

NoiseTransfer: Image Noise Generation with Contrastive Embeddings

Seunghwan Lee and Tae Hyun Kim

Dept. of Computer Science, Hanyang University, Seoul, Korea
{seunghwanlee, taehyunkim}@hanyang.ac.kr

Abstract. Deep image denoising networks have achieved impressive success with the help of a considerably large number of synthetic train datasets. However, real-world denoising is a still challenging problem due to the dissimilarity between distributions of real and synthetic noisy datasets. Although several real-world noisy datasets have been presented, the number of train datasets (*i.e.*, pairs of clean and real noisy images) is limited, and acquiring more real noise datasets is laborious and expensive. To mitigate this problem, numerous attempts to simulate real noise models using generative models have been studied. Nevertheless, previous works had to train multiple networks to handle multiple different noise distributions. By contrast, we propose a new generative model that can synthesize noisy images with multiple different noise distributions. Specifically, we adopt recent contrastive learning to learn distinguishable latent features of the noise. Moreover, our model can generate new noisy images by transferring the noise characteristics solely from a single reference noisy image. We demonstrate the accuracy and the effectiveness of our noise model for both known and unknown noise removal.

Keywords: Image denoising · Image noise generation.

1 Introduction

Noise is a common artifact in imaging systems; thus, accurately modeling the noise is an important task for many image-processing and computer vision applications. To remove the noise in an image, several statistical noise models have been adopted in the literature. The most simple, widely used noise models are additive white Gaussian noise and Poisson noise. However, in a real-world scenario, image noise does not follow Gaussian and Poisson distributions [31, 3], and these simple statistical noise models cannot accurately capture the noise characteristic of real noise which includes signal-dependent and signal-independent components. Moreover, developing a noise model that can simulate complex real-world noise process is very difficult because complicated processing steps of an imaging pipeline include various noise models, such as photon noise, read

Code is available at <https://github.com/shlee0/NoiseTransfer>

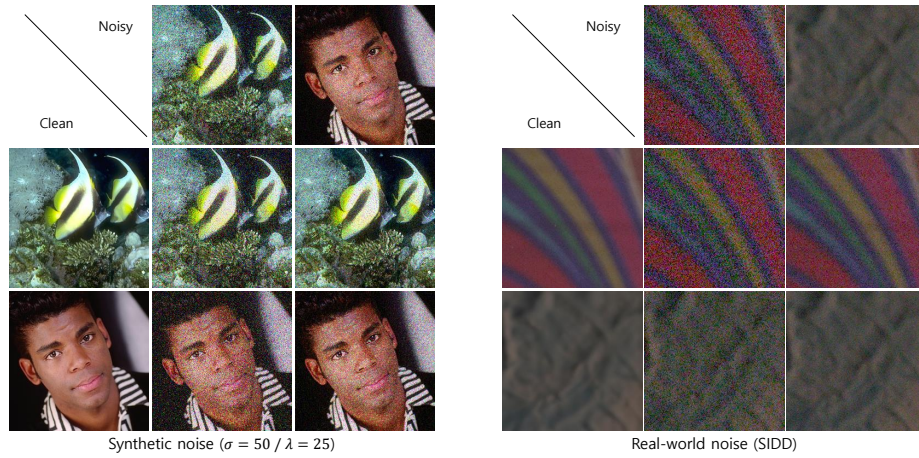


Fig. 1: Examples of generated noisy images. Our proposed model can generate a noisy version of a clean image by transferring noise information in the reference noisy image. **(Left)** Synthetic noise (*i.e.*, Gaussian noise ($\sigma = 50$) (2nd column) and Poisson noise ($\lambda = 25$) (3rd column)) generation results. **(Right)** Real-world noise generation results from the Smartphone Image Denoising Dataset (SIDD) [3]. Noisy images generated in an unpaired manner can have different noise levels that do not exist in the original dataset.

noise, and spatially correlated noise. Conventional denoising networks that show promising results in removing noise from known distribution (*e.g.*, Gaussian) frequently fail in dealing with real noise from an unknown distribution due to these limitations.

Collecting real-world datasets that include pairs of clean and real-noisy images can solve these problems. However, the noise distributions of conventional cameras are different from one another, so we need to acquire a large amount of labelled real-world dataset, which is very time-consuming. This problem stimulates the need for synthetic, but realistic noise generation system to avoid taking pairs of clean and noisy pictures. Recently, several generative adversarial network (GAN)-based noise models have been proposed to model the complex real-world noise in a data-driven manner better. Since Chen *et al.* [12] has proposed a generative model to synthesize zero-mean noise, recent models [24, 2, 19, 11, 36, 9, 20, 10] made many attempts to generate signal-dependent noise by considering a clean image as a conditional input.

