

CVF-SID: Cyclic multi-Variate Function for Self-Supervised Image Denoising by Disentangling Noise from Image

Reyhaneh Neshatavar^{1*} Mohsen Yavartanoo^{1*} Sanghyun Son¹ Kyoung Mu Lee^{1,2}

¹Dept. of ECE & ASRI, ²IPAI, Seoul National University, Seoul, Korea

{reyhanehneshat, myavartanoo, thstkdgus35, kyoungmu}@snu.ac.kr

Abstract

Recently, significant progress has been made on image denoising with strong supervision from large-scale datasets. However, obtaining well-aligned noisy-clean training image pairs for each specific scenario is complicated and costly in practice. Consequently, applying a conventional supervised denoising network on in-the-wild noisy inputs is not straightforward. Although several studies have challenged this problem without strong supervision, they rely on less practical assumptions and cannot be applied to practical situations directly. To address the aforementioned challenges, we propose a novel and powerful self-supervised denoising method called CVF-SID based on a Cyclic multi-Variate Function (CVF) module and a self-supervised image disentangling (SID) framework. The CVF module can output multiple decomposed variables of the input and take a combination of the outputs back as an input in a cyclic manner. Our CVF-SID can disentangle a clean image and noise maps from the input by leveraging various self-supervised loss terms. Unlike several methods that only consider the signal-independent noise models, we also deal with signal-dependent noise components for real-world applications. Furthermore, we do not rely on any prior assumptions about the underlying noise distribution, making CVF-SID more generalizable toward realistic noise. Extensive experiments on real-world datasets show that CVF-SID achieves state-of-the-art self-supervised image denoising performance and is comparable to other existing approaches. The code is publicly available from this [link](#).

1. Introduction

Image denoising is an active research topic and has attracted increasing attention due to its practicality in computer vision. The fundamental idea of image denoising is to remove unwanted noise signals from a given input and restore a noise-free clean image. Following the recent ad-

*equal contribution

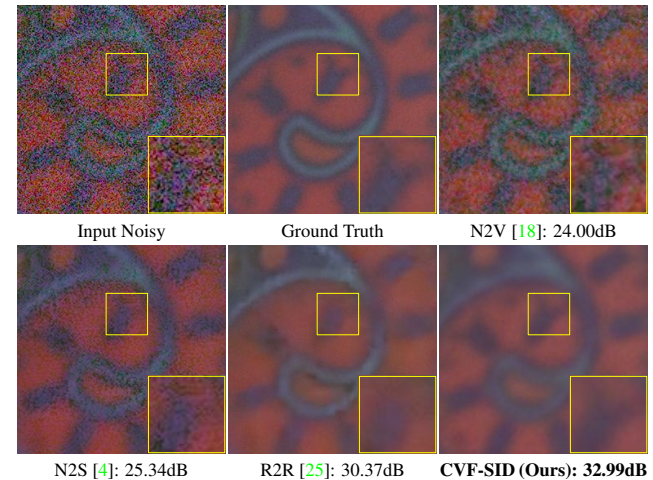


Figure 1. **Real-world image denoising results on the SIDD validation dataset.** In contrast to R2R, our CVF-SID is directly applicable to sRGB images. N2V and N2S fail to restore the clean image, and R2R loses the details evidenced in the yellow patch.

vances in convolutional neural networks (CNN), the latest denoising methods have achieved dramatic performance compared to the traditional algorithms. Specifically, those methods resort to supervised learning on the large-scale synthetic dataset, where noise is simply modeled with additive white Gaussian (AWGN) [9, 11, 20, 33, 35].

Nevertheless, recent studies [3, 12] have observed that the denoising models learned on synthetic images do not generalize well on practical examples. The primary reason for this issue is that real-world noise distribution differs from the synthetic AWGN. To deal with this limitation, few attempts have been made to acquire realistic noisy-clean image pairs [1] in the wild. Still, this process is challenging and sometimes unavailable as it requires multiple shots under the same static scene with several constraints.

Advanced methods overcome the lack of paired images by adopting novel un-/self-supervised frameworks. The generation-based approaches [13, 16] utilize unsupervised adversarial training. They first generate noisy samples from a set of clean images by imitating the noise distribution of

the target dataset. Then, a denoising model can be trained in a supervised manner with the generated noisy-clean image pairs. Nevertheless, they require appropriate clean images that are not always available due to domain differences between noisy and clean images. Rather than using clean ground-truth data, Noise2Noise (N2N) [21] uses two noisy images taken from the same scene and configuration. While N2N shows comparable performance with the supervised methods, it is less practical as multiple noisy images under the same scene are required.

As an alternative solution, several strategies [4, 14, 18, 24, 25, 27] have been proposed to train their methods on noisy images only. To generate feasible input-target pairs from a single noisy image, these recent approaches try to synthesize two independent noisy images from the input. However, these methods are usually constructed by assuming a specific distribution, *e.g.*, AWGN, or less practical configurations for the underlying noise. Such an assumption limits their practical applications where the assumption does not hold. For instance, recent Recorrupted-to-Recorrupted (R2R) [25] is not applicable to sRGB inputs directly, while digital images are usually stored using sRGB color space. Also, this method requires additional prior knowledge, *e.g.*, a pre-trained model with provided noise level function (NLF) [22] by Raw-RGB images, on real-world noise, which is not trainable using sRGB images only.

