generally time-consuming. In addition, its recovering performance is limited due to heuristically defined noise models.

### 2.2. Discriminative Learning Based Denoising

Learning based methods, which exploit large modeling capacity of deep CNNs, have significantly improved denoising performances by learning the underlying mapping of clean-noisy image pairs. In particular, DnCNN [33] achieves state-of-the-art results by applying residual learning for a deep CNN with batch normalization. DnCNN is

often referred to as a benchmark of learning based methods due to its simple structure and impressive performance. However, their performances are not robust in many cases to the data outside training domains and are often outperformed by image prior based methods when applied to real datasets of noisy images [1]. Several works have been conducted to relax the above limitation by feeding information about the noise level to the denoising network. FFDNet [34] non-blindly removes additive white Gaussian noise (AWGN) with a tunable noise level map and is one of the most efficient networks. CBDNet [15] can blindly remove actual noises by estimating the noise levels through a noise estimation subnetwork. Nevertheless, flexibility and computational efficiency of these networks are still inadequate because they feature a single denoising network.

### 2.3. Mixture of Experts

MoE is an approach that replaces a single network by a weighted sum of expert networks for further improvement of prediction performance [18, 19]. In MoE, a trainable gating network plays a pivotal role; it activates experts and computes which experts to be weighted and how much. Instead of activating all experts uniformly, selective activation of necessary experts at the gate enables to improve the network efficiency even more [3, 30]. MoE is effective if given dataset can be classified into several different source domains. However, appropriate partitioning of unlabeled datasets into subsets for each expert remains an open problem, and the subject relates to unsupervised domain adaptation [14].

### 2.4. Competition of Experts

Some recent studies have employed competition of experts [22, 26]. Lee *et al.* [22] proposed Stochastic Multiple Choice Learning (sMCL) which minimizes oracle loss in ensembles of experts via competition of experts. They applied sMCL for image classification, segmentation, and caption generation tasks. While all experts are assigned for the prediction in sMCL, we use a single expert for a prediction to save the computational cost. Moreover, the expert initialization and the data sampling method are also different. In the context of causality, Parascandolo *et al.* [26] proposed an algorithm to recover inverse mechanisms from transformed data points through competition of experts. Their algorithm implicitly assumes that the training data is generated by known transformations, the series of procedures referred to in their paper, whose computations of transformation are distinctly different from each other. Thus it does not fit with a number of practical applications.

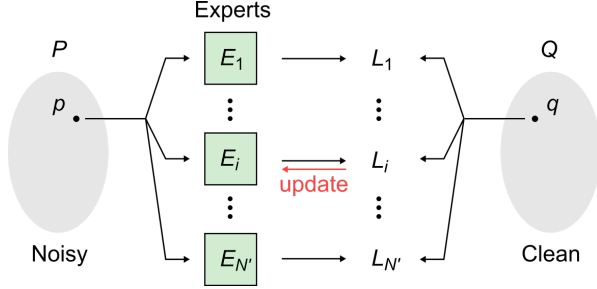


Figure 1: Schematic illustration of training procedure with competition of experts. After preparing experts well trained for dataset  $P$  and  $Q$ , noisy image  $p$  is fed to all experts independently, and only parameters of an expert minimizing the loss are updated.

### 3. Method

#### 3.1. Formulation

The proposed training procedure is schematically illustrated in Fig. 1. With given noisy images  $P = \{p_i\}_{i=1}^n$  and corresponding ground-truth images  $Q = \{q_i\}_{i=1}^n$ , our objective is learning a generalized mapping from distribution  $\mathcal{D}_P$  to distribution  $\mathcal{D}_Q$  without any *a priori* information about the noise, where  $P \subset \mathcal{D}_P$  and  $Q \subset \mathcal{D}_Q$ . Additionally, it is required for practical applications that the mapping is executed with as few parameters as possible. To achieve this, we train MoE via competition of experts and employ a single network for the evaluation. The training with the competition consists of a three-step procedure:

1. Train all experts  $E_1, \dots, E_{N'}$  for a whole of the dataset  $P$  and  $Q$  until convergence.
2. For each iteration, feed noisy images  $p$  to all experts independently and update only parameters of an expert minimizing

$$\mathcal{L}_j = \mathcal{L}(E_j(p), q), \quad (1)$$

where  $\mathcal{L}$  is a given loss function.

3. Train a gating network with the index of a winning expert  $\tilde{j} = \arg \min_{1 \leq j \leq N'} \mathcal{L}_j$  as a target label. Note that this step can be run in parallel with step 2 because the gating network is independent of each expert.

With this manner, noisy images  $P$  are automatically divided into  $N$  clusters as a result of the specialization of experts through the catastrophic forgetting [10]. We note that  $N' (\geq N)$  is a parameter to be arbitrarily set in advance. This training procedure is potentially applicable to any network trained with any task dependent loss.