Despite this encouraging progress, there are still some steps to move forward for image noise generation. Typically, generative models have difficulty in controlling the specific type of noise during synthesizing. In other words, which types of noise will be realized is not predictable at the inference if generator is trained with a large of noise distributions. In addition, this randomness increases if the training dataset includes several different noise types. A naïve, straightforward

solution would be to train multiple generators independently to handle multiple noise models. Alternatively, image metadata such as camera-ISO and the raw Bayer pattern can be utilized to avoid this hassle. However, this external data is not always available (*e.g.*, images from unknown resources).

In this work, we propose a novel generative noise model, which can allow multiple different types of noise models. We transfer the noise characteristics within a given reference noisy image to corrupt freely available clean images, and we synthesize new noisy images in this manner. Moreover, our model requires only the noisy image itself without demanding any external information (*e.g.*, metadata). Specifically, we train our discriminator to distinguish the distribution of each noise from the others in a self-supervised manner by adopting a contrastive learning. Then, our generator learns to synthesize a new noisy image using the noise information extracted from the discriminator. With this strategy, we can perform noise generation with paired or unpaired images, and Fig. 1 presents some examples. We demonstrate that our generative noise model can handle a wide range of noise distributions, and the conventional denoising networks trained with our newly synthesized noisy images can remove the real noise much better than existing generative noise models. The main contributions of our work are summarized as follows:

- We propose a novel generative noise model that can handle diverse noise distributions with a single noise generator without additional meta information.
- Our model exploits the representation power of the contrastive learning. To the best of our knowledge, our model is the first approach which utilizes contrastive noise embedding to control the type of noise to be generated.
- Extensive experiments demonstrate that our model achieves state-of-the-art performance in noise generation and is applicable for image denoising.

2 Related Work

2.1 Contrastive learning

The contrastive learning mechanism introduced by [16] learns similar/dissimilar representations in a self-supervised manner from positive/negative pairs. In the works of instance discrimination [8, 35], a query and a key form a positive pair if they originate from the same image and form a negative pair if otherwise. It is known that more negative samples can yield better representation ability, and a large number of negative samples can be maintained in a batch [13] or dynamic dictionary updated by a momentum-based key encoder [18].

After the contrastive learning has shown powerful representation ability in several downstream tasks, it has been integrated with the GAN framework as an auxiliary task. For instance, contrastive learning could relieve forgetting problem of discriminator [14, 26], and improve image translation quality by maximizing mutual information of corresponding patches in different domains [30, 17]. ContraGAN [22] improved image generation quality by incorporating data-to-data

relations as well as data-to-class relations into discriminator. ContraD [21] empirically showed that training the GAN discriminator jointly with the augmentation techniques used in the literature of contrastive learning benefits the task of the discriminator. Moreover, contrastive learning can learn content-invariant degradation representation by constructing image pairs with the same degradation as positive examples. Recently, DASR [33] and AirNet [27] utilized learned degradation representation for image restoration. Different from previous works, our work studies image noise synthesis conditional on degradation representation learned through contrastive learning.

2.2 Generative noise model

To address the limitations of simple synthetic noise models, considerable effort has been devoted to numerous generative noise models to synthesize complex noise for the real-world image denoising problem. Particularly, recent generative noise models yield signal-dependent noise given a clean image. Some approaches require metadata (*e.g.*, smartphone code, ISO level, and shutter speed) as an additional input to generate noise from a specific distribution [24, 2, 11]. However, these approaches assume that the metadata is available, which might not be common in the real scenario (*e.g.*, internet images and pictures), and the use of this additional information limits the usage of the generative noise model in practice. Unlike existing generative models, our model extracts noise representation from an input noisy image itself without relying on the metadata, and thus allows us to use any noisy image as a reference. Then, our generator synthesizes new noisy images based on noise information of the reference noisy image, such that we can easily predict which type of noise will be realized.

3 Proposed Method: NoiseTransfer

Our generative noise model synthesizes new noisy images by transferring the noise of a reference noisy image to other clean images. Specifically, our discriminator takes a single reference noisy image as an input and outputs noise embeddings that represent noise characteristics of the reference noisy image. Then, our generator synthesizes new noisy images by corrupting clean images available for free using the given noise embeddings that we dub NoiseTransfer. Fig. 2 depicts the overview of the proposed NoiseTransfer scheme.