To mitigate the limitations mentioned above, we present a self-supervised denoising method for the real-world sRGB images. To this end, we propose a novel cyclic multi-variate function (CVF), which disentangles its input into several sub-components and retakes a combination of its outputs as an input. We utilize CVF to design our self-supervised image disentangling model (CVF-SID) for sRGB image denoising. Under various self-supervised training objectives, our CVF-SID can learn to disentangle the noise-free image, signal-dependent and signal-independent noises from a given noisy sRGB input. Furthermore, we propose a self-supervised data augmentation strategy for CVF-SID to effectively increase the number of training samples. Our main contributions can be summarized as follows:

- We introduce CVF-SID, a novel self-supervised method for image denoising based on our defined cyclic multi-variate function (CVF). CVF-SID disentangles a given real-world noisy input to clean image, signal-dependent, and signal-independent noises.
- For fully self-supervised CVF-SID, we propose various training objectives and an augmentation strategy.
- Experimental results demonstrate that our CVF-SID achieves superior denoising performance among several un-/self-supervised methods on real-world sRGB images as shown in Figure 1 and is comparable with the other approaches.




Figure 2. **The proposed Cyclic multi-Variate Function (CVF).** Our CVF f takes a combination $g(s_1X, s_2Y, s_3Z, \dots)$ of multiple variables as an input and outputs the decomposed variables.

2. Related work

Traditional approaches on image denoising such as NLM [5], BM3D [9], WNNM [11], NC [19], K-SVD [28], or EPLL [15] adopt non-learning-based formulation. Recently, deep learning has demonstrated impressive performance on image denoising. In general, they can be categorized based on the way how the network is trained.

2.1. Training on paired noisy-clean images

In general, supervised denoising networks are trained on synthetic noisy-clean images, where the noise is assumed to be additive white Gaussian (AWGN) of a certain level [6, 8]. DnCNN [33] is the first CNN-based approach for image denoising, which introduces residual learning and outperforms the traditional methods. FFDNet [34] further proposes a fast and flexible solution to handle various noise levels within a single model by taking a noise map as an additional input. However, the conventional methods do not generalize well on real-world applications due to domain discrepancy between realistic and synthetic noise. To overcome this limitation, several approaches such as CBDNet [12], RIDNet [3], or DIDN [32] train their methods on realistic noisy-clean pairs [1]. Nevertheless, gathering well-aligned noisy-clean pairs from the real-world scenes is challenging and not very practical, as it requires huge human labor under controlled environments [26].

2.2. Training on unpaired noisy-clean images

To overcome the limitations of the supervised approaches, generation-based methods aim to synthesize noisy samples from clean images in an unsupervised manner [10]. They first try to simulate realistic noise in the adversarial training framework and then train a denoising model on the generated noisy-clean pairs. GCBD [7] is the first generation-based method for blind denoising. However, it is not applicable to real-world scenarios since the method considers additive noise only, while real-world




Figure 3. **Overview of our CVF-SID framework.** (a) We employ the network f_θ to disentangle the clean image, signal-dependent, and signal-independent noises from a noisy input image. In our cyclic network, we feed the outputs of our network again to the network to ensure that each output is pure and does not have the information of the other outputs. (b) We also feed some self-supervised augmented images to our network to better estimate the distribution of each noise and improve the performance.

noise is not. Further, UIDNet [13] employs an image sharpening technique to estimate arbitrary noise distribution in real-world cases. Recent C2N [16] tries to explicitly consider signal-independent, dependent, and spatially-correlated noise in their generation framework. However, such methods require clean images to generate the corresponding noisy images, which is not applicable when the scene distribution of the noisy images is not matched to the existing clean samples.

2.3. Training on paired noisy-noisy images

To alleviate the issues of generation-based approaches, some researchers tried to train their network on paired noisy-noisy images instead of paired noisy-clean or unpaired noisy-noisy images in an un-/self-supervised manner. Noise2Noise (N2N) [21], as a weakly supervised learning on image denoising, proposes to use several noisy images instead of ground truth images. While it can achieve comparable results with supervised methods, taking several independent noisy images from the same scene is also very difficult in real-world cases. Noise2Void (N2V) [18] and Noise2Self (N2S) [4] use a blind-spot learning to avoid learning the identity function without availability of paired images. These approaches remove the center pixel of each receptive field and predict it with other neighboring pixels. However, ignoring some pixels leads to the loss of some useful information and reduces the performance [1].

Noisier2Noise [24] is designed to handle spatially correlated noises by learning on noisy-noisy pairs. Nevertheless, it requires additional information, *e.g.*, the underlying noise distribution, which is a less practical assumption to be applied to real-world inputs. Self2Self (S2S) [27] is proposed on blind denoising to generate paired data from a single noisy image by applying Bernoulli dropout. Later,

Neighbor2Neighbor [14] proposes to create sub-sampled paired images based on pixel-wise independent noise assumption. Recorrupted-to-Recorrupted (R2R) [25] expands the concept of Noisier2Noise [24] toward real-world scenarios. However, R2R resorts to the Gaussian noise assumption when no raw information is provided for a given noisy input image, which prevents its practical applications on digital noisy sRGB images. In contrast, our CVF-SID method can be trained on noisy sRGB images directly without generating pseudo noisy-noisy pairs.

3. Method

We introduce the concept of our CVF and construct our self-supervised denoising model, CVF-SID, on the formulation. For convenience, we denote clean and noisy images as $I_c, I_n \in \mathbb{R}^{H \times W}$, respectively, where the image has a spatial resolution of $H \times W$. Color channels, *i.e.*, RGB, are omitted for simplicity. We represent signal-dependent and independent noise maps as N_d and N_i , respectively, where they have the same dimension to I_n .