#### 3.2. Stabilizing Training

**Initialize method.** An appropriate preparation of experts is important to stabilize the competition. If we start competition without pre-training (*i.e.*, start from step 2

---

#### Algorithm 1 Training with the competition of experts

---

**Require:**  $\{p, q\}$ : an image pair sampled from input dataset  $P$  and target dataset  $Q$ ,  $\{E_i\}_{i=1}^{N'}$ : experts with weight parameters  $\{\theta_i\}_{i=1}^{N'}$  ( $N'$ : number of experts),  $N_p$ : number of patches sampled from an image,  $\mathcal{L}$ : loss function

- 1: **repeat**
  - 2:   Sample patches  $\{[p^1, \dots, p^{N_p}], [q^1, \dots, q^{N_p}]\}$  ( $:= \{\mathbf{p}, \mathbf{q}\}$ ) from  $\{p, q\}$
  - 3:   Compute  $\mathcal{L}_1 = \mathcal{L}(E_1(\mathbf{p}), \mathbf{q})$
  - 4:   Update  $\theta_1$  via backpropagating gradients of  $\mathcal{L}_1$
  - 5: **until** Converge
  - 6:  $\{\theta_i\}_{i=2}^{N'} \leftarrow \theta_1$
  - 7: **repeat**
  - 8:   Sample patches  $\{\mathbf{p}, \mathbf{q}\}$  from  $\{p, q\}$
  - 9:   Compute  $\{\mathcal{L}_i\}_{i=1}^{N'} = \{\mathcal{L}(E_i(\mathbf{p}), \mathbf{q})\}_{i=1}^{N'}$
  - 10:   Let  $l = \arg \min_{1 \leq i \leq N'} \mathcal{L}_i$   $\triangleright$  use a smallest index if duplicates
  - 11:   Update  $\theta_l$  via backpropagating gradients of  $\mathcal{L}_l$
  - 12: **until** Converge
- 

instead of step 1), the model fails to classify the input dataset. This is due to that only the expert winning at the initial stage of the competition continues to be updated throughout the subsequent training. To train a classifier and experts without such unbalanced updates, we need to prepare experts trained until the convergence for a given dataset (step 1). We firstly train an expert  $E_1$  until convergence, after that just copy its parameters to all other experts  $E_2, \dots, E_{N'}$  to save training cost. Note that if two or more experts output the same result, we only update the expert with the smallest index among them.

**Global identity assumption.** Further stabilizing the competition, we introduce an assumption that any region in a noisy image  $p$  belongs to the same domain in the dataset. We refer to this assumption as *global identity assumption*. With this assumption, we randomly sample several patches  $\{[p^1, \dots, p^{N_p}], [q^1, \dots, q^{N_p}]\}$  from an image pair  $\{p, q\}$  and group them into a mini-batch for each iterations, where  $N_p$  is a number of patches sampled from an image. Under the *global identity assumption*, Eq. 1 is rewritten as

$$\mathcal{L}_j = \sum_{k=1}^{N_p} \mathcal{L}(E_j(p^k), q^k). \quad (2)$$

This assumption is necessary to perform a mini-batch processing in a training with the competition because we assume the input dataset is unlabeled. We describe an overall procedure of the training with the competition of experts under the *global identity assumption* in Algorithm 1.

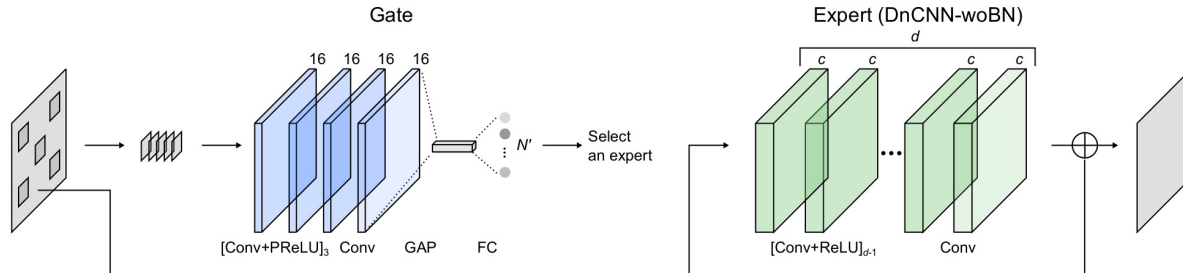


Figure 2: Overview of the evaluation network. Several parts of an input image, which grouped into a mini-batch, are fed into the gating network to select a suitable expert. After that, the whole of the input image is processed by the selected expert to remove the noise.

## 4. Implementation

### 4.1. Network Design

The expert network used in our experiment is based on DnCNN [33]. The each layer of DnCNN are combinations of  $3 \times 3$  convolution with zero-padding, ReLU [11], and Batch Normalization (BN) [17]. Unlike the original network, we do not utilize BN because we make a mini-batch from a single image in the training dataset (*i.e.*, *global identity assumption*), and the mini-batch size is rather small. In addition, we adopt several different combinations of network depth and number of filters for our experiments. Here, we refer to our modified version of the original DnCNN as DnCNN-woBN- $dXcY$ , where  $X$  and  $Y$  denote the network depth and the number of filters, respectively. Note that the original DnCNN consists of 17 and 20 convolution layers with 64 filters for specific noise and blind noise level tasks, respectively.

A gating network which selects a promising expert for an input image is required in the evaluation. A proposed gating network consists of a plain four-layer CNN followed by a global average pooling (GAP) [23] and a fully connected layer. Each convolution layer has 16 filters with a kernel size of  $3 \times 3$ , and the first three layers are followed by Parametric ReLU [16]. Since we adopt the *global identity assumption* in the training, it is also possible to sample parts of an input image in the evaluation, reducing the computational cost of the gating network. The gating and expert network architectures are illustrated in Fig. 2.