3.1 Noise discrimination with contrastive learning

Capturing different characteristics for different noises is essential to keep the noise information distinct. Therefore, we train our discriminator through contrastive learning to learn distinguishable noise embeddings of each noise, and we follow MoCo [18] framework: dynamic dictionary holding a large number of

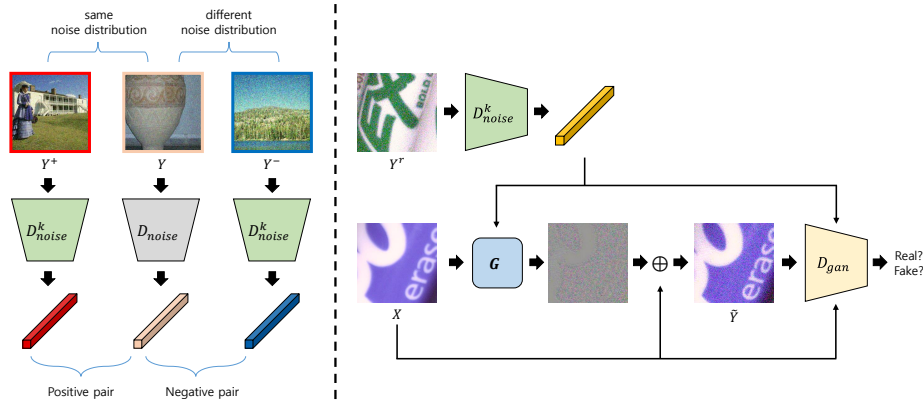


Fig. 2: Our discriminator consists of two branches with shared intermediate convolutional modules for each forward operation denoted by D_{noise} and D_{gan} respectively. **(Left)** An illustration of our noise representation learning scheme. Two noisy images sampled from the same noise distribution form a positive and, in different cases, a negative. **(Right)** Overall flow of the proposed NoiseTransfer. Our noise generator takes a clean image X and noise embeddings $D_{noise}^k(Y^r)$ where Y^r is a reference noisy image.

negative samples and momentum-based key network. Then, a form of a contrastive loss function, called InfoNCE [29], can be written with cosine similarity $s(u, v) = u \cdot v / \|u\|_2 \|v\|_2$ for encoded embeddings u and v as follows:

$$L_{Con}(q, k^+, Q) = -\log \frac{\exp(s(q, k^+)/\tau)}{\exp(s(q, k^+)/\tau) + \sum_{k^- \in Q} \exp(s(q, k^-)/\tau)}, \quad (1)$$

where q , k^+ , and k^- denote the embeddings of a query, positive key, and negative key, respectively; Q denotes a queue containing negative keys; and τ is a temperature hyperparameter. Eq. 1 pulls embeddings of the q close to those of the k^+ and pushes them apart from those of the k^- .

In our work, as shown in Fig. 2 (Left), we construct a positive pair of noisy images if they are sampled from the same noise distribution and a negative pair, otherwise. Then, the contrastive loss for noise discrimination can be formulated as follows:

$$L_{noise}^D = \mathbb{E}[L_{Con}(D_{noise}(Y), D_{noise}^k(Y^+), Q)], \quad (2)$$

where Y^+ denotes a noisy image that has the same noise distribution to that of another noisy image Y . Note that we encode the keys (k^+ and k^-) with momentum-based key network D_{noise}^k . We assume that embeddings in Q are from noisy images whose noise distributions are different from that of Y . Eq. 2 encourages our discriminator to learn distinguishable noise representation for each different noise.

Our final goal is to synthesize a new noisy image \tilde{Y} through a generator, which has the same noise distribution as the real one Y . Thus, we derive another contrastive loss for the generator as follows:

$$L_{noise}^G = \mathbb{E}[L_{Con}(D_{noise}(\tilde{Y}), D_{noise}^k(Y^+), Q)]. \quad (3)$$

Note that, in Eq. 3, our generated noisy image \tilde{Y} is encoded as a query. Moreover, we adopt a feature matching loss [32] to stabilize training as follows:

$$L_{noise}^{FM} = \|m_{noise}(Y) - m_{noise}(\tilde{Y})\|_1, \quad (4)$$

where $m_{noise}(\cdot)$ denotes the intermediate feature maps before pooling operation in the D_{noise} (please refer to the supplement for details).

3.2 Noise generation with contrastive embeddings

Given a clean image X and a reference noisy image Y^r , our generator learns to synthesize a new noisy image \tilde{Y} which is a noisy version of X and has the same noise distribution as Y^r . The generation process is described in Fig. 2 (Right). Reference noisy image Y^r is encoded by D_{noise}^k , and noise embeddings $D_{noise}^k(Y^r)$ that contain noise representation of the Y^r are fed to our generator. This approach enables our model to handle a wide range of noise distributions with a single generator. To generate realistic noisy images, our model performs adversarial training. The adversarial losses [15] for our model are defined as follows:

$$\begin{aligned} L_{gan}^D &= -\mathbb{E}[\log(D_{gan}(R))] - \mathbb{E}[\log(1 - D_{gan}(F))] \\ L_{gan}^G &= -\mathbb{E}[\log(D_{gan}(F))], \end{aligned} \quad (5)$$

where R denotes a set of $X, D_{noise}^k(Y^r)$, and Y , whereas F includes \tilde{Y} instead of Y . Our generator synthesizes noisy images with different kinds of noise distribution based on the $D_{noise}^k(Y^r)$, even with the same clean image X . Thus, our discriminator distinguishes whether the input noisy image is real or fake considering the X and $D_{noise}^k(Y^r)$. Similar to Eq. 4, we adopt feature matching loss for stable adversarial training as follows:

$$L_{gan}^{FM} = \|m_{gan}(Y) - m_{gan}(\tilde{Y})\|_1, \quad (6)$$

where $m_{gan}(\cdot)$ denotes feature maps before the last convolution layer in the D_{gan} . Finally, we utilize L_1 reconstruction loss $L_{recon} = \|GF(Y) - GF(\tilde{Y})\|_1$ with the Gaussian filter GF as used in [36] to enforce statistical features of noise distribution. Then, we define the final objective functions for our model as follows:

$$\begin{aligned} L_D &= L_{noise}^D + L_{gan}^D \\ L_G &= L_{noise}^G + L_{gan}^G + \\ &\quad \lambda_{noise}^{FM} L_{noise}^{FM} + \lambda_{gan}^{FM} L_{gan}^{FM} + \lambda_{recon} L_{recon}, \end{aligned} \quad (7)$$

where $\lambda_{noise}^{FM}, \lambda_{gan}^{FM}$, and λ_{recon} control the weights of the associated terms.




Fig. 3: Visual results of noise generation on the SIDD validation set. The corresponding noise is displayed below for each noisy image. **Left to Right:** CA-NoiseGAN [11], DANet [36], GDANet [36], NoiseGAN [9], C2N [20], NoiseTransfer (Ours), Noisy, Clean.

3.3 Discussion

Our model has several advantages compared with existing noise generators. [9] trained 17 different generators to handle numerous camera models and ISO levels. This solution could be straightforward to cover various noise distributions, but training multiple generators for each different noise lacks practicality. To compensate this, image metadata of a noisy image can be exploited to sample a specific noise type [24, 2, 11]. However, such external information is not always available in the real-world. PNGAN [10] requires pre-trained networks for training. Specifically, it uses a camera pipeline modeling network [37] to generate a noisy image that is further refined by generator. It also employs a pre-trained




Fig. 4: Visual results for real noise removal on the SIDD validation set (first three rows) and SIDD+ set (last four rows). **Left to Right:** RIDNet results trained by DANet [36], GDANet [36], C2N [20], CycleISP [37], NoiseTransfer (Ours), and Noisy and Clean image.

denoising network as the regularizer. This strategy makes the generated noisy image distribution dependent on pre-trained networks, which may not be suitable in several cases. Although C2N [20] takes random vector that determines the property of synthesized noise, we do not know which value should be used for the random vector when a particular type of noise is required. Compared with these models, our model can handle numerous noise distributions with a single generator, and does not need external resources, and synthesizes a desired noise by transferring the noise information from a reference noisy image.

4 Experiments

4.1 Implementation Details

We train our NoiseTransfer model by using various synthetic and real noisy images. First, for real-world noise, we use the SIDD-Medium dataset [3] following previous works [2, 11, 36, 20]. In this case, two different patches are randomly selected from the same noisy image to get Y and Y^r . For synthetic noise, we sample noise from Gaussian distribution ($\sigma \in [0, 70]$), Poisson distribution ($\lambda \in [5, 100]$), and the combined Poisson-Gaussian distribution ($\sigma \in [0, 70]$ and $\lambda \in [5, 100]$). Then, we acquire synthetic noisy images by corrupting clean images in the DIV2K training set [4] and SIDD-Medium set using the noise from these synthetic distributions. We use 32 mini-batch of 96×96 patches for training. Each mini-batch includes 16 patches from the SIDD-Medium set and 16 patches corrupted with synthetic noise distributions. We apply data augmentation (flip and rotation) to diversify train images. The noise embedding vector (*i.e.*, outcome by D_{noise}) has 128-dimension, the size of the queue is set to 4096, and temperature parameter τ is set to 0.1 [23]. λ_{noise}^{FM} , λ_{gan}^{FM} , and λ_{recon} are equally set to 100. We use Adam optimizer [25] with an learning rate of 1e-4, $\beta_1 = 0.5$, and $\beta_2 = 0.99$. We also apply L2 regularization with regularization factor 1e-7. Our discriminator and generator are updated 2,000 times during one epoch, and training for 200 epochs takes approximately a week on two Tesla V100 GPUs. We provide more details including network configurations and additional experimental results in the supplement.