3.1. Cyclic multi-Variate Function

We define a cyclic multi-variate function (CVF) f as a mapping from $g(X, Y, Z, \dots)$ to $[X, Y, Z, \dots]$, where X, Y, Z, \dots are vectors and g is a combination function. Therefore, the function can take its outputs as an input again by combining the output values, as shown in Figure 2. Also, for a set of scalar values $[s_1, s_2, s_3, \dots]$, a decomposition of $g(s_1X, s_2Y, s_3Z, \dots)$ should be $[s_1X, s_2Y, s_3Z, \dots]$. By utilizing the aforementioned attributes of CVF, we aim to learn a denoising model in a self-supervised manner.

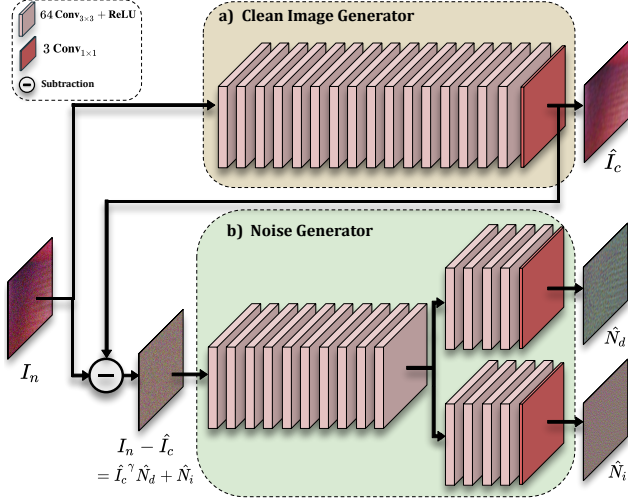


Figure 4. Overview of our network architecture f_θ .

3.2. Self-supervised image denoising using CVF

In general, a noisy image I_n can be expressed as a function of the clean image I_c , signal-dependent noise map N_d , and the signal-independent noise map N_i [29] as follows:

$$I_n = I_c + I_c^\gamma N_d + N_i, \quad (1)$$

where γ is a parameter regarding the correlation between the signal and the corresponding noise term. Previous methods [27, 33] have designed their model to take a single noisy image and reconstruct its clean counterpart, *i.e.*, $f(I_n) = I_c$. In contrast, we utilize the concept of CVF to design our network f_θ with learnable parameters θ for disentangling the given noisy image I_n into the aforementioned three components I_c , N_d , and N_i as shown in Figure 3a. Following the notation in Section 3.1, we denote the noisy image I_n as $g(s_1 I_c, s_2 N_d, s_3 N_i)$, where $s_1 = s_2 = s_3 = 1$.

After decomposing the given noisy image into the clean image and the noise components, we exploit their statistical properties to construct a self-supervised cycle. For practical reasons, we assume that the noise maps N_d and N_i are zero-mean [23] and spatially invariant with unknown distributions. Moreover, we suppose that two elements in each pairs of (I_c, N_d) , (I_c, N_i) , and (N_d, N_i) are independent. In other words, I_c , N_d , and N_i do not contain the information of each other. We note that N_d is a spatially invariant noise map that acts as a scaling factor in the signal-dependent noise term $I_c^\gamma N_d$.

Under such assumptions, f_θ decomposes a given noisy image I_n as follows:

$$\begin{aligned} f_\theta(I_n) &= \left[f_\theta^{\text{clean}}(I_n), f_\theta^{\text{dep}}(I_n), f_\theta^{\text{indep}}(I_n) \right] \\ &= \left[\hat{I}_c, \hat{N}_d, \hat{N}_i \right], \end{aligned} \quad (2)$$

where \hat{I}_c , \hat{N}_d , and \hat{N}_i denotes predicted clean image, signal-dependent noise, and signal-independent noise terms, respectively. Since we do not use any noisy-clean image pairs, it is not possible to apply direct supervisions to the outputs of the function, and it is not guaranteed that f_θ disentangle these components perfectly.

Therefore, to train our model in a self-supervised manner, we feed the outputs \hat{I}_c , \hat{N}_d , and \hat{N}_i again to the same network f_θ with shared parameters θ . Then, we constrain the second outputs, *e.g.*, $f_\theta(\hat{I}_c)$, based on their desired properties. Since a predicted noise-free image \hat{I}_c should not contain any noise, regardless of the signal dependency, it can be modeled as $\hat{I}_c = g(\hat{I}_c, \mathbf{0}, \mathbf{0})$, where $s_1 = 1$ and $s_2 = s_3 = 0$. Here, we use $\mathbf{0}$ to represent an $H \times W$ array of zeros. Therefore, given a predicted clean image \hat{I}_c as an input, our f_θ has to generate the following outputs:

$$f_\theta(\hat{I}_c) = \left[\hat{I}_c, \hat{\mathbf{0}}, \hat{\mathbf{0}} \right], \quad (3)$$

where the predicted output noise maps \hat{N}_d and \hat{N}_i should be zeros which are denoted as $\hat{\mathbf{0}}$.

Our another assumption is derived from the observation that an image corrupted by the signal-dependent noise only should be decomposed to \hat{I}_c and \hat{N}_d or equivalently $\hat{I}_c + \hat{I}_c^\gamma \hat{N}_d = g(\hat{I}_c, \hat{N}_d, \mathbf{0})$, where $s_1 = s_2 = 1$ and $s_3 = 0$. Then, the network f_θ should predict zero as a signal-independent term as follows:

$$f_\theta(\hat{I}_c + \hat{I}_c^\gamma \hat{N}_d) = f_\theta(\hat{I}_n^{\text{dep}}) = \left[\hat{I}_c, \hat{N}_d, \hat{\mathbf{0}} \right], \quad (4)$$

where \hat{I}_n^{dep} is a predicted signal-dependent noisy image.