### 4.2. Datasets

We used 400 images from BSDS500 [25] and 4,744 images of WED [24] for the training. The noisy images were generated by synthesizing AWGNs or JPEG compression noises. During training, one of the two noise sources was randomly selected with equal probability for an iteration. The noise level  $\sigma$  of AWGN was uniformly selected from a range of  $[0, 55]$ , and the quality factor  $q$  of JPEG compression was uniformly selected from a range of  $[5, 100]$ .

For the evaluation, 136 images consist of two sets of BSD68 [28] was utilized. One of the test sets was gener-

ated by adding AWGNs to the BSD68 with the noise level of  $\sigma = \{n \times 55/68\}_{n=0}^{67}$ , where  $n$  is the index of an image in BSD68 (*i.e.*, the images have different noise levels at an interval of  $55/68$ ). Similarly for the another test set, quality factors of the JPEG compression were arranged as  $q = \{5+n \times 95/68\}_{n=0}^{67}$ . Here the compression qualities are rounded to integer values. We refer a test dataset consists of the above two sets of images (*i.e.*, a total of 136 images) to BSD68-GJ. Note that some of the following evaluations were performed on the plain BSD68 dataset with specific noise levels.

### 4.3. Training and Inference

We cropped 16 patches (the size of each patch is  $64 \times 64$ ) from an image and grouped them into a mini-batch to train the experts (*global identity assumption*). After initializing each expert with an identical pre-trained weight, the experts were specialized via a competition, *i.e.*, only the expert which led to the lowest error for each mini-batch was updated. During the training, the gating network was simultaneously trained by using the index of the winner expert as the target label. Owing to this, we can conduct validation at any training stage. The MSE loss and cross entropy loss were used for training experts and gating network, respectively. We used Adam optimizer [20] with a fixed learning rate of  $1 \times 10^{-4}$  for both gating and expert networks. It should be noted that our DnCNN-woBN- $d17c64$  trained for a specific level AWGN ( $\sigma = 25$ ) without competition shows similar peak signal-to-noise ratio (PSNR) with the original work: PSNR (ours) = 29.18 dB, PSNR (original work [33]) = 29.23 dB. For a fair comparison, we also trained our model only with BSD400 dataset as in the original work [33]. The performance drop was negligible ( $-0.01$  dB).

For the evaluation, five non-overlapping patches (the size of each patch was  $64 \times 64$ ) cropped from a test image were fed into the gating network. Then a suitable expert was selected according to an average of the outputs of the gate. After that, the selected expert processes the original image to recover a corresponding ground-truth image. The overall evaluation network is illustrated in Fig. 2.

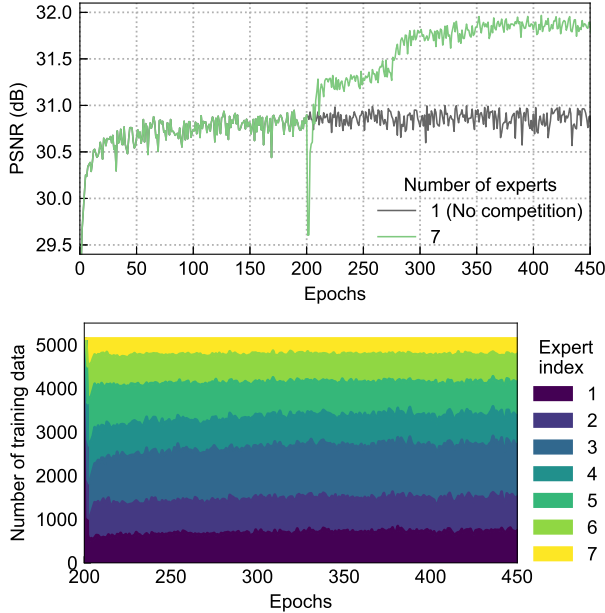


Figure 3: **Top**: Recorded performance of DnCNN-woBN- $d5c16$  for BSD68-GJ obtained with (green) and without (gray) competition. For the condition with the competition, the number of experts is set to seven and the competition starts from 201st epoch after original expert converges well. **Bottom**: Recorded transitions of winning numbers of the seven experts within the latest epoch. The experts differentiate into each branch after starting the competition, and the winning numbers finally settle down to a constant ratio.

## 5. Experiments

### 5.1. Details of Competition of Experts

Figure 3 (Top) shows the historical average PSNR during the training with and without competition. The DnCNN-woBN- $d5c16$  was used as the expert networks, and the PSNR was computed for the BSD68-GJ at each training epoch. For the condition with the competition, the number of experts  $N'$  was set to seven and the competition started from 201st epoch after the original expert converged adequately. The result demonstrates the effectiveness of the proposed method by about 1 dB improvement in the average PSNR. Figure 3 (Bottom) shows historical transitions of the winning numbers of each expert within the latest epoch. The experts differentiated into respective branches within 10 epochs after starting the competition. After that, the winning numbers gradually transited, and finally settle down to a constant ratio. The PSNR increased in the regime where the ratio of the winning numbers changed, indicating that the training data was spontaneously separated into appropriate domains to enhance the denoising performance. We note that such a transition behavior during training was reproduced for several executions, showing the stability of the proposed method.

