

Rethinking Transformer-Based Blind-Spot Network for Self-Supervised Image Denoising

Junyi Li, Zhilu Zhang, and Wangmeng Zuo*

Harbin Institute of Technology, Harbin, China
nagejacob@gmail.com, cszlzhang@outlook.com, wnzuo@hit.edu.cn

Abstract

Blind-spot networks (BSN) have been prevalent neural architectures in self-supervised image denoising (SSID). However, most existing BSNs are conducted with convolution layers. Although transformers have shown the potential to overcome the limitations of convolutions in many image restoration tasks, the attention mechanisms may violate the blind-spot requirement, thereby restricting their applicability in BSN. To this end, we propose to analyze and redesign the channel and spatial attentions to meet the blind-spot requirement. Specifically, channel self-attention may leak the blind-spot information in multi-scale architectures, since the downsampling shuffles the spatial feature into channel dimensions. To alleviate this problem, we divide the channel into several groups and perform channel attention separately. For spatial self-attention, we apply an elaborate mask to the attention matrix to restrict and mimic the receptive field of dilated convolution. Based on the redesigned channel and window attentions, we build a Transformer-based **B**lind-**S**pot Network (TBSN), which shows strong local fitting and global perspective abilities. Furthermore, we introduce a knowledge distillation strategy that distills TBSN into smaller denoisers to improve computational efficiency while maintaining performance. Extensive experiments on real-world image denoising datasets show that TBSN largely extends the receptive field and exhibits favorable performance against state-of-the-art SSID methods.

Code — <https://github.com/nagejacob/TBSN>

Introduction

Image denoising is a fundamental low-level vision task that aims to recover implicit clean images from their noisy observations. With the development of convolutional neural networks, learning-based methods (Mao, Shen, and Yang 2016; Zhang et al. 2017; Tai et al. 2017; Lin et al. 2024) have shown significant improvements against traditional ones (Buades, Coll, and Morel 2005; Dabov et al. 2007). In order to facilitate network training, it is common to synthesize noisy-clean image pairs with additive white Gaussian noise (AWGN) for supervised learning. Since the distribution gap between AWGN and camera noise, they exhibit degraded denoising performance in real-world scenarios. One

*Corresponding author.




Figure 1: Visualization of effective receptive field for representative blind-spot networks (BSNs). Benefiting from proposed dilated window and channel attention mechanisms, our TBSN shows stronger local fitting and global information aggregation capability with respect to previous BSNs.

feasible solution is to capture datasets (Plotz and Roth 2017; Abdelhamed, Lin, and Brown 2018) with strictly aligned noisy-clean pairs for network training (Guo et al. 2019; Kim et al. 2020). However, the data collection process requires rigorously controlled environments and much human labor, which is less practical.

Recently, self-supervised image denoising (SSID) (Krull, Buchholz, and Jug 2019; Lee, Son, and Lee 2022; Li et al. 2023; Wang et al. 2023b; Jang et al. 2023a) has been introduced to circumvent the requirement of the paired datasets. The pioneering work Noise2Void randomly masks some locations of a noisy input and trains the network to reconstruct them from their surrounding ones. In the case of pixel-wise independent noise, the network learns to predict the masked pixels without random noise, *i.e.*, the clean pixels, thereby achieving self-supervised denoising. Blind-spot networks (BSN) (Laine et al. 2019; Wu et al. 2020; Byun, Cha, and Moon 2021) take a step further to implement the mask schema with dedicated designed network architectures that exclude the corresponding input pixel from the receptive field of each output location, which shows superiority in both performance and training efficiency. For spatially correlated noise removal in real-world scenarios, some works (Wu et al. 2020; Zhou et al. 2020) suggest first breaking the noise correlation with pixel-shuffle downsampling (PD), and then denoising with BSN. In particular, asymmetric PD strategy (Lee, Son, and Lee 2022; Wang et al. 2023b; Jang et al. 2023a) for training and inference has shown a better trade-off between noise removal and detail preserving.

Existing BSN architectures are mostly convolutional neu-

ral networks (CNNs). However, convolution operations have limited capability in capturing long-range dependencies and the static weights of convolution filters could not flexibly adapt to the input content. As highlighted by the success of the image restoration models (Liang et al. 2021; Zamir et al. 2022; Chen et al. 2023b), such limitations could be mitigated with transformer models (Vaswani et al. 2017). Nonetheless, transformer operators may violate the blind-spot requirement and lead to overfitting to the noisy inputs. Despite the difficulty, few attempts have been made to apply transformers into BSNs. For instance, LG-BPN (Wang et al. 2023b) incorporates channel-wise self-attention (Zamir et al. 2022) for global feature enhancement but still uses convolution layers for local information integration. SwinIA (Papkov and Chizhov 2023) implements a swin-transformer (Liu et al. 2021b) based BSN with modified window attention. However, limited by the requirement for blind spots, it can only exploit shallow features of the noisy inputs in the attention layer, thus showing inferior performance. It can be seen that it is very challenging to bring out the effective capabilities of transformers in BSN.

In this paper, we propose to analyze the spatial and channel self-attention mechanisms and redesign them to meet the blind-spot requirement. For channel-wise self-attention, we observe that simply applying it may leak the blind-spot information, especially in multi-scale architectures. The deep features of such architectures have been downsampled multiple times and the spatial information is shuffled to the channel dimension. The interaction between channels may leak spatial information at the blind-spot, leading to overfitting to the noisy input. We empirically find that this effect appears when the channel dimension is larger than the spatial resolution. To eliminate the undesirable effect, we divide the channels into groups and perform channel attention on each group separately, where the group channel number is controlled less than spatial resolution. For spatial self-attention, we elaborately redesign window attention by restricting its receptive field to maintain the blind-spot requirement. Specifically, a fixed mask is applied to the attention matrix so that each pixel can only attend to pixels at even coordinates. Combining the designed spatial and channel self-attention mechanisms, we propose a dilated transformer attention block (DTAB). We embed DTAB into the encoder-decoder based U-Net architecture, thus presenting a transformer-based blind-spot network (TBSN).