For a given pure signal-independent part \hat{N}_i , we can regard that the corresponding clean image part is zero. In other word, we can rewrite $\hat{N}_i = g(\mathbf{0}, *, \hat{N}_i)$, where $s_1 = 0$ and $s_3 = 1$. Here, $(*)$ denotes that we do not care about the signal-dependent part. Therefore, our f_θ should predict the same noise \hat{N}_i for the signal-independent path, and also zero for the clean image branch as follows:

$$f_\theta(\hat{N}_i) = \left[\hat{\mathbf{0}}, *, \hat{N}_i \right], \quad (5)$$

where we cannot identify the signal-dependent part $(*)$ as the predicted clean image should be zero.

Lastly, we simulate virtual synthetic noisy images by combination of predicted outputs \hat{I}_c , \hat{N}_d , and \hat{N}_i with the various scalar factors (s_1, s_2, s_3) as shown in Figure 3b. Therefore, we generate the augmented inputs by setting $s_1 = 1$ and selecting s_2 and s_3 from $\{-1, 0, 1\}$ and apply the f_θ as follows:

$$f_\theta(\hat{I}_n^{\text{aug}}) = \left[\hat{I}_c, s_2 \hat{N}_d, s_3 \hat{N}_i \right], \quad (6)$$

where $\hat{I}_n^{\text{aug}} = \hat{I}_c + s_2 \hat{I}_c^\gamma \hat{N}_d + s_3 \hat{N}_i$. This approach operates like self-supervised data augmentation, where no extra samples are required.

3.3. Network architecture

We employ a CNN model with sequential layers as the learnable f_θ as shown in Figure 4. Our model consists of two parts: the clean image generator and the noise generator. Given a noisy image, we employ DnCNN [33] without the skip connection and batch normalization layers (BN) as the clean generator which aims to reconstruct the corresponding noise-free output \hat{I}_c . Then, we subtract the output from the noisy input image and feed this noise signal, *i.e.*, $I_n - \hat{I}_c$, to sequential convolutional layers with two branches to estimate N_d and N_i , respectively. We provide more details in the supplementary material.

3.4. Loss functions for self-supervised learning

To train our network f_θ , we define a set of loss functions based on statistical behaviors of general noise. First, we define the consistency loss \mathcal{L}^{con} to ensure the combination g of outputs \hat{I}_c , \hat{N}_d and \hat{N}_i converges to noisy input as follows:

$$\mathcal{L}^{\text{con}} = \|I_n - g(f_\theta(I_n))\|. \quad (7)$$

For simplicity, we use $\|\cdot\|$ to represent the L^2 norm.

Moreover, we construct the identity loss \mathcal{L}^{id} based on our inter-dependency assumption as follows:

$$\begin{aligned} \mathcal{L}^{\text{id}} = & \left\| \hat{I}_c - f_\theta^{\text{clean}}(\hat{I}_c) \right\| + \left\| \hat{I}_c - f_\theta^{\text{clean}}(\hat{I}_n^{\text{dep}}) \right\| \\ & + \left\| \hat{N}_d - f_\theta^{\text{dep}}(\hat{I}_n^{\text{dep}}) \right\| + \left\| \hat{N}_i - f_\theta^{\text{indep}}(\hat{N}_i) \right\|. \end{aligned} \quad (8)$$

When our network takes a denoised image \hat{I}_c as an input, we expect the model to predict the same image as output without any noise terms. We also define similar loss training objectives for images corrupted with pure signal-dependent noise \hat{I}_n^{dep} and the predicted signal-independent noise \hat{N}_i to construct our identity loss \mathcal{L}^{id} .

On the other hand, we design the zero loss $\mathcal{L}^{\text{zero}}$ to satisfy the constraints in Eq. 3, Eq. 4, and Eq. 5 as follows:

$$\begin{aligned} \mathcal{L}^{\text{zero}} = & \left\| f_\theta^{\text{dep}}(\hat{I}_c) \right\| + \left\| f_\theta^{\text{indep}}(\hat{I}_c) \right\| \\ & + \left\| f_\theta^{\text{clean}}(\hat{N}_i) \right\| + \left\| f_\theta^{\text{indep}}(\hat{I}_n^{\text{dep}}) \right\|. \end{aligned} \quad (9)$$

While Eq. 7, Eq. 8, and Eq. 9 provide several constraints for our self-supervised framework, we further introduce a regularization term to avoid trivial solutions, *e.g.*, zero-valued noise. Our assumption is that the variance of the predicted noise should be positive. Inspired by the patch-based local estimation [23], we first approximate the noise variance from the given image by averaging the variances of M small patches I_n^j as follows:

$$\begin{aligned} \text{Var}(I_n^j) & \approx \text{Var}(\hat{I}_c^j + \hat{I}_c^{j\gamma} \hat{N}_d^j) + \text{Var}(\hat{N}_i^j) \\ & = C_j^{2\gamma} \hat{\sigma}_{j,d}^2 + \hat{\sigma}_{j,i}^2 = C_j^{2\gamma} \hat{\sigma}_d^2 + \hat{\sigma}_i^2, \end{aligned} \quad (10)$$

where $\hat{\sigma}_{j,*}^2 = \text{Var}(\hat{N}_*^j)$ for $* = d$ or i . We assume that an ideal clean image is approximately constant for a j -th patch, *i.e.*, $I_c^j \approx C_j$ and $\text{Var}(I_c^j) \approx 0$, for a *small local region*. Also, we assume that N_d and N_i are spatially uncorrelated, which means that $\hat{\sigma}_{j,*}^2 = \hat{\sigma}_*^2$ is a constant in the image space. Accordingly, we define the regularization loss \mathcal{L}^{reg} to prevent the estimated noise map from having zero-variance, *i.e.*, trivial solution, as follows:

$$\mathcal{L}^{\text{reg}} = \frac{1}{M} \left\| \sum_{j=1}^M \text{Var}(I_n^j) - \sum_{j=1}^M C_j^{2\gamma} \hat{\sigma}_d^2 - M \hat{\sigma}_i^2 \right\|. \quad (11)$$

Our regularization loss \mathcal{L}^{reg} in Eq. 11 is constructed based on an assumption that outputs of the network f_θ are independent to each other.