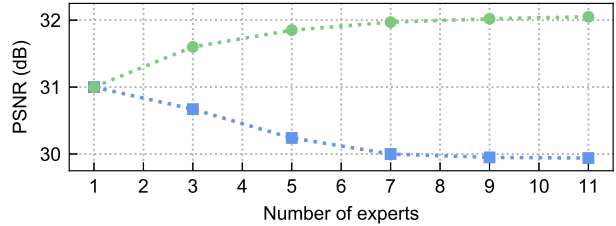


Figure 4: Performances of DnCNN-woBN- $d5c16$  for BSD68-GJ as a function of the number of experts  $N'$ . The results of the single expert selected by the gate are shown with the green circle. The average PSNR steadily increases as more experts are utilized. Ensemble results averaging all the experts' outputs are shown with the blue square. Performance drops by the ensemble indicate the success of the expert specialization.

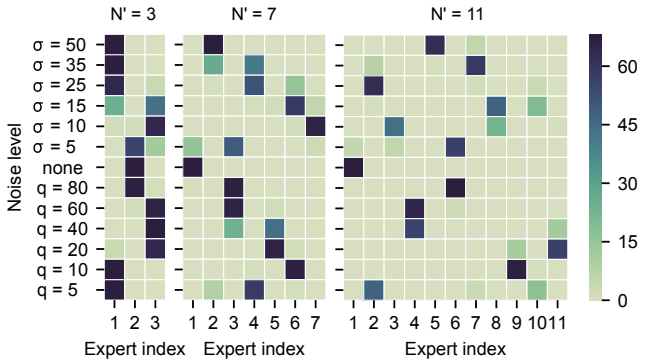


Figure 5: Visualization of expert assignment for different noise sources/levels in three conditions of  $N' = 3, 7, 11$ . The test images of each row are generated from BSD68 by adding a specific level AWGN ( $\sigma = 5, 10, 15, 25, 35, 50$ ) or compressing via a specific JPEG quality ( $q = 80, 60, 40, 20, 10, 5$ ). The test images are assigned to an index of expert which maximizes PSNR.

### 5.2. Number of Experts

Figure 4 shows the average PSNR of DnCNN-woBN- $d5c16$  for the BSD68-GJ as a function of the number of experts  $N'$ . The average PSNR steadily increases as more experts are utilized, and it started to seem saturated as the number of experts exceeds seven. Figure 5 visualizes which expert maximized performance for different noise sources/levels in three conditions of  $N' = 3, 7, 11$ . The test images of each row were generated from BSD68 by adding a specific level AWGN ( $\sigma = 5, 10, 15, 25, 35, 50$ ) or compressing via a specific JPEG quality ( $q = 80, 60, 40, 20, 10, 5$ ). The test images were assigned to an index of expert which maximized PSNR. In our experiment, training data did not have distinct domains because the noise level was gradated via the random uniform selections. In consequence, as shown in Fig. 5, the more experts were implemented, the more finely the training domains were divided. Such a fine segmentation, however, might trigger overfitting and/or lead to failure in the selection of experts, and thus it might limit denoising performance (see Fig. 4).

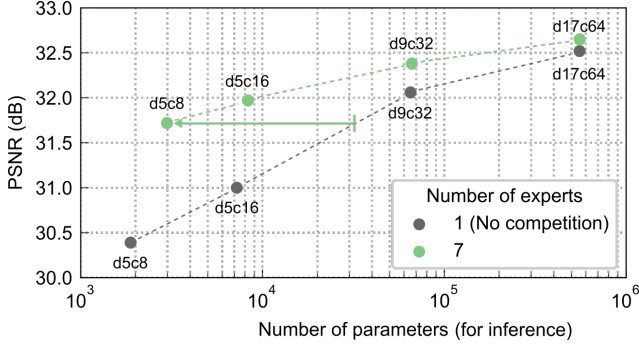


Figure 6: PSNR vs. network complexity curve. Four different expert networks DnCNN-woBN-*d5c8*, *d5c16*, *d9c32*, and *d17c64* are utilized and evaluated for BSD68-GJ with (green) and without (gray) competition. The number of experts is seven for the condition with the competition. As depicted by the arrow, the total complexity is reduced by up to approximately 90% without sacrifice of the denoising performance.

### 5.3. Computational Efficiency

Figure 6 shows PSNR vs. network complexity curve (complexity means the total number of parameters). The four different expert networks DnCNN-woBN-*d5c8*, *d5c16*, *d9c32*, and *d17c64* were utilized in this case and evaluated for BSD68-GJ with and without competition. The number of experts was seven for the condition with the competition. The total complexity of this model is defined as a sum of the complexity of the gating network and the expert. Here the gate complexity is calculated as the product of the original gate complexity and the area ratio of sampled patches to the whole of the input image. As depicted by the arrow in Fig. 6, a total complexity was saved up to approximately 90% by the competition of experts while the accuracy was maintained. The competition of experts is more efficient and effective than models without competition in the low-complexity regime. This result indicates our method is suitable for fast processing applications. Nevertheless, the proposed approach is also beneficial for the *d17c64* corresponding to the standard DnCNN (by +0.13 dB of PSNR).