4.2 Noisy Image Generation

We first measure the accuracy of the generated noisy images. To do so, we use Average KL Divergence (AKLD) value [36] and Kolmogorov-Smirnov (KS) test value¹ [9] for quantitative evaluation, and compare the results with DANet [36], GDANet [36], C2N [20], and CycleISP [37] in Tab. 1. Note that, DANet trained only with the SIDD-Medium dataset outperforms GDANet [36] trained with three different real-noise datasets (SIDD-Medium, Poly [34], and RENOIR [5]). The results demonstrate that GDANet does not handle the specific noise better in the SIDD dataset than DANet. CycleISP [37] samples random noise considering specific camera settings; hence, it is unlikely that distribution of the randomly sampled noise matches to that of noise within a specific noisy image. By contrast, our NoiseTransfer which is trained with multiple different noise models can deal with the specific noise by transferring the noise characteristics within the reference noisy image, because our model utilizes noisy images as the reference \tilde{Y} as well as clean images. This advantage allows our model to obtain the best performance among the compared models. However, it is worth mentioning that better AKLD/KS values do not always imply higher denoising

¹ Histograms are computed with 256 bins evenly distributed in $[-256, 256]$.

performance as will be described in Sec. 4.3. AKLD/KS values compute the distance of pixel value distributions of two images, thus, we cannot predict how realistic generated noisy image is with only those values.

Fig. 3 presents visual comparisons of generated noisy images. The visual results show that our NoiseTransfer can synthesize more realistic and desirable noise than other models which frequently generate unexpected patterns. Note that CA-NoiseGAN [11] conducts noise generation with raw images², and NoiseGAN [9] only covers four ISO levels (400 \sim 3200), thus, we provide only visual comparisons with these approaches.

Table 1: AKLD/KS test values on the SIDD validation and SIDD+ datasets. The best values are highlighted in bold.

Dataset	SIDD validation		SIDD+	
	AKLD↓	KS↓	AKLD↓	KS↓
DANet	0.2117	0.0732	0.4144	0.1468
GDANet	0.2431	0.1079	0.4744	0.2557
C2N	0.3138	0.1499	0.2882	0.1330
CycleISP	0.6881	0.1743	0.9412	0.1743
Ours	0.1655	0.0617	0.2324	0.1537

4.3 Real Noise Denoising

To more accurately validate the quality of generated noisy images, we evaluate the applicability of our NoiseTransfer in real-world image denoising. In this work, we choose lightweight yet effective RIDNet [6] as a baseline denoising network. For a fair comparison, all generative noise models are evaluated by measuring the denoising performance of RIDNet. Following previous works [2, 11, 36, 20], we use images from SIDD [3] for training and validation. The ground-truth clean images and the corresponding generated noisy images are used to train RIDNet. We do not include the ground-truth noisy images in the dataset when training the denoiser to evaluate generative noise models. For our NoiseTransfer, we randomly select a clean image X and choose another random noisy image as the reference Y^r from the SIDD-Medium dataset to render a new noisy patch \tilde{Y} .

We evaluate the real noise removal performance on the SIDD validation, SIDD+ [1], and DND benchmark [31] datasets. In Tab. 2, we measure the denoising performance in terms of PSNR and SSIM values, and denoising results by RIDNet trained with generated noisy images from DANet, GDANet, C2N, CycleISP, and our NoiseTransfer are compared. We also present the result when the ground-truth noisy images are used instead of generated images (RIDNet+GT). Note that C2N [20] got the best KS value on SIDD+ in Tab. 1, but PSNR/SSIM

² For visualization, the camera pipeline matlab code (<https://github.com/AbdoKamel/simple-camera-pipeline>) is used.

values are lower than other models. This result shows AKLD/KS values do not always hint at higher denoising performance as stated in Sec. 4.2. Our outstanding denoising results show that RIDNet trained with generated noisy images by our NoiseTransfer is generally applicable for real-world denoising. Notably, our method got comparable denoising performance with ‘RIDNet+GT’, especially on SIDD+, and this result is not surprising in this field. For example, NoiseFlow [2] achieved better performance when using generated images during training rather than the ground-truth real images for raw image denoising (refer to Table. 3 in [2]). This is due to the small number of real GT samples in the training dataset. Fig. 4 shows visual denoising results on the SIDD validation and SIDD+ datasets.

Moreover, we plot changes of PSNR values by RIDNet during training on the SIDD validation and SIDD+ datasets in Fig. 5. Particularly, we observe that when the RIDNet is trained with noisy images generated by either GDANet or DANet, the denoiser is overfitted after some iterations and PSNR values drops. We believe this overfitting problem can be caused by unrealistic patterns that GDANet and DANet produce as shown in Fig. 3. Note that CycleISP is not a generative model and instead injects synthetic realistic noise, so RIDNet trained with noisy images synthesized by CycleISP does not suffer from the overfitting problem. However it provides limited performance because CycleISP considers predetermined shot/read noise factors for specific camera settings to inject random noise, which may not follow distribution of real noise. In contrast, RIDNet trained with images by our NoiseTransfer results show promising denoising results on several datasets (more than 0.5dB compared with CycleISP on average).

Table 2: Denoising results in terms of PSNR/SSIM values on the various real-world noise datasets. ‘GT’ denotes the ground-truth noisy images. **Red** and **Blue** denote the best and second values.