Additionally, BSN architectures are mostly computationally inefficient due to the additional design for satisfying the blind-spot requirement. It becomes even worse with increasing model size and complicated post-refinement process (Lee, Son, and Lee 2022). However, some simple and efficient supervised denoisers have the potential to reach the performance of state-of-the-art SSID methods. In this work, we take advantage of this property to explore a knowledge distillation strategy for reducing the computation cost during inference. Specifically, we regard the results of pre-trained TBSN as pseudo ground-truths, and take them as supervision to train a plain U-Net, namely TBSN2UNet.

Extensive experiments are conducted on real-world denoising datasets (Abdelhamed, Lin, and Brown 2018; Plotz

and Roth 2017) to assess the effectiveness of TBSN and TBSN2UNet. As shown in Fig. 1, benefiting from proposed spatial and channel self-attention mechanisms, TBSN enhances the local adaptivity and largely expands the receptive field. TBSN behaves favorably against state-of-the-art SSID methods in terms of both quantitative metrics and perceptual quality. Moreover, TBSN2UNet maintains the performance of TBSN while significantly reducing inference costs.

Our main contributions can be summarized as follows:

- We propose a transformer-based blind-spot network (TBSN) that contains spatial and channel self-attentions for self-supervised image denoising.
- For channel self-attention, we find it may leak the blind-spot information when the channel number becomes large, we thus perform it on each divided channel group separately to eliminate this adverse effect. For spatial self-attention, we introduce masked window attention where an elaborate mask is applied to the attention matrix to maintain the blind-spot requirement.
- Extensive experiments demonstrate that TBSN achieves state-of-the-art performance on real-world image denoising datasets, while our U-Net distilled from TBSN effectively reduces the computation cost during inference.

Related Work

Deep Image Denoising

The development of learning-based methods (Zhang et al. 2017) has shown superior performance against traditional patch-based ones (Buades, Coll, and Morel 2005; Dabov et al. 2007) on synthetic Gaussian denoising. More advanced deep neural architectures (Mao, Shen, and Yang 2016; Tai et al. 2017; Liu et al. 2018) are further proposed to improve the denoising ability. NBNNet (Cheng et al. 2021) proposes a noise basis network by learning a set of reconstruction basis in the feature space. InvDN (Liu et al. 2021a) proposes a lightweight denoising network based on normalizing flow architectures. Recently, transformers that were first introduced for sequence modeling in natural language processing (Vaswani et al. 2017) have been successfully applied to vision tasks (Dosovitskiy et al. 2020; Liu et al. 2021b). For image denoising, transformers are studied with large-scale image pre-training (Chen et al. 2021) and Swin Transformer architectures (Liang et al. 2021). Restormer (Zamir et al. 2022) and Uformer (Wang et al. 2022) propose multi-scale hierarchical network designs, which achieve better trade-offs between performance and efficiency. However, there are limited efforts for adapting transformers into self-supervised image denoising (Wang et al. 2023b; Papkov and Chizhov 2023) due to the blind-spot requirement.

Self-Supervised Image Denoising

Self-supervised image denoising (SSID) seeks to utilize the information from the noisy images themselves as supervision (Krull, Buchholz, and Jug 2019; Batson and Royer 2019). In order to prevent trivial solutions such as overfitting to the identity mapping, blind-spot networks (BSN) (Wu et al. 2020) exclude the corresponding noisy

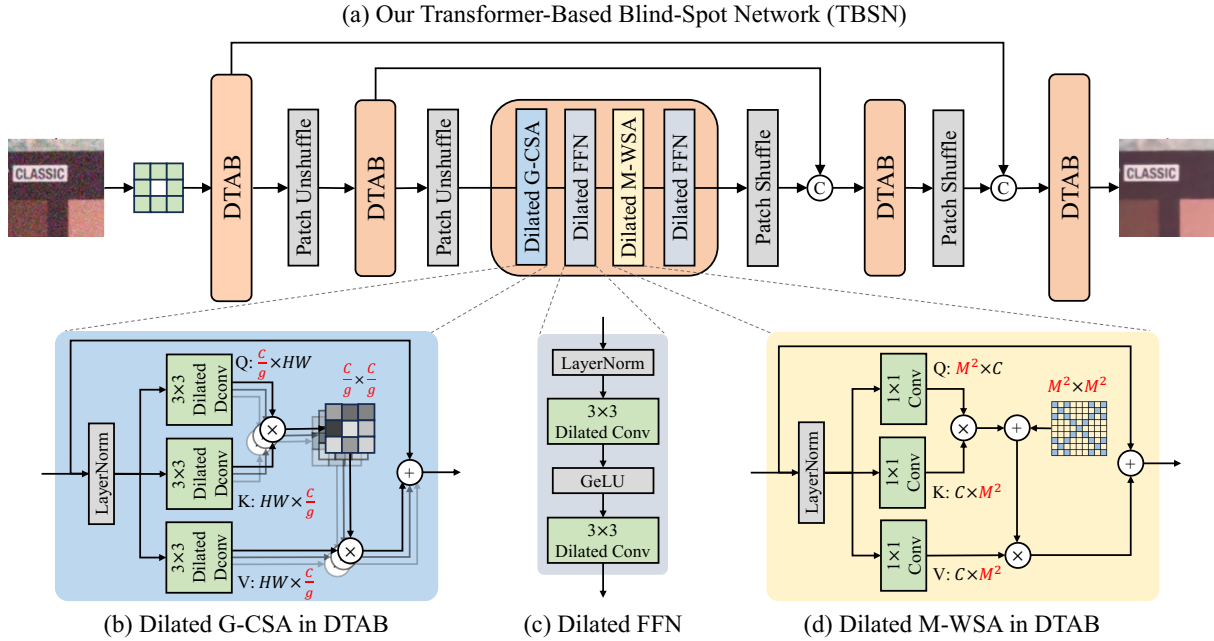


Figure 2: Overview of transformer-based blind-spot network (TBSN). It deploys multi-scale dilated transformer attention blocks (DTAB) to mitigate the shortcomings of convolutions. The core modules of DTAB are dilated counterparts of grouped channel-wise self-attention (G-CSA), masked window-based self-attention (M-WSA), and feed-forward network (FFN), respectively.