### 5.4. Visual Comparison

Figure 7 shows visual comparisons between with and without competition along with different levels of AWGNs and JPEG compression noises. The employed architecture was DnCNN-woBN-*d5c16*, and the number of experts was seven for the competition. For almost all noise levels, the results obtained through competition demonstrate better perceptual quality than that of the single model. In particular, the model trained via competition preserves fine details in the ground-truth images for the low-noise conditions while the single model overly denoises images and tends to blur characteristic details of the ground-truth images.

Mode	Non-Blind		Blind (for AWGN and JPEG)	
Method	NL-means <sup>1</sup>	BM3D <sup>2</sup>	d5c16	d9c32
$\sigma = 15$	29.78	31.10	30.60	<b>31.15</b>
$\sigma = 25$	27.43	28.63	28.25	<b>28.64</b>
$\sigma = 50$	24.62	25.69	25.25	<b>25.79</b>
Time (s)	0.595	4.853	<b>0.137</b>	0.952

Mode	Non-Blind				
Method	WNNM [13]	EPLL [35]	MLP [5]	TNRD [7]	DnCNN [33]
$\sigma = 15$	31.37	31.21	-	31.42	<b>31.73</b>
$\sigma = 25$	28.83	28.68	28.96	28.92	<b>29.23</b>
$\sigma = 50$	25.87	25.67	26.03	25.97	<b>26.23</b>
Time (s)	1572.22	353.37	100.02	<b>8.007</b>	28.14

Table 1: **Top:** Average PSNR and execution time for BSD68 (image size is  $481 \times 321$ ) synthesized with specific level AWGNs. The number of experts is seven for our methods DnCNN-woBN-*d5c16* and *d9c32*. All of the algorithms are executed by a single-threaded CPU, and the test images are processed one by one. **Bottom:** PSNR scores reported in some more non-blind baselines. Run times are estimated from the run time of BM3D on our machine (top table) and the reported run times in Abdelhamed *et al.* [1].

Mode	Blind (for JPEG)	Blind (for AWGN and JPEG)	
Method	Knusperli	d5c8	d5c16
$q = 10$	27.14 / 0.7824	27.46 / 0.7874	<b>27.82 / 0.7948</b>
$q = 20$	29.08 / 0.8536	29.86 / 0.8675	<b>30.18 / 0.8719</b>
$q = 40$	31.22 / 0.9041	32.31 / <b>0.9176</b>	<b>32.38 / 0.9162</b>
Time (s)	0.108	<b>0.068</b>	0.138

Table 2: Average PSNR, SSIM, and execution time for BSD68 (image size is  $481 \times 321$ ) compressed by specific JPEG qualities. The number of experts is seven for our methods DnCNN-woBN-*d5c8* and *d5c16*. All of those algorithms are executed by a single-threaded CPU, and the test images are processed one by one.

### 5.5. Comparison with Prior Arts

We compare our model ( $N' = 7$ ) trained for both of the AWGN and JPEG noise removal with several existing denoising methods. As a common benchmark, we used plain BSD68 dataset with specific noise levels. For AWGN removal, two non-blind patch-based methods were employed as the benchmark: NL-means algorithm implemented in Scikit-image [29] and C++ implementation of BM3D [21]. For JPEG compression noise removal, Knusperli [12] a deblocking JPEG decoder recently developed at Google Open Source was employed as the benchmark. The execution times of all algorithms were based on a single thread on the laptop machine equipped with an Intel Core i5 CPU @ 2.3GHz with 16GB of memory.

<sup>1</sup>We applied fast-mode NL-means of the Scikit-image with the known noise level  $\sigma$  (*i.e.*, non-blind mode). For each noise levels, the  $h$  parameter computing the decay in patch weights was hand-tuned from  $0.3 \times \sigma$  to  $0.9 \times \sigma$  at interval of  $0.1 \times \sigma$  to demonstrate the approximately best-case performance.

<sup>2</sup>We used the DCT transform for the step 1 algorithm and the Bior transform for the step 2 algorithm without using any other options. We found this setting demonstrated the best performance in terms of computational efficiency.