Dataset	SIDD validation	SIDD+	DND
RIDNet + DANet	35.91 / 0.8762	34.64 / 0.8600	37.54 / 0.9292
RIDNet + GDANet	36.58 / 0.8851	35.22 / 0.8844	37.13 / 0.9345
RIDNet + C2N	31.87 / 0.7274	34.22 / 0.8290	35.41 / 0.9032
RIDNet + CycleISP	38.09 / 0.9095	35.49 / 0.9006	38.95 / 0.9485
RIDNet + NoiseTransfer	38.57 / 0.9112	36.30 / 0.9095	39.15 / 0.9492
RIDNet + GT	39.14 / 0.9155	36.21 / 0.9057	39.35 / 0.9502

4.4 Synthetic Noise Denoising

Finally, we evaluate the applicability of our NoiseTransfer in removing synthetic noise. To do so, we first generate noisy images which include noise from known distributions using our NoiseTransfer. Specifically, we use randomly selected clean image in the DIV2K training set as X , and add synthetic noise into another clean image for the reference noisy image Y^r . We add one of Gaussian noise ($\sigma \in [0, 50]$) and Poisson noise ($\lambda \in [5, 50]$) following N2G [28]. Additionally, we also add noise from Poisson-Gaussian distribution ($\sigma \in [0, 50], \lambda \in [5, 50]$) to




Fig. 5: PSNR value changes during training on the SIDD validation set (Left) and SIDD+ (Right). Denoising performance trained by GT noisy images, Noise-Transfer (Ours), CycleISP, GDANet, DANet, and C2N are compared.

Table 3: PSNR/SSIM results of synthetic noise removal on the BSDS500 dataset. ‘GT’ denotes the ground-truth noisy images. ‘N2G_g’ and ‘N2G_p’ denote two independently trained networks for Gaussian and Poisson noise respectively. For a random noise level, we report an average of 10 trials. Red and Blue denote the best and second values respectively.

Noise Type	Gaussian		Poisson		Poisson-Gaussian	
Noise level	Fixed ($\sigma = 25$)	Random ($\sigma \in [0, 50]$)	Fixed ($\lambda = 30$)	Random ($\lambda \in [5, 50]$)	Fixed ($\lambda = 30, \sigma = 25$)	Random ($\lambda \in [5, 50], \sigma \in [0, 50]$)
RIDNet + N2G _g	31.21 / 0.8806	32.82 / 0.8818	30.28 / 0.8702	29.52 / 0.8464	28.80 / 0.8276	28.15 / 0.8070
RIDNet + N2G _p	31.17 / 0.8792	32.74 / 0.8825	30.31 / 0.8703	29.56 / 0.8477	28.93 / 0.8301	28.32 / 0.8106
RIDNet + NoiseTransfer	31.27 / 0.8834	32.76 / 0.8872	30.58 / 0.8735	29.95 / 0.8530	29.16 / 0.8344	28.71 / 0.8190
RIDNet + GT	31.48 / 0.8868	33.20 / 0.8913	30.74 / 0.8774	30.20 / 0.8608	29.31 / 0.8375	28.85 / 0.8212

confirm that our NoiseTransfer can generate diverse noises well. The new noisy image \tilde{Y} is synthesized with X and Y^r , and RIDNet is trained with pairs of clean image X and generated noisy image \tilde{Y} . In Fig. 6, we present examples of our generated noisy images from known distributions, and we see that our model can synthesize signal-independent noise as well as signal-dependent one.

To quantitatively measure the denoising performance, we use BSDS500 dataset [7] as testset and degrade images in the BSDS500 by adding noise from known distributions, and then put them into RIDNet as input. Denoising performance of RIDNet trained with noisy images by N2G [28], our NoiseTransfer, and the ground-truth noisy images are compared in Tab. 3. Note that N2G does not generate noise, but instead extracts noise by denoising the input noisy image. It also adopts a bernoulli random mask to destroy residual structure in the noise, and then the masked noise is used to corrupt other clean images. In Tab. 3, N2G exhibits favorable performance when the noise type in N2G training matches the noise type in the test noisy image, but it reveals a slight performance drop against other types of noise. In other words, N2G_g shows slightly better denoising results for Gaussian noise removal and N2G_p for Poisson noise removal. This result implies that we may need to train multiple networks separately for each noise distribution (*e.g.*, separated two networks for Gaussian and Poisson). By

contrast, our NoiseTransfer shows consistently better denoising performance for several types of noise with a single generator, thus, our method does not require multiple generators independently trained for each noise type.