pixel from the receptive field at every spatial location. Probabilistic inference (Laine et al. 2019) and regular loss functions (Huang et al. 2021; Zhang et al. 2022) are further introduced to recover the missing information at the blind-spot. For real-world RGB image denoising, the noise is spatially correlated due to the demosaic operation in image signal processing (ISP) pipeline (Guo et al. 2019). It will easily fit the input noise when deploying BSN designed for spatially independent noise removal. One feasible solution is to break the noisy correlation with pixel-shuffle downsampling (Zhou et al. 2020), then apply BSN to the downsampled images (Lee, Son, and Lee 2022; Wang et al. 2023b; Pan et al. 2023; Jang et al. 2023b,a). In addition, CVF-SID (Neshatavar et al. 2022) learns a cyclic function to decompose the noisy image into clean and noisy components. SASL (Li et al. 2023) detects flat and textured areas then constructs supervisions for them separately. Although much effort has been made for developing SSID algorithms (Cheng, Liu, and Tan 2023; Lin et al. 2023; Chen et al. 2023a; Zou, Yan, and Fu 2023; Wang et al. 2023a), there is still a lack of further exploration for BSN architectures. In this work, we adapt the transformer mechanism to BSN to further unleash the potential of blind-spot manners.

Method

Overview of the Network Architecture

As shown in Fig. 2, TBSN follows dilated BSN (Wu et al. 2020) to apply 3×3 centrally masked convolution at the first layer and dilated transformer attention blocks (DTABs) in the remaining layers. The network architecture is U-

Net and adopts patch-unshuffle/shuffle (Jang et al. 2023a) based downsampling/upsampling operations to maintain the blind-spot requirement. The building block, *i.e.*, DTAB, is formed with dilated counterparts of grouped channel-wise self-attention (G-CSA), masked window-based self-attention (M-WSA), and feed-forward network (FFN), respectively. Thus, TBSN benefits from both the global interaction of channel attention and the local fitting ability of window attention. We will provide detailed illustrations of the network design in the following subsections.

Grouped Channel-Wise Self-Attention (G-CSA)

Channel attention (Hu, Shen, and Sun 2018) recalibrates the channel-wise feature responses by explicitly modeling the interdependencies between channels. Given an input feature $\mathbf{X} \in \mathbb{R}^{H \times W \times C}$, channel attention can be formalized as,

$$\text{CA}(\mathbf{X}) = \mathbf{X} * \phi(\mathbf{X}). \quad (1)$$

where function $\phi(\cdot)$ aggregates the spatial information in each channel, and $*$ is channel-wise multiplication operation. For instance, NAFNet (Chen et al. 2022) achieves $\phi(\cdot)$ by global average pooling, while Restormer (Zamir et al. 2022) applies transposed matrix multiplication in the channel dimension. However, in SSID task, channel attention may leak blind-spot information as $\phi(\cdot)$ aggregates the content of all the spatial locations, which is ignored in previous methods.

In this work, we systematically analyze the effects of channel attention (CA) in BSN and empirically find it depends on the channel number versus spatial resolution. For

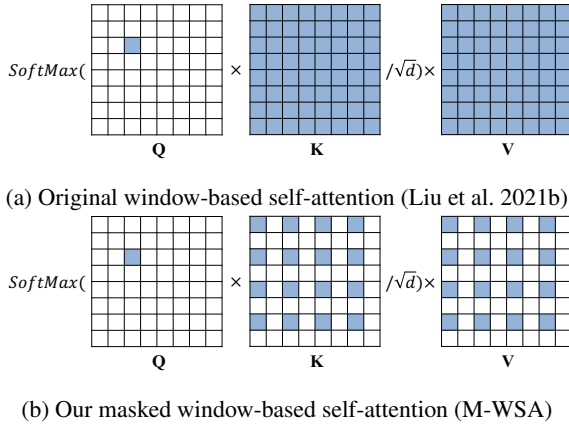


Figure 3: Illustration of receptive field. (a) Original window attention performs all-to-all interactions between query and key/value. (b) Our M-WSA applies an elaborate mask to the attention matrix (see Fig. 2(d)), so the blue pixels of query only attend to the locations of key/value at even coordinates, which mimics the receptive field of dilated convolutions.

single-level architectures (Wang et al. 2023b), spatial information is largely compressed by global average pooling, thus CA is beneficial for performance. For multi-scale architectures (Jang et al. 2023a), the spatial information is shuffled to various channels by downsampling operations. Thus, CA may be partially equivalent to spatial interaction, leaking the blind-spot values. To this end, we propose to control the channel number smaller than spatial resolution. Specifically, we introduce grouped channel-wise self-attention (G-CSA) to divide the deep feature into multiple channel groups and perform CA separately. Our G-CSA could be formulated as,

$$\text{G-CSA}(\mathbf{X}) = \text{Concat}(\mathbf{X}_1 * \phi(\mathbf{X}_1), \dots, \mathbf{X}_G * \phi(\mathbf{X}_G)). \quad (2)$$

where $\mathbf{X} = \text{Concat}(\mathbf{X}_1, \dots, \mathbf{X}_G)$, G is the group number. We set the channel number of each group, *i.e.*, $\frac{C}{G}$, to be small enough to avoid the leakage of spatial information. In the implementation, we adapt MDTA (Zamir et al. 2022) to our G-CSA with Eq. (2) for global interaction. We also replace the 3×3 depth convolutions with their dilated counterparts to achieve blind-spot requirement, as shown in Fig. 2(b).