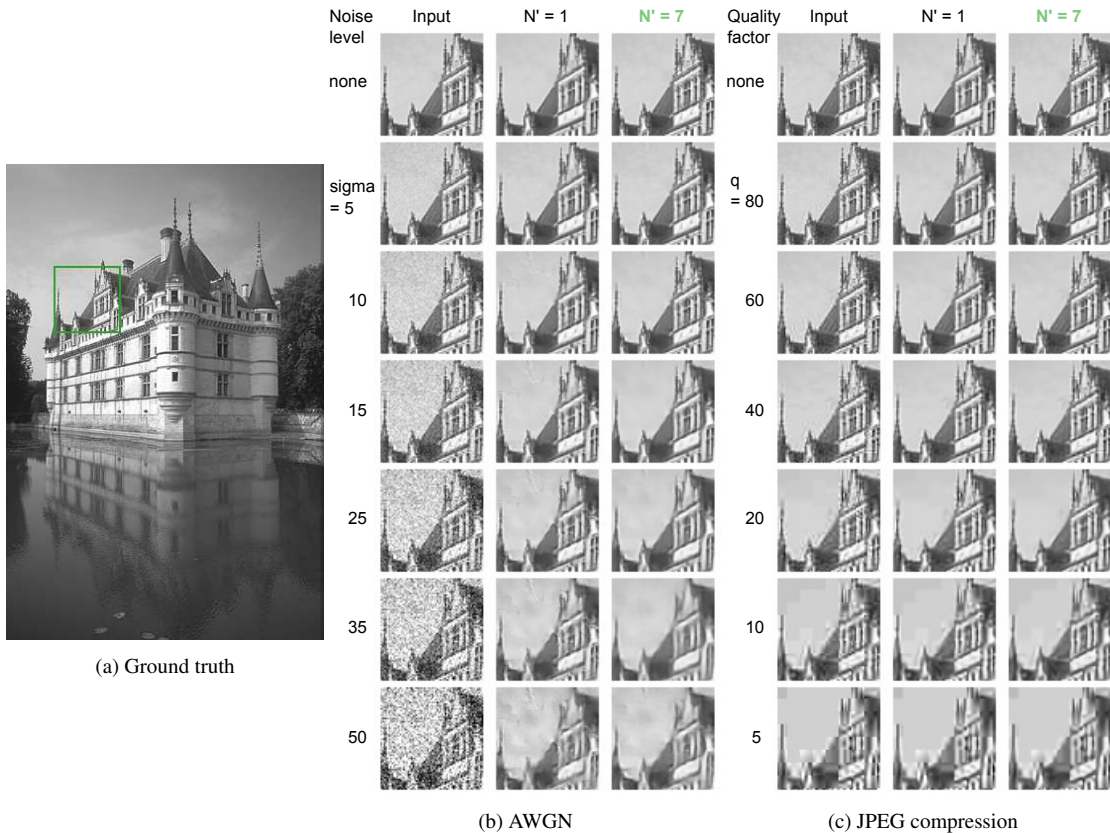


Figure 7: Visual comparisons between with and without competition. The employed architecture is DnCNN-woBN-*d5c16*, and the number of experts is seven for the competition. (a) A ground-truth image from BSD68. (b) From top to bottom: inputs and denoising results with and without competition for different AWGN levels 0 (without noise), 5, 10, 15, 25, 35, and 50, respectively. (c) From top to bottom: inputs and denoising results with and without competition for different JPEG compression qualities null (without compression), 80, 60, 40, 20, 10, and 5, respectively.

Table 1 (Top) summarizes the average PSNR and execution time for BSD68 synthesized with specific level AWGNs. Our DnCNN-woBN-*d9c32* shows the highest PSNR scores for all of the noise levels, and its execution time is approximately five times faster than that of BM3D. Furthermore, our DnCNN-woBN-*d5c16* is more than four times faster than fast-mode NL-means while demonstrating overperformance in PSNR by approximately 0.8dB. Table 1 (Bottom) shows scores reported in some more non-blind methods. Though they show higher PSNR than our *d9c32* network, they are one to three orders of magnitude larger in computations. Moreover, they evaluated in the non-blind condition which is much easier than our blind condition.

Table 2 summarizes the average PSNR, structural similarity (SSIM) [32], and execution time for BSD68 compressed by specific JPEG qualities. Both of the DnCNN-woBN-*d5c8* and *d5c16* indicate higher performance than the model of Knusperli, and the execution of DnCNN-woBN-*d5c8* is 37% faster than Knusperli. Visual comparisons between our method and prior arts for AWGNs and JPEG compression noises are exhibited in Fig. 8 and Fig. 9.

## 6. Conclusion

In this paper, we proposed an efficient ensemble network trained via a competition of expert networks, aiming application for low-level vision tasks. We apply the proposed method to image blind denoising and realize automatic division of unlabeled noisy datasets into clusters respectively optimized to enhance denoising performance. This method enables to save up to approximately 90% of computational cost without sacrifice of the denoising performance compared to single network models with identical architectures. We also compare the proposed method with several existing algorithms and observe significant outperformance over the standard denoising algorithm BM3D in terms of computational efficiency.

The training procedure of our method relies on an assumption that any region in an input image belongs to the same domain in the dataset. Consequently, it is difficult to deal with spatially variant low-frequency features of images. A further extension of our method would be the adaptation for other low-level vision tasks.