Finally, we define the augmentation loss \mathcal{L}^{aug} for each of the augmented data in Eq. 6 as follows:

$$\begin{aligned} \mathcal{L}^{\text{aug}} = & \left\| f_\theta^{\text{clean}}(\hat{I}_n^{\text{aug}}) - \hat{I}_c \right\| + \left\| f_\theta^{\text{dep}}(\hat{I}_n^{\text{aug}}) - s_2 \hat{N}_d \right\| \\ & + \left\| f_\theta^{\text{indep}}(\hat{I}_n^{\text{aug}}) - s_3 \hat{N}_i \right\|, \end{aligned} \quad (12)$$

where we calculate Eq. 12 for all possible combinations of s_2 and s_3 in the augmented sample \hat{I}_n^{aug} . Our total training objective $\mathcal{L}^{\text{total}}$ is defined by the summation of all aforementioned loss functions as follows:

$$\mathcal{L}^{\text{total}} = \mathcal{L}^{\text{con}} + \mathcal{L}^{\text{id}} + \mathcal{L}^{\text{zero}} + \mathcal{L}^{\text{reg}} + \lambda_{\text{aug}} \mathcal{L}^{\text{aug}}. \quad (13)$$

4. Experiments

In this section, we first discuss the datasets as well as detailed configurations used to train our CVF-SID framework. We then describe comprehensive experimental results and extensive comparisons with the other methods.

4.1. Dataset

We train and evaluate our method on two datasets, including real-world noisy images: SIDD [1] and DND [26]. **Smartphone Image Denoising Dataset (SIDD)** [1] is one of the representative real-world datasets which contains well-aligned noisy-clean image pairs for training. We use sRGB images from the SIDD-Medium dataset for training, including 320 noisy-clean image pairs. For evaluation, validation and benchmark splits are used that each contains 32 patches of size 256×256 from 40 images where no ground-truth clean images are provided for the benchmark track.

Darmstadt Noise Dataset (DND) [26] benchmark consists of 50 noisy images captured with consumer-grade cameras of various sensor sizes. Each image is cropped into 20 patches of size 512×512 , where total $50 \times 20 = 1,000$ samples are provided for evaluation. Compared to the SIDD dataset, images in the DND dataset are captured under normal lighting conditions and therefore contain weaker noise.

Type of supervision	Training data	Method	SIDD		DND	
			PSNR	SSIM	PSNR	SSIM
Supervised	Paired noisy/clean	MLP [6]	24.71	0.641	34.23	0.833
		TNRD [8]	24.73	0.643	33.65	0.830
		DnCNN [33]	23.66	0.583	32.43	0.790
		DnCNN+ [33]	32.59	0.861	37.90	0.943
		CBDNet [12]	33.28	0.868	38.05	0.942
		RIDNet [3]	38.70	0.950	39.25	0.952
		DIDN [32]	39.82	0.973	39.62	0.954
Unsupervised	Unpaired noisy/clean	GCBD [7]	-	-	35.58	0.922
		UIDNet [13]	32.48	0.897	-	-
		C2N [16]	35.35	0.937	36.38	0.887
	Paired noisy/noisy	R2R [25]	34.78	0.844	-	-
Self-supervised	Paired noisy/noisy	N2V [18]	27.68	0.668	-	-
		N2S [4]	29.56	0.808	-	-
		NAC [30]	-	-	36.20	0.925
	Single noisy	CVF-SID (T) (Ours)	34.43	0.912	36.31	0.923
		CVF-SID (S) (Ours)	34.51	0.916	36.49	0.924
		CVF-SID (S²) (Ours)	34.71	0.917	36.50	0.924

Table 1. **Quantitative comparison of real-world sRGB image denoising on SIDD and DND benchmark datasets.** We compare CVF-SID with other denoising methods in terms of PSNR and SSIM. T, S, and S² refer to different training strategies discussed in Section 4.3.




Figure 5. **Qualitative comparison of different denoising methods on SIDD benchmark.**

4.2. Training details

During the training, we construct a mini-batch of size 64, which contains 40×40 random crops of training images. Random flip and rotation augmentation is applied to increase the number of effective training samples. The learning rate is set to 10^{-4} , where ADAM [17] optimizer is used to update the learnable parameters. For all of our experiments, we fix $\gamma = 1$ in Eq. 1. To apply the regres-

sion loss \mathcal{L}^{reg} in Eq. 11, we densely extract 6×6 patches for each pixel in input images to calculate the approximated variance. In Eq. 13, we set $\lambda_{\text{aug}} = 0.1$ to calculate the total loss $\mathcal{L}^{\text{total}}$. We adopt peak signal-to-noise ratio (PSNR) and structural similarity (SSIM) as evaluation metrics for quantitative comparison. All experiments are done using PyTorch 1.7.1 and Quadro RTX 8000 GPUs. We note that the inference time is about 10ms on average for a given 256×256 input image with only the clean image generator.