Masked Window-Based Self-Attention (M-WSA)

Window-based self-attention (Liu et al. 2021b) has been wildly applied in image restoration (Liang et al. 2021; Chen et al. 2023b). In this work, we mimic the behavior of dilated convolutions (Wu et al. 2020; Wang et al. 2023b; Jang et al. 2023a) to propose a masked window-based self-attention (M-WSA) for SSID, which can be plug-and-played into any layer and exploit current deep features as *query/key/value*. As shown in Fig. 2(d), we ingeniously design a fixed attention mask adding to the attention matrix to restrict the interactions between *query* and *key/value* tokens. From Fig. 3(a), in original window attention, each *query* token interacts with *key/value* tokens at all spatial locations within the window. In our M-WSA, the *query* token attends to the spatial locations at even coordinates (see Fig. 3(b)). Therefore, M-WSA

exhibits the same functionality as dilated convolutions for building BSN, but with a larger receptive field and stronger local fitting capability.

Here we formally illustrate our attention mask. In window attention, within a local window of size $M \times M$, the current feature is first projected to *query*, *key* and *value* tokens as $\mathbf{Q}, \mathbf{K}, \mathbf{V} \in \mathbb{R}^{M^2 \times d}$, respectively. Then the original window attention can be formulated as,

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{SoftMax}(\mathbf{Q}\mathbf{K}^T / \sqrt{d})\mathbf{V}. \quad (3)$$

where d is the feature dimension. In our M-WSA, our attention mask $\mathbf{M} \in \mathbb{R}^{M^2 \times M^2}$ is applied to the attention matrix that restricts each *query* only attends to *key/value* at even coordinates, as shown in Fig. 2(d). Thus, Eq. (3) can be modified as,

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{SoftMax}(\mathbf{Q}\mathbf{K}^T / \sqrt{d} + \mathbf{M})\mathbf{V}, \quad (4)$$

$$\mathbf{M}(i, j) = \begin{cases} 0, & \text{if } x_i - x_j \equiv y_i - y_j \equiv 0 \pmod{2} \\ -\infty, & \text{otherwise} \end{cases}. \quad (5)$$

Specifically, \mathbf{M} is a two-valued matrix that masks out certain locations according to the relative position of *query* (at i) and *key/value* (at j) tokens. (x_i, y_i) and (x_j, y_j) are the spatial locations of i and j . When i and j are with even distance on both axes, $\mathbf{M}(i, j) = 0$, the attention value is unchanged. Otherwise, $\mathbf{M}(i, j) = -\infty$ and the attention value becomes 0 after the softmax operation, thereby being masked out. Inspired by relative position embedding (Liu et al. 2021b), $\mathbf{M} \in \mathbb{R}^{M^2 \times M^2}$ can be calculated from a smaller-sized binary matrix $\hat{\mathbf{M}} \in \mathbb{R}^{(2M-1) \times (2M-1)}$ according to the relative position of i and j to improve the efficiency, *i.e.*,

$$\mathbf{M}(i, j) = \begin{cases} 0, & \text{if } \hat{\mathbf{M}}(x_i - x_j, y_i - y_j) = 0 \\ -\infty, & \text{if } \hat{\mathbf{M}}(x_i - x_j, y_i - y_j) = 1 \end{cases}, \quad (6)$$

$$\hat{\mathbf{M}}(x, y) = \begin{cases} 0, & \text{if } x \equiv y \equiv 0 \pmod{2} \\ 1, & \text{otherwise} \end{cases}. \quad (7)$$

In the implementation perspective, we adopt overlapping cross-attention (Chen et al. 2023b) that calculates *key/value* tokens from a larger field to further expand receptive field.

Discussion. The proposed G-CSA and M-WSA are distinct from the transformer operators in previous BSNs. As shown in Fig. 4(a)(b), channel attention in LGBPN (Wang et al. 2023b) is risked to leak blind-spot information when applied in multi-scale architectures, while our G-CSA performs channel attention in separate groups to alleviate this problem. From Fig. 4(c)(d), window attention in SwinIA (Papkov and Chizhov 2023) masks the main diagonal of the attention matrix to maintain the blind-spot requirement. Its *key/value* tokens are limited to be from pixel-wise shallow features of the noisy input, thus showing inferior results. In contrast, our M-WSA applies a dedicated designed mask to mimic the behavior of dilated convolution, which can be flexibly performed on the deep features.

Knowledge Distillation for Efficient Inference

Self-supervised image denoising methods usually require high computational cost due to complicated network designs (Wu et al. 2020), increased network size (Jang et al.

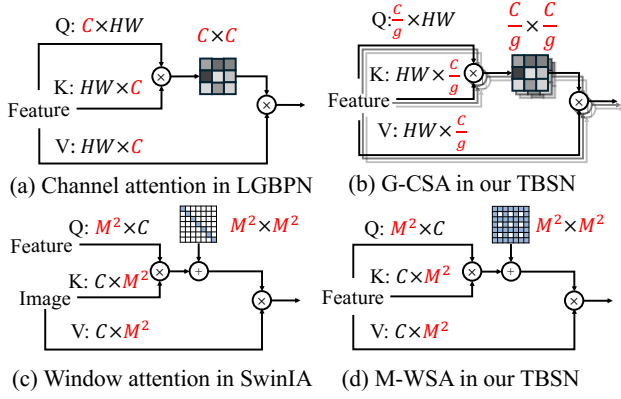


Figure 4: Channel and window attentions in LG-BPN (Wang et al. 2023b) and SwinIA (Papkov and Chizhov 2023). Our G-CSA and M-WSA are distinct from these modules.