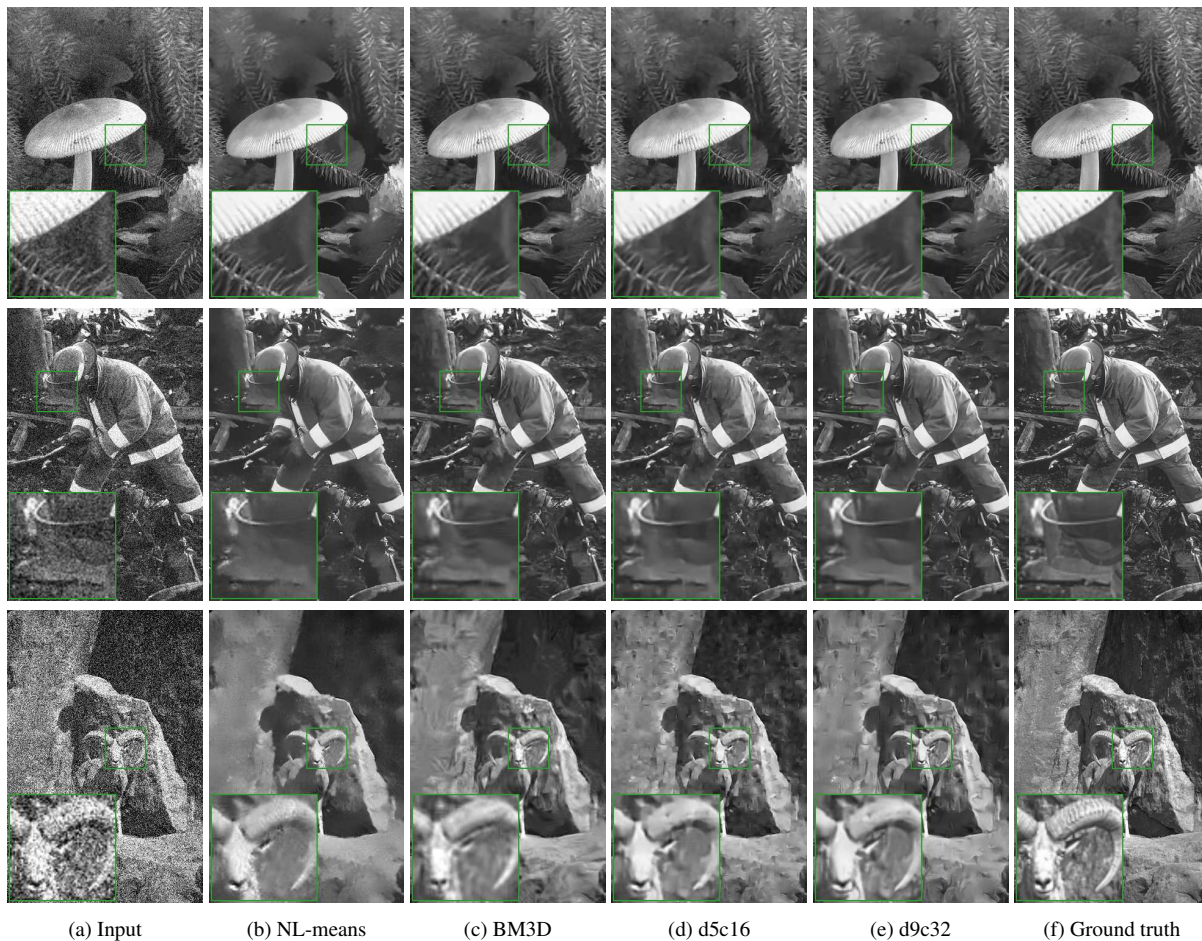


Figure 8: Visual comparisons between our method and prior arts NL-means and BM3D for three test images synthesized with specific level AWGNs. From top to bottom:  $\sigma = 15, 25, 50$ . The number of experts is seven for our methods DnCNN-woBN-*d5c16* and *d9c32*.

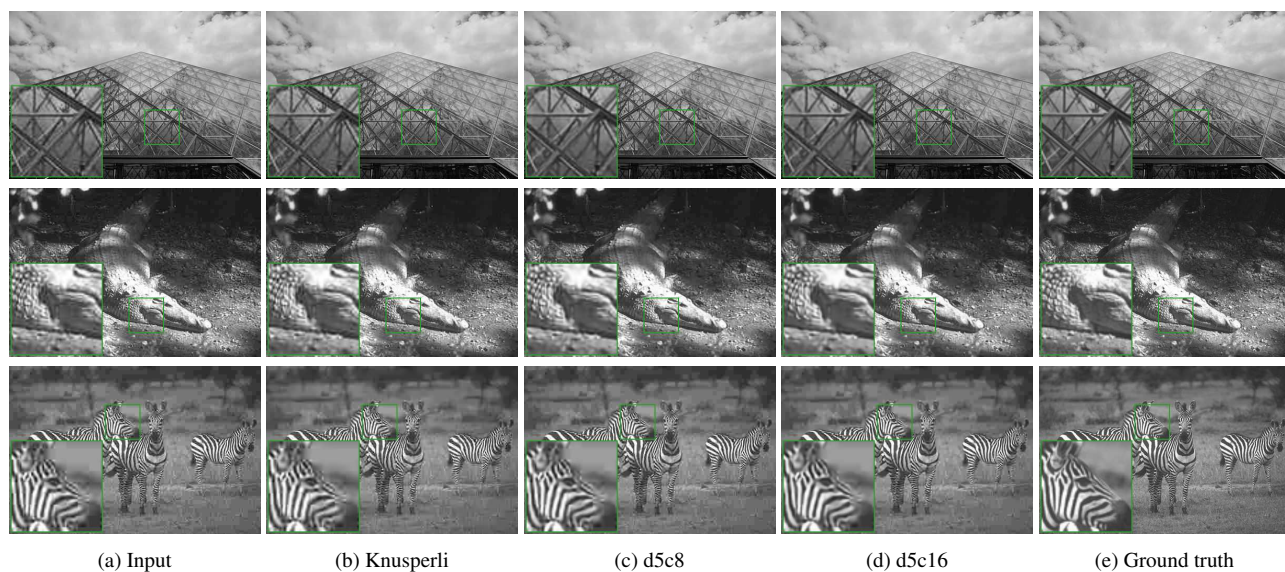


Figure 9: Visual comparisons between our method and a prior art Knusperli for three test images compressed by specific JPEG qualities. From top to bottom:  $q = 40, 20, 10$ . The number of experts is seven for our methods DnCNN-woBN-*d5c8* and *d5c16*.