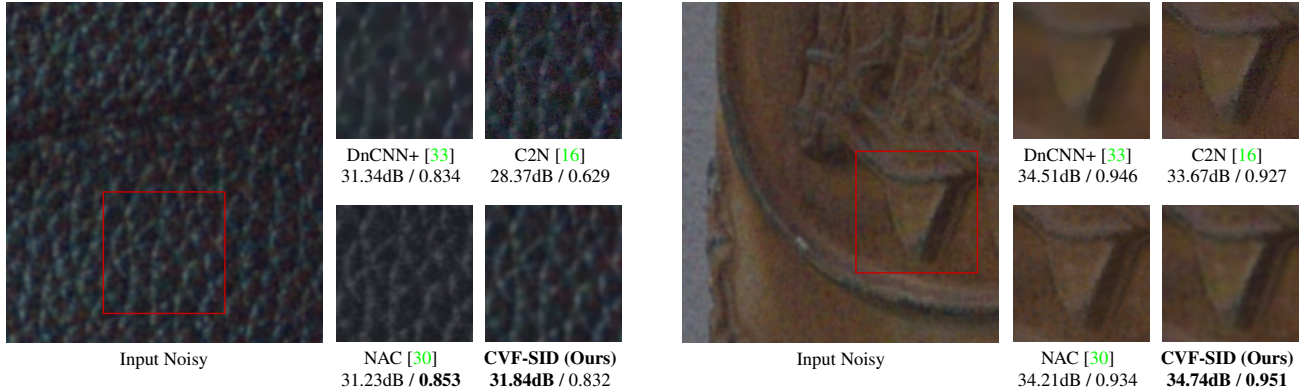


Figure 6. **Visual comparison of different denoising methods on DND benchmark.** DnCNN+ [33], C2N [16], and NAC [30] are supervised, unsupervised, and self-supervised methods, respectively. We report PSNR/SSIM of each result w.r.t. clean ground-truth image.

4.3. Evaluation on real-world sRGB datasets

We evaluate our CVF-SID on real-world sRGB noisy images from SIDD validation, SIDD benchmark, and DND benchmark datasets. For the SIDD and DND benchmarks, we submit the denoising results to websites for server-side evaluation. To evaluate our model on the three different datasets, we leverage three different training strategies. We refer each of them to **T**, **S**, and **S²**, respectively.

Training on SIDD-Medium dataset (T). In the first scenario, we train our CVF-SID on noisy images from the SIDD-Medium dataset. The learned model is then evaluated on three datasets.

Training on test dataset (S). Since CVF-SID is fully self-supervised, we can train our model on the same dataset used for evaluation. Therefore, we train our method on three evaluation datasets, respectively, and test each of them using the same dataset. Since the noise distribution of training and test samples are matched, such a strategy leads to better adaptation to the evaluation datasets.

Double-denoising on test dataset (S²). After training on the test dataset (S), we apply our learned CVF-SID to the test images to acquire denoised images. Then, we use the denoised images as a new dataset to double-train the second CVF-SID. The final denoised results are restored by two successive CVF-SID models on the original noisy images.

Table 1 shows extensive comparisons between several supervised, un-/self-supervised denoising methods on SIDD and DND benchmarks. We note that (T), (S), and (S²) denote our different training strategies described above, respectively. Interestingly, CVF-SID (S) slightly outperforms CVF-SID (T), which is trained on a large SIDD-Medium dataset. This observation validates the advantage of our method, which can be directly trained on test sRGB images without requiring Raw-RGB data or a large number of training samples. Moreover, evaluation of CVF-SID (T) on DND (trained on SIDD, tested on DND) shows the generalization ability of CVF-SID for out-of-domain (or cross-domain) image denoising. Furthermore, we show that CVF-

SID can be trained in a recursive fashion (S²) on the actual test dataset to achieve better denoising performance. Without using any clean images, CVF-SID (S²) outperforms several un-/self-supervised methods on both SIDD and DND benchmarks. Figure 5 and 6 show qualitative comparisons between different denoising methods on these datasets.

The first row of Figure 5 shows that N2S and DnCNN cannot reconstruct characters, while BM3D and N2V cannot perfectly remove noise. In contrast, our CVF-SID can remove unpleasant noise while preserving text details. In the fourth row, we can see that NC and R2R cannot preserve detailed textures while ours can. On the left side of Figure 6, the proposed CVF-SID can reconstruct detailed textures while removing the noise. We note that NAC cannot preserve the original colors compared to the other methods. On the right side, our CVF-SID can preserve edges while suppressing noise from the input.

We also evaluate CVF-SID on SIDD validation dataset using three different training strategies as shown in Table 2. Our approach performs much better than the existing self-supervised methods and unsupervised method C2N [16] and even achieves comparable performance to recent unsupervised R2R [25]. A major merit of CVF-SID compared to R2R [25] and other self-supervised approaches is that we do not generate any paired auxiliary noisy images. Also, we do not assume any specific distribution regarding the unknown noise signals, making our method more generalizable.

As a result, our approach can be applied to sRGB images directly while R2R requires Raw-RGB images for pre-training. We note that Raw-RGB color space contains more information than sRGB, and thus using Raw-RGB samples usually yields better performance than the pure sRGB configuration [1, 2, 26]. Since most digital images are stored in the sRGB format, our CVF-SID can handle more general inputs than R2R. In Figure 7, we visualize how CVF-SID decomposes the given noisy image into the clean image, signal-dependent, and signal-independent noises. We attach more visual comparisons in the supplementary material.