2021), and post-processing operation (Lee, Son, and Lee 2022). The computation burden largely limits their applicability in certain situations, *e.g.*, on mobile devices. Nonetheless, the performance of SSID methods still falls behind the corresponding supervised ones. Even lightweight supervised methods may achieve better performance than complex self-supervised ones. In other words, the lightweight network may be fully sufficient to fit the results from some complex self-supervised methods. Taking advantage of this, we suggest a knowledge distillation strategy to reduce the inference cost while maintaining the performance.

Specifically, we adopt the efficient U-Net (Ronneberger, Fischer, and Brox 2015) architecture as our student network, which is distilled from the self-supervised learned TBSN (namely TBSN2UNet),

$$\mathcal{L}_{distill} = \|sg(\mathbf{TBSN}(y)) - \mathbf{U-Net}(y)\|_1 \quad (8)$$

where y is the noisy image, $sg(\cdot)$ is the stop gradient operation. Note that we aim to reduce the computation cost during inference. It is different from the methods that apply knowledge distillation for better performance (Wu et al. 2020; Jang et al. 2021; Li et al. 2023).

Experiments

Implementation Details

Datasets. We conduct experiments on two widely used real-world image denoising datasets, *i.e.*, SIDD (Abdelhamed, Lin, and Brown 2018) and DND (Plotz and Roth 2017). The noisy-clean pairs of SIDD dataset are collected from five smartphone cameras, where each noisy image is captured multiple times and the average image serves as ground truth. It contains 320 training images, 1280 validation patches and 1280 benchmark patches, respectively. We train our networks on the noisy images of train split, and test on the benchmark split. DND is a benchmark dataset collected from DSLR cameras. The noisy images are captured with a short exposure time while the corresponding clean images are captured with a long exposure time. It contains 50 pairs

	Method	SIDD Bench PSNR / SSIM	DND Bench PSNR / SSIM
Non-learning Based	BM3D	25.65 / 0.685	34.51 / 0.851
	WNNM	25.78 / 0.809	34.67 / 0.865
Supervised (Synthetic pairs)	DnCNN	26.25 / 0.599	32.43 / 0.790
	CBDNet	33.28 / 0.868	38.05 / 0.942
	Zhou <i>et al.</i>	34.00 / 0.898	38.40 / 0.945
Supervised (Real pairs)	DnCNN	37.61 / 0.941	38.73 / 0.945
	VDN	39.26 / 0.955	39.38 / 0.952
	Restormer	40.02 / 0.960	40.03 / 0.956
	NAFNet	40.30 / 0.961	- / -
Unpaired	GCBD	-	35.58 / 0.922
	UIDNet	32.48 / 0.897	-
	C2N	35.35 / 0.937	37.28 / 0.924
	DBSN	-	37.93 / 0.937
	Noise2Void	27.68 / 0.668	-
Self-Supervised	Noise2Self	29.56 / 0.808	-
	NAC	-	36.20 / 0.925
	R2R	34.78 / 0.898	-
	CVF-SID	34.71 / 0.917	36.50 / 0.924
	AP-BSN	36.91 / 0.931	38.09 / 0.937
	SASL	37.41 / 0.934	38.18 / 0.938
	LG-BPN	37.28 / 0.936	38.43 / 0.942
	PUCA	37.54 / 0.936	38.83 / 0.942
	AT-BSN	<u>37.78</u> / 0.944	38.68 / 0.942
Self-Supervised (Ours)	TBSN	<u>37.78</u> / 0.940	39.08 / 0.945
	TBSN2UNet	37.79 / 0.940	<u>39.01</u> / 0.945

Table 1: Quantitative comparison on SIDD and DND benchmark datasets. In self-supervised category, the first and second place results are highlighted in **bold** and underlined.

for test only. We train and test our networks on the test images in a fully self-supervised manner.

Training Details. For self-supervised training of TBSN, we follow AP-BSN (Lee, Son, and Lee 2022) to apply pixel-shuffle downsampling (PD) to break the noise correlation, and adopt asymmetric PD factors during training and inference to trade-off the denoising effect and detail preserving. We also improve the denoising results with random replacement refinement (R3) strategy. The batch size and patch size are set to 4 and 128×128 , respectively. We adopt ℓ_1 loss and AdamW (Loshchilov and Hutter 2018) optimizer to train the network. The learning rate is initially set to 3×10^{-4} , and is decreased by 10 every 40k iterations with total 100k training iterations. For knowledge distillation, the training settings are the same as self-supervised learning. All the experiments are conducted on PyTorch framework and Nvidia RTX2080Ti GPUs.

Comparison with State-of-the-Art Methods

Quantitative comparison. Tab. 1 shows the quantitative results of proposed TBSN and state-of-the-art self-supervised methods: Noise2Void (Krull, Buchholz, and Jug 2019), Noise2Self (Batson and Royer 2019), NAC (Xu et al. 2020), R2R (Pang et al. 2021), CVF-SID (Neshatavar et al. 2022), AP-BSN (Lee, Son, and Lee 2022), SASL (Li et al. 2023),

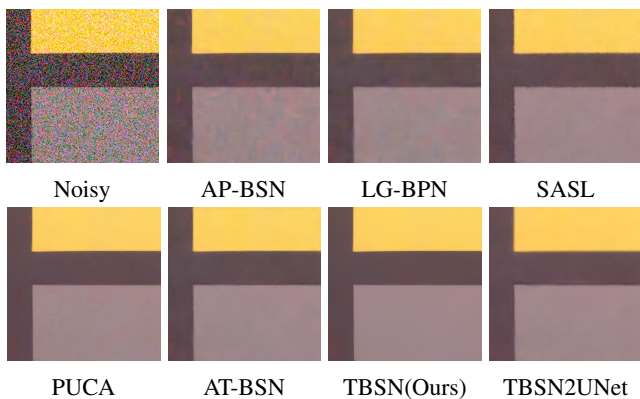


Figure 5: Qualitative comparison on SIDD dataset, please zoom in for better observation.