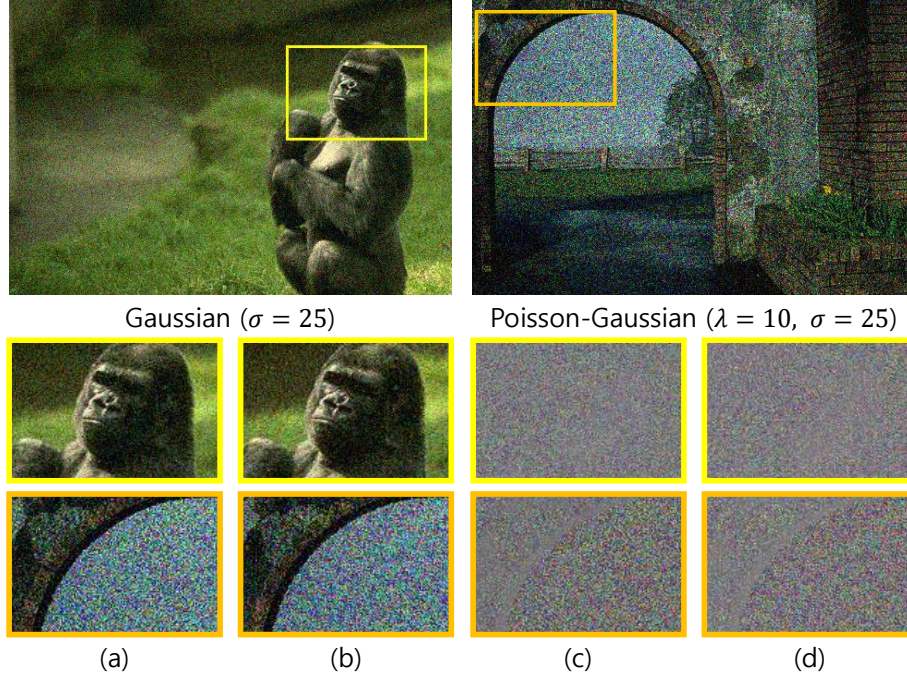


Fig. 6: Examples of synthetic noise generation. (a) Ground-truth noisy image. (b) Generated noisy image by NoiseTransfer (Ours). (c)-(d) Noise added in (a) and (b), respectively.

4.5 Ablation Study

In this work, we introduced contrastive losses for our NoiseTransfer. Thus, we provide ablation study with and without using the additional contrastive losses during training. We compare the accuracy of the generated noisy images in terms of AKLD and KS value. Fig. 7 obviously shows the effect of contrastive learning for image noise generation of our model³. First, without L_{noise}^D which is the crux of our approach, we found the noise generation performance is very poor, and the model diverged after 14 epochs (Green). This result demonstrates that learning distinguishable noise representation is crucial for our single generator

³ It compares the results for the first 60 training epochs, but is sufficient to confirm the effect of the proposed losses, allowing fair and efficient ablations.

to cover different kinds of noise distributions. Next, with L_{noise}^D , we could see much better training results, but still have unstable performance early in training (Blue). Finally, we can achieve more training stability and better performance when we explicitly guide our generator to synthesize a new noisy image with the same noise distribution to that of the Y^r with additional losses L_{noise}^G and L_{noise}^{FM} (Red).




Fig. 7: AKLD and KS test values for the first 60 training epochs. (Green) Trained without L_{noise}^D . This model diverged after 14 epochs. (Blue) Trained without L_{noise}^G and L_{noise}^{FM} . (Red) Our final NoiseTransfer.

5 Conclusion

In this work, we proposed a novel noisy image generator trained with contrastive learning. Different from existing works, our discriminator learns distinguishable noise representation for each different noise, which is the core of our method. Thus, ours can extract noise characteristics from an input reference noisy image and generate new noisy images by transferring the specific noise to clean images. This approach enables our generator to synthesize noisy images based on the noise information both in a paired or unpaired manner. Consequently, our model can handle multiple noise distributions with a single generator. Experiments demonstrate that the proposed generative noise model can produce more accurate noisy images than conventional methods and the applicability for image denoising.

Acknowledgements This work was supported by Samsung Electronics Co., Ltd, and Samsung Research Funding Center of Samsung Electronics under Project Number SRFCIT1901-06.