## References

- [1] A. Abdelhamed, S. Lin, and M. S. Brown. A high-quality denoising dataset for smartphone cameras. In *CVPR*, 2018. 1, 2, 6
- [2] A. Beck and M. Teboulle. Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. *IEEE Transactions on Image Processing*, 2009. 1
- [3] Y. Bengio, N. Léonard, and A. Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 2
- [4] A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. In *CVPR*, 2005. 1, 2
- [5] H. C. Burger, C. J. Schuler, and S. Harmeling. Image denoising: Can plain neural networks compete with bm3d? In *CVPR*, 2012. 6
- [6] J. Chen, J. Chen, H. Chao, and M. Yang. Image blind denoising with generative adversarial network based noise modeling. In *CVPR*, 2018. 1
- [7] Y. Chen and T. Pock. Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration. *IEEE transactions on pattern analysis and machine intelligence*, 2016. 6
- [8] K. Dabov, A. Foi, V. Katkovich, and K. Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on image processing*, 2007. 1, 2
- [9] S. Diamond, V. Sitzmann, S. Boyd, G. Wetzstein, and F. Heide. Dirty pixels: Optimizing image classification architectures for raw sensor data. *arXiv preprint arXiv:1701.06487*, 2017. 1
- [10] R. M. French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 1999. 3
- [11] X. Glorot, A. Bordes, and Y. Bengio. Deep sparse rectifier neural networks. In *AISTATS*, 2011. 4
- [12] Google. Knusperli: A deblocking jpeg decoder. <https://github.com/google/knusperli>. Accessed: 2019-03-18. 6
- [13] S. Gu, L. Zhang, W. Zuo, and X. Feng. Weighted nuclear norm minimization with application to image denoising. In *CVPR*, 2014. 1, 2, 6
- [14] J. Guo, D. J. Shah, and R. Barzilay. Multi-source domain adaptation with mixture of experts. *arXiv preprint arXiv:1809.02256*, 2018. 2
- [15] S. Guo, Z. Yan, K. Zhang, W. Zuo, and L. Zhang. Toward convolutional blind denoising of real photographs. *arXiv preprint arXiv:1807.04686*, 2018. 1, 2
- [16] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, 2015. 4
- [17] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 4
- [18] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, G. E. Hinton, et al. Adaptive mixtures of local experts. *Neural computation*, 1991. 1, 2
- [19] M. I. Jordan and R. A. Jacobs. Hierarchical mixtures of experts and the em algorithm. *Neural computation*, 1994. 1, 2
- [20] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 4
- [21] M. Lebrun. An analysis and implementation of the bm3d image denoising method. *Image Processing On Line*, 2012. 6
- [22] S. Lee, S. Purushwalkam, M. Cogswell, V. Ranjan, D. Crandall, and D. Batra. Stochastic multiple choice learning for training diverse deep ensembles. In *NIPS*, 2016. 2
- [23] M. Lin, Q. Chen, and S. Yan. Network in network. In *ICLR*, 2013. 4
- [24] K. Ma, Z. Duanmu, Q. Wu, Z. Wang, H. Yong, H. Li, and L. Zhang. Waterloo exploration database: New challenges for image quality assessment models. *IEEE Transactions on Image Processing*, 2017. 4
- [25] 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 *ICCV*, 2001. 4
- [26] G. Parascandolo, N. Kilbertus, M. Rojas-Carulla, and B. Schölkopf. Learning independent causal mechanisms. In *ICML*, 2018. 2
- [27] T. Plotz and S. Roth. Benchmarking denoising algorithms with real photographs. In *CVPR*, 2017. 1
- [28] S. Roth and M. J. Black. Fields of experts: A framework for learning image priors. In *CVPR*, 2005. 4
- [29] Scikit-image. Non-local means denoising for preserving textures. [https://github.com/scikit-image/scikit-image/blob/master/skimage/restoration/non\\_local\\_means.py](https://github.com/scikit-image/scikit-image/blob/master/skimage/restoration/non_local_means.py). Accessed: 2019-03-18. 6
- [30] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, and J. Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *ICLR*, 2017. 2
- [31] Y. Tang and C. Eliasmith. Deep networks for robust visual recognition. In *ICML*, 2010. 1
- [32] Z. Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli, et al. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 2004. 7
- [33] 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*, 2017. 1, 2, 4, 6
- [34] K. Zhang, W. Zuo, and L. Zhang. FFDNet: Toward a fast and flexible solution for cnn-based image denoising. *IEEE Transactions on Image Processing*, 2018. 1, 2
- [35] D. Zoran and Y. Weiss. From learning models of natural image patches to whole image restoration. In *ICCV*, 2011. 6