Figure 6: Qualitative comparison on DND dataset.

LG-BPN (Wang et al. 2023b), PUCA (Jang et al. 2023a) and AT-BSN (Chen et al. 2024). Among these, blind-spot techniques designed for spatial independent noise (Noise2Void, Noise2Self, and R2R) exhibit little denoising effect on real-world noisy images. Although pixel-shuffle downsampling (PD) breaks the noise correlation and successfully removes the noise (Lee, Son, and Lee 2022), the performance is still limited by its plain convolutional BSN architecture. Some recent works tackle this problem by searching for advanced BSN architectures. For instance, LG-BPN incorporates transformer block (Zamir et al. 2022) into BSN for global information and brings 0.37dB improvement over AP-BSN baseline. PUCA designs a multi-scale BSN with channel attention and achieves 0.63dB improvement. Nonetheless, benefiting from channel and window attention mechanisms, our TBSN boosts the improvement to 0.87dB on SIDD benchmark dataset. AT-BSN introduces multi-teacher distillation strategy, in comparison, our TBSN2UNet achieves 0.01dB and 0.33dB improvement against AT-BSN in SIDD and DND Benchmark datasets respectively, showing the superiority of our method.

Qualitative comparison. The qualitative results of self-supervised image denoising methods are shown in Fig. 5 and Fig. 6. The denoising of the color chart in Fig. 5 depends on

Method	SASL	LG-BPN	PUCA	TBSN	TBSN2UNet
PSNR (dB)	37.41	37.28	37.54	37.78	37.79
#Param (M)	1.08	6.61	12.78	12.97	1.08
#FLOPs (G)	35.0	6699.6	2644.2	5463.9	35.0
Time (ms)	4.8	5229.4	444.7	1032.1	4.8

Table 2: Model complexity comparison on SIDD benchmark dataset. #FLOPs and time are measured on 256×256 size.

the global information, where former methods fail to wipe the noise completely. Benefiting from the channel-wise self-attention, our TBSN removes the spatially correlated noise smoothly. The cup and the wall in Fig. 5 and Fig. 6 show that TBSN could maintain the details due to its local fitting capability of window attention.

Comparison of Model Complexity

The channel-wise (Zamir et al. 2022) and window-based (Liu et al. 2021b) self-attentions in TBSN are efficient transformer modules designed for image restoration task. In addition, TBSN adopts hierarchical multi-scale architecture to further improve its efficiency. As shown in Tab. 2, TBSN maintains similar computation complexity as the convolutional counterparts PUCA, and is more efficient than LG-BPN. In addition, SASL shows attractive #Param and #FLOPs results due to its U-Net architecture. For a fair comparison, the U-Net distilled from our TBSN exhibits the same computation complexity as SASL but with higher performance, which demonstrates the superiority of our knowledge distillation strategy.

Ablation Study

Visualization of the Receptive Field

The expansion of receptive field is a major factor towards the success of transformers (Zamir et al. 2022). In this subsection, we plot the input pixels for recovering the center pixel of output to access the effective receptive field of TBSN. Specifically, we pass an input image through the network, select the center pixel of the output image, and calculate its gradient with respect to the input image. The pseudo code in PyTorch format is as follows,

```
import torch as t

image = t.rand((1, 3, H, W), requires_grad=True)
output = model(image)
center_pixel = t.mean(output[...], H//2, W//2)
center_pixel.backward()
gradient = t.sum(t.abs(x.grad), dim=1, keepdim=True)
```

Such gradient indicates how much the output pixel changes with a disturbance on each input pixel. We sum the absolute value of the gradient along the channel axis for visualization. From Fig 1, TBSN shows a significantly wider receptive field than former BSNs (Lee, Son, and Lee 2022; Wang et al. 2023b; Jang et al. 2023a), which is a possible explanation for the appealing performance of TBSN.

Model	Window Attention	Channel Attention	PSNR
(1)	-	-	36.90
(2)	Dilated M-WSA	-	37.17
(3)	-	Dilated G-CSA	37.55
(4)	SwinIA	Dilated G-CSA	37.12
(5)	Swin Trans.	Dilated G-CSA	37.56
(6)	Dilated M-WSA	LG-BPN	37.66
(7)	Dilated M-WSA	SE	37.43
(8)	Dilated M-WSA	SCA	37.63
TBSN	Dilated M-WSA	Dilated G-CSA	37.71

Table 3: Ablation study of dilated transformer attention block (DTAB) on SIDD validation dataset. ‘-’ refers to plain dilated convolutions. (1~3) indicates our dilated M-WSA and G-CSA are essential components of DTAB. (4~8) shows they are not replaceable with analogous operators.

Analysis on DTAB

Tab. 3 analyzes the effectiveness of the components in our dilated transformer attention block (DTAB). We begin with a base model (1) that degenerates the window and channel attentions to dilated convolutions (Wu et al. 2020). In contrast, our dilated M-WSA (2) enhances the base model with local fitting capability and provides 0.27dB improvement. Our dilated G-CSA (3) exhibits global interaction and shows 0.65dB improvement. TBSN achieves a total improvement of 0.81 with combined channel and window attentions, which demonstrates the complementarity of the local and global operations. In addition, we assess the effects of other window and channel attention implementations in Fig. 4. Replacing M-WSA with SwinIA (4) leads to 0.59dB performance drop while replacing G-CSA with LG-BPN (6) leads to 0.05dB performance drop. The other attention mechanisms, Swin Transformer (Liu et al. 2021b) (5), SE (Hu, Shen, and Sun 2018) (7) and SCA (Chen et al. 2022) (8) also shows inferior performance. To summarize, the ablation study results suggest that DTAB is the optimal network choice.