References

1. Abdelhamed, A., Afifi, M., Timofte, R., Brown, M.S.: Ntire 2020 challenge on real image denoising: Dataset, methods and results. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (June 2020)
2. Abdelhamed, A., Brubaker, M.A., Brown, M.S.: Noise Flow: Noise Modeling with Conditional Normalizing Flows. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2019)
3. Abdelhamed, A., Lin, S., Brown, M.S.: A high-quality denoising dataset for smart-phone cameras. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
4. Agustsson, E., Timofte, R.: Ntire 2017 challenge on single image super-resolution: Dataset and study. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (2017)
5. Anaya, J., Barbu, A.: Renoir—a dataset for real low-light image noise reduction. *Journal of Visual Communication and Image Representation* **51**, 144–154 (2018)
6. Anwar, S., Barnes, N.: Real image denoising with feature attention. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2019)
7. Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. *PAMI* **33**(5), 898–916 (May 2011). <https://doi.org/10.1109/TPAMI.2010.161>
8. Bachman, P., Hjelm, R.D., Buchwalter, W.: Learning representations by maximizing mutual information across views. In: Advances in Neural Information Processing Systems (NIPS) (2019)
9. Bernardo Henz, Eduardo S. L. Gastal, M.M.O.: Synthesizing camera noise using generative adversarial networks. *IEEE Transactions on Visualization and Computer Graphics* **27**(3), 2123–2135 (2021). <https://doi.org/10.1109/TVCG.2020.3012120>
10. Cai, Y., Hu, X., Wang, H., Zhang, Y., Pfister, H., Wei, D.: Learning to generate realistic noisy images via pixel-level noise-aware adversarial training. In: Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) *Advances in Neural Information Processing Systems (NIPS)* (2021)
11. Chang, K.C., Wang, R., Lin, H.J., Liu, Y.L., Chen, C.P., Chang, Y.L., Chen, H.T.: Learning camera-aware noise models. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 343–358. Springer (2020)
12. Chen, J., Chen, J., Chao, H., Yang, M.: Image blind denoising with generative adversarial network based noise modeling. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3155–3164 (2018)
13. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *International Conference on Machine Learning (ICML)*. pp. 1597–1607. PMLR (2020)
14. Chen, T., Zhai, X., Ritter, M., Lucic, M., Houlsby, N.: Self-supervised gans via auxiliary rotation loss. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12154–12163 (2019)
15. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in Neural Information Processing Systems (NIPS)*. vol. 27. Curran Associates, Inc. (2014)
16. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). vol. 2, pp. 1735–1742. IEEE (2006)

17. Han, J., Shoenby, M., Petersson, L., Armin, M.A.: Dual contrastive learning for unsupervised image-to-image translation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (2021)
18. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9729–9738 (2020)
19. Hong, Z., Fan, X., Jiang, T., Feng, J.: End-to-end unpaired image denoising with conditional adversarial networks. In: Association for the Advancement of Artificial Intelligence (AAAI). vol. 34, pp. 4140–4149 (2020)
20. Jang, G., Lee, W., Son, S., Lee, K.M.: C2n: Practical generative noise modeling for real-world denoising. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 2350–2359 (October 2021)
21. Jeong, J., Shin, J.: Training GANs with stronger augmentations via contrastive discriminator. In: Proceedings of the International Conference on Learning Representations (ICLR) (2021)
22. Kang, M., Park, J.: Contragan: Contrastive learning for conditional image generation. In: Advances in Neural Information Processing Systems (NIPS). vol. 33, pp. 21357–21369. Curran Associates, Inc. (2020)
23. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. In: Advances in Neural Information Processing Systems (NIPS). vol. 33, pp. 18661–18673. Curran Associates, Inc. (2020)
24. Kim, D.W., Ryon Chung, J., Jung, S.W.: Grdn: Grouped residual dense network for real image denoising and gan-based real-world noise modeling. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 0–0 (2019)
25. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Proceedings of the International Conference on Learning Representations (ICLR) (2015)
26. Lee, K.S., Tran, N.T., Cheung, N.M.: Infomax-gan: Improved adversarial image generation via information maximization and contrastive learning. In: IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 3942–3952 (2021)
27. Li, B., Liu, X., Hu, P., Wu, Z., Lv, J., Peng, X.: All-in-one image restoration for unknown corruption. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 17452–17462 (June 2022)
28. Lin, H., Zhuang, Y., Huang, Y., Ding, X., Liu, X., Yu, Y.: Noise2grad: Extract image noise to denoise. In: Zhou, Z.H. (ed.) IJCAI. pp. 830–836. IJCAI (8 2021). <https://doi.org/10.24963/ijcai.2021/115>, main Track
29. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
30. Park, T., Efros, A.A., Zhang, R., Zhu, J.Y.: Contrastive learning for unpaired image-to-image translation. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020)
31. Plotz, T., Roth, S.: Benchmarking denoising algorithms with real photographs. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1586–1595 (2017)
32. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X., Chen, X.: Improved techniques for training gans. In: Advances in Neural Information Processing Systems (NIPS). vol. 29. Curran Associates, Inc. (2016)

33. Wang, L., Wang, Y., Dong, X., Xu, Q., Yang, J., An, W., Guo, Y.: Unsupervised degradation representation learning for blind super-resolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
34. Xu, J., Li, H., Liang, Z., Zhang, D., Zhang, L.: Real-world noisy image denoising: A new benchmark. arXiv preprint arXiv:1804.02603 (2018)
35. Ye, M., Zhang, X., Yuen, P.C., Chang, S.F.: Unsupervised embedding learning via invariant and spreading instance feature. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
36. Yue, Z., Zhao, Q., Zhang, L., Meng, D.: Dual adversarial network: Toward real-world noise removal and noise generation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 41–58. Springer (2020)
37. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L.: Cycleisp: Real image restoration via improved data synthesis. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020)