	Method	PSNR	SSIM
Non-learning based	BM3D [9]	25.65	0.475
	WNNM [11]	26.20	0.693
	NC [19]	31.31	0.725
	MCWNNM [31]	33.40	0.815
Unsupervised	C2N [16]	34.08	-
	R2R [25]	35.04	0.844
Self-supervised	N2V [18]	29.35	0.651
	N2S [4]	30.72	0.787
	CVF-SID (T)	34.51	0.941
	CVF-SID (S)	34.67	0.943
	CVF-SID (S²)	34.81	0.944

Table 2. Quantitative comparison of real-world denoising on sRGB images in SIDD validation dataset.

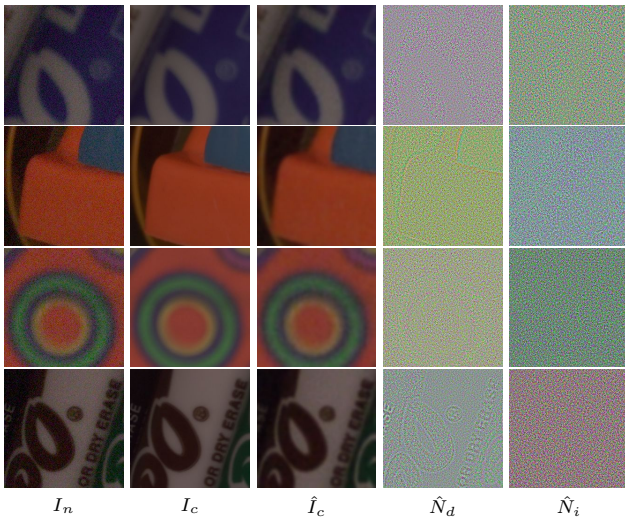


Figure 7. Decomposition results of our CVF-SID (S) on the SIDD validation dataset. For better visualization, the noise maps \hat{N}_d and \hat{N}_i are normalized to $[0, 1]$.

4.4. Ablation study

In this section, we conduct some ablation studies to evaluate the performance of our proposed method better.

Ablation on the loss function. As we have discussed in Section 3.4, we use various types of self-supervised loss terms to train our self-supervised CVF-SID. Table 3 identifies the effect of each loss function for the training. While the consistency term \mathcal{L}^{con} is necessary to train our model, the identity loss \mathcal{L}^{id} brings a significant performance gain and stabilize the learning process.

Effect of the augmentation. To validate the effect of the proposed self-supervised augmentation strategies in Eq. 12, we also train our CVF-SID without and with the augmentation sets A and B in Figure 3b. Table 4a shows the effect of each augmentation set, where set B brings about 0.2dB improvements in the denoising performance.

Ablation on γ correlation. Following Torricelli *et al.* [29], we set the correlation parameter γ in Eq. 1 to 1 to represent

\mathcal{L}^{con}	Loss			PSNR	SSIM
	\mathcal{L}^{id}	$\mathcal{L}^{\text{zero}}$	\mathcal{L}^{reg}		
✓	—	—	—	33.10	0.923
✓	✓	—	—	34.24	0.942
✓	✓	✓	—	34.29	0.940
✓	✓	✓	✓	34.43	0.942

Table 3. Effect of loss terms for our CVF-SID (S) on the SIDD validation dataset. We note that \mathcal{L}^{aug} in Eq. 13 is not used. Please refer to Section 3.4 for more details about each training objective.

Aug.		PSNR	SSIM	γ	PSNR	SSIM
A	B					
—	—	34.43	0.942	0.25	34.45	0.942
✓	—	34.48	0.943	0.50	34.46	0.942
✓	✓	34.67	0.943	1.00	34.67	0.943
✓	✓	34.67	0.943	1.50	34.66	0.943

(a) Effect of the augmentations. (b) Effect of the correlation γ .

Table 4. Effects of different hyperparameters for CVF-SID (S) on the SIDD validation dataset. (a) An overview of the augmentation sets A and B is illustrated in Figure 3b. (b) We find the best correlation parameter γ by grid search.

pure multiplicative noise. Since real-world noise may exhibit more complex behavior, we conduct an ablation study regarding appropriate value for γ . Table 4b shows that CVF-SID achieves the best under the pure multiplicative assumption, *i.e.*, $\gamma = 1$, while increasing the value does not change the performance much.

5. Conclusion

We propose CVF, a novel cyclic multi-variate function that decomposes an input under the cyclic procedure. Then, we utilize CVF to design our self-supervised CVF-SID denoising framework, which aims to learn a CNN to disentangle the signal-dependent, signal-independent noises and clean image from a real-world noisy sRGB input. The proposed approach does not rely on any prior information about the noise distribution, thus more generalizable than previous self-supervised denoising methods. Extensive studies demonstrate several strengths and superiority of our formulation compared to the others. One remaining limitation is that we resort to a fixed correlation parameter γ in our framework, while the correlation may vary for different images in real-world applications. This results in a sub-optimal decomposition, as shown in some examples of Figure 7, where there exists little correlation between the image and signal-dependent noise term. In our future work, we will also aim to learn the correlation parameter in a self-supervised manner while extending the concept of CVF toward various computer vision tasks.

Acknowledgement. This work was supported in part by IITP grant funded by the Korea government (MSIT) [No. 2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University)].