Analysis on Channel Attention

As illustrated in the method, channel-wise self-attention (CSA) (Wang et al. 2023b) may leak blind-spot information in multi-scale architectures. Tab. 4 analyzes the effect of CSA with downsampling scales from 1 to 5. In correspondence, the channel number at the deepest layers grows from 48 to 768, and the spatial resolution reduces from 128^2 to 8^2 , respectively. From the middle lines of Tab. 4, the plain dilated CSA provides positive effects at the scales less equal to 3, but leads to obvious performance drop at the 4- and 5-scales. This is owing to the channel dimension being larger than the spatial resolution at 4- and 5-scales so it leaks the blind-spot information. Instead, our grouped dilated G-CSA divides the channels into several groups and performs channel attention separately. As the channel dimension within each group is controlled smaller than the spatial resolution, dilated G-CSA provides constant improvement at all scales.

#Scale	1	2	3	4	5
Max. #Channel	48	96	192	384	768
Min. #Spatial	128^2	64^2	32^2	16^2	8^2
w/o Dilated CSA	37.15	37.29	37.55	37.07	37.02
w/ Dilated CSA	37.33	37.51	37.66	36.84	11.35
#Group	1	2	4	8	16
#Channel / #Group	48	48	48	48	48
w/ Dilated G-CSA	37.34	37.52	37.71	37.38	37.31

Table 4: Ablation study of channel-wise self-attention (CSA) on SIDD validation dataset. Applying CSA to the 4- or 5-scale architectures causes performance drop as it leaks the blind-spot information. Our grouped CSA (G-CSA) eliminates the negative effects by splitting the channels into groups and performing CSA separately.

Training	Model	PSNR (dB)	#Param (M)	#FLOPs (G)	Time (ms)
Supervised	Network: U-Net	38.92	1.08	35.0	4.8
Knowledge Distillation	Teacher: TBSN	37.71	12.97	607.1	1032.1
	Student: U-Net	37.70	1.08	35.0	4.8

Table 5: Ablation study of knowledge distillation on SIDD validation dataset.

Analysis on Knowledge Distillation

As shown in Tab. 5, we conduct experiments with U-Net architecture to assess the effectiveness of knowledge distillation. Despite fewer parameters and FLOPs, U-Net trained in a supervised manner achieves 1.21dB higher performance than TBSN on SIDD validation dataset. It demonstrates lightweight U-Net has enough learning capacity to receive the denoising performance of TBSN. Consequently, the distilled student U-Net achieves comparable results as the teacher TBSN. The results show that knowledge distillation is a feasible way to reduce the model size and computation cost during inference in SSID.

Conclusion

In this paper, we propose a transformer-based blind-spot network, namely TBSN, for self-supervised image denoising. Key designs are introduced to adapt the spatial and channel self-attention operators for constructing BSN. For spatial attention, an elaborate mask is applied to the window attention matrix, thus restricting its receptive field to mimic the dilated convolutions. For the spatial information leakage problem of channel attention, we propose to perform channel attention in separate groups to eliminate its harmful effects. Moreover, a knowledge distillation strategy is introduced to reduce the computation cost during inference. Extensive experiments on real-world denoising datasets demonstrate that TBSN largely expands the effective receptive field and achieves state-of-the-art performance.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grants 62371164 and U22B2035.

References

- Abdelhamed, A.; Lin, S.; and Brown, M. S. 2018. A high-quality denoising dataset for smartphone cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1692–1700.
- Batson, J.; and Royer, L. 2019. Noise2self: Blind denoising by self-supervision. In *International Conference on Machine Learning*, 524–533. PMLR.
- Buades, A.; Coll, B.; and Morel, J.-M. 2005. A non-local algorithm for image denoising. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, volume 2, 60–65. Ieee.
- Byun, J.; Cha, S.; and Moon, T. 2021. Fbi-denoiser: Fast blind image denoiser for poisson-gaussian noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5768–5777.
- Chen, H.; Qu, C.; Zhang, Y.; Chen, C.; and Jiao, J. 2023a. Multi-view self-supervised disentanglement for general image denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12281–12291.
- Chen, H.; Wang, Y.; Guo, T.; Xu, C.; Deng, Y.; Liu, Z.; Ma, S.; Xu, C.; Xu, C.; and Gao, W. 2021. Pre-trained image processing transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12299–12310.
- Chen, L.; Chu, X.; Zhang, X.; and Sun, J. 2022. Simple baselines for image restoration. In *European Conference on Computer Vision*, 17–33. Springer.
- Chen, S.; Zhang, J.; Yu, Z.; and Huang, T. 2024. Exploring Efficient Asymmetric Blind-Spots for Self-Supervised Denoising in Real-World Scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2814–2823.
- Chen, X.; Wang, X.; Zhou, J.; Qiao, Y.; and Dong, C. 2023b. Activating more pixels in image super-resolution transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22367–22377.
- Cheng, J.; Liu, T.; and Tan, S. 2023. Score priors guided deep variational inference for unsupervised real-world single image denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12937–12948.
- Cheng, S.; Wang, Y.; Huang, H.; Liu, D.; Fan, H.; and Liu, S. 2021. Nbnnet: Noise basis learning for image denoising with subspace projection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4896–4906.
- Dabov, K.; Foi, A.; Katkovnik, V.; and Egiazarian, K. 2007. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8): 2080–2095.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Guo, S.; Yan, Z.; Zhang, K.; Zuo, W.; and Zhang, L. 2019. Toward convolutional blind denoising of real photographs. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1712–1722.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7132–7141.
- Huang, T.; Li, S.; Jia, X.; Lu, H.; and Liu, J. 2021. Neighbor2neighbor: Self-supervised denoising from single noisy images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14781–14790.
- Jang, G.; Lee, W.; Son, S.; and Lee, K. M. 2021. C2n: Practical generative noise modeling for real-world denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2350–2359.
- Jang, H.; Park, J.; Jung, D.; Lew, J.; Bae, H.; and Yoon, S. 2023a. PUCA: patch-unshuffle and channel attention for enhanced self-supervised image denoising. *Advances in Neural Information Processing Systems*, 36.
- Jang, Y. I.; Lee, K.; Park, G. Y.; Kim, S.; and Cho, N. I. 2023b. Self-supervised image denoising with downsampled invariance loss and conditional blind-spot network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12196–12205.
- Kim, Y.; Soh, J. W.; Park, G. Y.; and Cho, N. I. 2020. Transfer learning from synthetic to real-noise denoising with adaptive instance normalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3482–3492.
- Krull, A.; Buchholz, T.-O.; and Jug, F. 2019. Noise2void-learning denoising from single noisy images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2129–2137.
- Laine, S.; Karras, T.; Lehtinen, J.; and Aila, T. 2019. High-quality self-supervised deep image denoising. *Advances in Neural Information Processing Systems*, 32.
- Lee, W.; Son, S.; and Lee, K. M. 2022. AP-BSN: Self-Supervised Denoising for Real-World Images via Asymmetric PD and Blind-Spot Network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17725–17734.
- Li, J.; Zhang, Z.; Liu, X.; Feng, C.; Wang, X.; Lei, L.; and Zuo, W. 2023. Spatially Adaptive Self-Supervised Learning for Real-World Image Denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9914–9924.
- Liang, J.; Cao, J.; Sun, G.; Zhang, K.; Van Gool, L.; and Timofte, R. 2021. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1833–1844.

- Lin, J.; Zhang, Z.; Wei, Y.; Ren, D.; Jiang, D.; Tian, Q.; and Zuo, W. 2024. Improving image restoration through removing degradations in textual representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2866–2878.
- Lin, X.; Ren, C.; Liu, X.; Huang, J.; and Lei, Y. 2023. Un-supervised image denoising in real-world scenarios via self-collaboration parallel generative adversarial branches. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12642–12652.
- Liu, P.; Zhang, H.; Zhang, K.; Lin, L.; and Zuo, W. 2018. Multi-level wavelet-CNN for image restoration. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 773–782.
- Liu, Y.; Qin, Z.; Anwar, S.; Ji, P.; Kim, D.; Caldwell, S.; and Gedeon, T. 2021a. Invertible denoising network: A light solution for real noise removal. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13365–13374.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021b. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Loshchilov, I.; and Hutter, F. 2018. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- Mao, X.; Shen, C.; and Yang, Y.-B. 2016. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. *Advances in neural information processing systems*, 29.
- Neshatavar, R.; Yavartanoo, M.; Son, S.; and Lee, K. M. 2022. CVF-SID: Cyclic multi-Variate Function for Self-Supervised Image Denoising by Disentangling Noise from Image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17583–17591.
- Pan, Y.; Liu, X.; Liao, X.; Cao, Y.; and Ren, C. 2023. Random Sub-Samples Generation for Self-Supervised Real Image Denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Pang, T.; Zheng, H.; Quan, Y.; and Ji, H. 2021. Recorrupted-to-recorrupted: unsupervised deep learning for image denoising. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2043–2052.
- Papkov, M.; and Chizhov, P. 2023. SwinIA: Self-Supervised Blind-Spot Image Denoising with Zero Convolutions. *arXiv preprint arXiv:2305.05651*.
- Plotz, T.; and Roth, S. 2017. Benchmarking denoising algorithms with real photographs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1586–1595.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241. Springer.
- Tai, Y.; Yang, J.; Liu, X.; and Xu, C. 2017. Memnet: A persistent memory network for image restoration. In *Proceedings of the IEEE international conference on computer vision*, 4539–4547.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, J.; Di, S.; Chen, L.; and Ng, C. W. W. 2023a. Noise2info: Noisy image to information of noise for self-supervised image denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 16034–16043.
- Wang, Z.; Cun, X.; Bao, J.; Zhou, W.; Liu, J.; and Li, H. 2022. Uformer: A general u-shaped transformer for image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 17683–17693.
- Wang, Z.; Fu, Y.; Liu, J.; and Zhang, Y. 2023b. LG-BPN: Local and Global Blind-Patch Network for Self-Supervised Real-World Denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18156–18165.
- Wu, X.; Liu, M.; Cao, Y.; Ren, D.; and Zuo, W. 2020. Un-paired learning of deep image denoising. In *European conference on computer vision*, 352–368. Springer.
- Xu, J.; Huang, Y.; Cheng, M.-M.; Liu, L.; Zhu, F.; Xu, Z.; and Shao, L. 2020. Noisy-as-clean: Learning self-supervised denoising from corrupted image. *IEEE Transactions on Image Processing*, 29: 9316–9329.
- Zamir, S. W.; Arora, A.; Khan, S.; Hayat, M.; Khan, F. S.; and Yang, M.-H. 2022. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5728–5739.
- Zhang, K.; Zuo, W.; Chen, Y.; Meng, D.; and Zhang, L. 2017. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7): 3142–3155.
- Zhang, Z.; Wang, R.; Zhang, H.; Chen, Y.; and Zuo, W. 2022. Self-supervised learning for real-world super-resolution from dual zoomed observations. In *European Conference on Computer Vision*, 610–627. Springer.
- Zhou, Y.; Jiao, J.; Huang, H.; Wang, Y.; Wang, J.; Shi, H.; and Huang, T. 2020. When awgn-based denoiser meets real noises. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 13074–13081.
- Zou, Y.; Yan, C.; and Fu, Y. 2023. Iterative denoiser and noise estimator for self-supervised image denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13265–13274.