References

- [1] Abdelrahman Abdelhamed, Stephen Lin, and Michael S. Brown. A high-quality denoising dataset for smartphone cameras. In *CVPR*, 2018. 1, 2, 3, 5, 7
- [2] Abdelrahman Abdelhamed, Radu Timofte, and Michael S. Brown. NTIRE 2019 challenge on real image denoising: Methods and results. In *CVPR Workshops*, 2019. 7
- [3] Saeed Anwar and Nick Barnes. Real image denoising with feature attention. In *CVPR*, 2019. 1, 2, 6
- [4] Joshua Batson and Loic Royer. Noise2Self: Blind denoising by self-supervision. In *ICML*, 2019. 1, 2, 3, 6, 8
- [5] A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. In *CVPR*, 2005. 2
- [6] Harold C Burger, Christian J Schuler, and Stefan Harmeling. Image denoising: Can plain neural networks compete with BM3D? In *CVPR*, 2012. 2, 6
- [7] Jingwen Chen, Jiawei Chen, Hongyang Chao, and Ming Yang. Image blind denoising with generative adversarial network based noise modeling. In *CVPR*, 2018. 2, 6
- [8] Yunjin Chen and Thomas Pock. Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration. *IEEE TPAMI*, 2016. 2, 6
- [9] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE TIP*, 2007. 1, 2, 6, 8
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [11] Shuhang Gu, Lei Zhang, Wangmeng Zuo, and Xiangchu Feng. Weighted nuclear norm minimization with application to image denoising. In *CVPR*, 2014. 1, 2, 8
- [12] Shi Guo, Zifei Yan, Kai Zhang, Wangmeng Zuo, and Lei Zhang. Toward convolutional blind denoising of real photographs. In *CVPR*, 2019. 1, 2, 6
- [13] Zhiwei Hong, Xiaocheng Fan, Tao Jiang, and Jianxing Feng. End-to-end unpaired image denoising with conditional adversarial networks. In *AAAI*, 2020. 1, 3, 6
- [14] Tao Huang, Songjiang Li, Xu Jia, Huchuan Lu, and Jianzhuang Liu. Neighbor2Neighbor: Self-supervised denoising from single noisy images. In *CVPR*, 2021. 2, 3
- [15] Samuel Hurault, Thibaud Ehret, and Pablo Arias. EPLL: an image denoising method using a gaussian mixture model learned on a large set of patches. *IPOL*, 2018. 2
- [16] Geonwoon Jang, Wooseok Lee, Sanghyun Son, and Kyoung Mu Lee. C2N: Practical generative noise modeling for real-world denoising. In *CVPR*, 2021. 1, 3, 6, 7, 8
- [17] Diederik P Kingma and J Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [18] Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2Void - learning denoising from single noisy images. In *CVPR*, 2019. 1, 2, 3, 6, 8
- [19] Marc Lebrun, Miguel Colom, and Jean-Michel Morel. The noise clinic: a blind image denoising algorithm. *IPOL*, 2015. 2, 6, 8
- [20] Stamatios Lefkimmiatis. Non-local color image denoising with convolutional neural networks. In *CVPR*, 2017. 1
- [21] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, Timo Aila, et al. Noise2noise. In *ICML*, 2018. 2, 3
- [22] Ce Liu, William T Freeman, Richard Szeliski, and Sing Bing Kang. Noise estimation from a single image. In *CVPR*, 2006. 2
- [23] Xinhao Liu, Masayuki Tanaka, and Masatoshi Okutomi. Practical signal-dependent noise parameter estimation from a single noisy image. *IEEE TIP*, 2014. 4, 5
- [24] Nick Moran, Dan Schmidt, Yu Zhong, and Patrick Coady. Noisier2Noise: Learning to denoise from unpaired noisy data. In *CVPR*, 2020. 2, 3
- [25] Tongyao Pang, Huan Zheng, Yuhui Quan, and Hui Ji. Recorrupted-to-Recorrupted: Unsupervised deep learning for image denoising. In *CVPR*, 2021. 1, 2, 3, 6, 7, 8
- [26] Tobias Plotz and Stefan Roth. Benchmarking denoising algorithms with real photographs. In *CVPR*, 2017. 2, 5, 7
- [27] Yuhui Quan, Mingqin Chen, Tongyao Pang, and Hui Ji. Self2Self with dropout: Learning self-supervised denoising from single image. In *CVPR*, 2020. 2, 3, 4
- [28] Meyer Scetbon, Michael Elad, and Peyman Milanfar. Deep K-SVD denoising. *IEEE TIP*, 2021. 2
- [29] Gionatan Torricelli, Fabrizio Argenti, and Luciano Alparone. Modelling and assessment of signal-dependent noise for image de-noising. In *EUSIPCO*, 2002. 4, 8
- [30] Jun Xu, Yuan Huang, Ming-Ming Cheng, Li Liu, Fan Zhu, Zhou Xu, and Ling Shao. Noisy-as-Clean: Learning self-supervised denoising from corrupted image. *IEEE TIP*, 2020. 6, 7
- [31] Jun Xu, Lei Zhang, David Zhang, and Xiangchu Feng. Multi-channel weighted nuclear norm minimization for real color image denoising. In *ICCV*, 2017. 8
- [32] Songhyun Yu, Bumjun Park, and Jechang Jeong. Deep iterative down-up cnn for image denoising. In *CVPR Workshops*, 2019. 2, 6
- [33] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE TIP*, 2017. 1, 2, 4, 5, 6, 7
- [34] Kai Zhang, Wangmeng Zuo, and Lei Zhang. FFDNet: Toward a fast and flexible solution for CNN based image denoising. *IEEE TIP*, 2018. 2
- [35] Daniel Zoran and Yair Weiss. From learning models of natural image patches to whole image restoration. In *ICCV*, 2011. 1