

Uformer: A General U-Shaped Transformer for Image Restoration

Zhendong Wang^{1*} Xiaodong Cun^{2†} Jianmin Bao Wengang Zhou¹ Jianzhuang Liu³ Houqiang Li¹

¹University of Science and Technology of China

²University of Macau ³Huawei Noah's Ark Lab

Abstract

In this paper, we present Uformer, an effective and efficient Transformer-based architecture for image restoration, in which we build a hierarchical encoder-decoder network using the Transformer block. In Uformer, there are two core designs. First, we introduce a novel locally-enhanced window (LeWin) Transformer block, which performs non-overlapping window-based self-attention instead of global self-attention. It significantly reduces the computational complexity on high resolution feature map while capturing local context. Second, we propose a learnable multi-scale restoration modulator in the form of a multi-scale spatial bias to adjust features in multiple layers of the Uformer decoder. Our modulator demonstrates superior capability for restoring details for various image restoration tasks while introducing marginal extra parameters and computational cost. Powered by these two designs, Uformer enjoys a high capability for capturing both local and global dependencies for image restoration. To evaluate our approach, extensive experiments are conducted on several image restoration tasks, including image denoising, motion deblurring, defocus deblurring and deraining. Without bells and whistles, our Uformer achieves superior or comparable performance compared with the state-of-the-art algorithms. The code and models are available at <https://github.com/ZhendongWang6/Uformer>.

1. Introduction

With the rapid development of consumer and industry cameras and smartphones, the requirements of removing undesired degradation (e.g., noise, blur, rain, and so on) in images are constantly growing. Recovering genuine images from their degraded versions, i.e., image restoration, is a classic task in computer vision. Recent state-of-the-art methods [9, 47, 74, 76, 77] are mostly ConvNets-based, which achieve impressive results but show a limitation in capturing

*Part of this work was done during an internship in Huawei Noah's Ark Lab.

†Corresponding author.

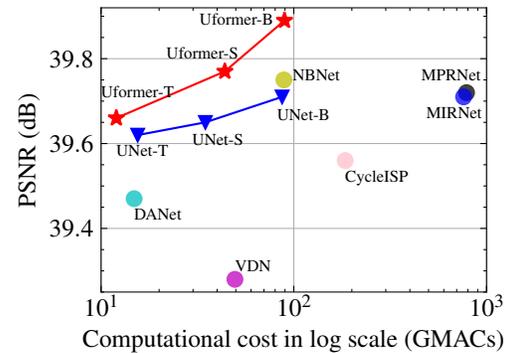


Figure 1. PSNR vs. computational cost on the SIDD dataset [1].

long-range dependencies. To address this problem, several recent works [31, 39, 83] start to employ single or few self-attention layers in low resolution feature maps due to the self-attention computational complexity being quadratic to the feature map size.

In this paper, we aim to leverage the capability of self-attention in feature maps at multi-scale resolutions to recover more image details. To this end, we present Uformer, an effective and efficient Transformer-based structure for image restoration. Uformer is built upon an elegant architecture UNet [49], where we modify the convolution layers to Transformer blocks while keeping the same overall hierarchical encoder-decoder structure and the skip-connections.

We propose two core designs to make Uformer suitable for image restoration tasks. First, we propose the **Locally-enhanced Window (LeWin) Transformer block**, which is an efficient and effective basic component. The LeWin Transformer block performs non-overlapping window-based self-attention instead of global self-attention, which significantly reduces the computational complexity on high resolution feature maps. Since we build hierarchical feature maps and keep the window size unchanged, the window-based self-attention at low resolution is able to capture much more global dependencies. On the other hand, local context is essential for image restoration, we further introduce a depth-wise convolutional layer between two fully-connected layers of the feed-forward network in the Transformer block to better capture local context. We also notice that recent works [35, 72]

use the similar design for different tasks.

Second, we propose a learnable multi-scale restoration modulator to handle various image degradations. The modulator is formulated as a multi-scale spatial bias to adjust features in multiple layers of the Uformer decoder. Specifically, a learnable window-based tensor is added to features in each LeWin Transformer block to adapt the features for restoring more details. Benefiting from the simple operator and window-based mechanism, it can be flexibly applied for various image restoration tasks in different frameworks.

Based on the above two designs, without bells and whistles, *e.g.*, the multi-stage or multi-scale framework [77, 78] and the advanced loss function [28, 29], our simple U-shaped Transformer structure achieves state-of-the-art performance on multiple image restoration tasks. For denoising, Uformer outperforms the previous state-of-the-art method (NBNet [9]) by 0.14 dB and 0.09 dB on the SIDD [1] and DND [46] benchmarks, respectively. For the motion blur removal task, Uformer achieves the best (GoPro [45], RealBlur-R [48], and RealBlur-J [48]) or competitive (HIDE [52]) performance, displaying its strong capability of deblurring. Uformer also shows the potential on the defocus deblurring task [3] and outperforms the previous best model [54] by 1.04 dB. Also, on the SPAD dataset [61] for deraining, it obtains 47.84 dB on PSNR, an improvement of 3.74 dB over the previous state-of-the-art method [47]. We expect our work will encourage further research to explore Transformer-based architectures for image restoration.

Overall, we summarize the contributions of this paper as follows:

- We present Uformer, a general and superior U-shaped Transformer for various image restoration tasks. Uformer is built on the basic LeWin Transformer block that is both efficient and effective.
- We present an extra light-weight learnable multi-scale restoration modulator to adjust on multi-scale features. This simple design significantly improves the restoration quality.
- Extensive experiments show that Uformer establishes new state-of-the-arts on various datasets for image restoration tasks.

2. Related Work

Image Restoration Architectures. Image restoration aims to restore the clean image from its degraded version. A popular solution is to learn effective models using the U-shaped structures with skip-connection to capture multi-scale information hierarchically for various image restoration tasks, including image denoising [9, 74, 77], deblurring [3, 28, 29], and demoreing [38, 55]. Some image restoration methods are inspired by the key insight from the rapid devel-

opment of image classification [17, 21, 27]. For example, ResNet-based structure has been widely used for general image restoration [42, 83] as well as for specific tasks in image restoration such as super-resolution [37, 84] and image denoising [16, 81]. There are also DenseNet-based methods [40, 41, 84]. More CNN-based image restoration architectures can be found in the recent surveys [32, 57, 65] and the NTIRE Challenges [2].

Until recently, some works start to explore the attention mechanism to boost the performance. For example, squeeze-and-excitation networks [20] and non-local neural networks [63] inspire a branch of methods for different image restoration tasks, such as super-resolution [36, 82], deraining [33, 61], denoising [76, 77], deblurring [47, 77], deshadow [11, 12], and so on. Our Uformer also applies the hierarchical structure to build multi-scale features while using the newly introduced LeWin Transformer block as the basic building block.

Vision Transformers. Transformer [59] shows a significant performance in natural language processing (NLP). Different from the design of CNNs, Transformer-based network structures are naturally good at capturing long-range dependencies in the data by the global self-attention. The success of Transformer in the NLP domain also inspires the computer vision researchers. The pioneering work of ViT [15] directly trains a pure Transformer-based architecture on the medium-size (16×16) flattened patches. With large-scale data pre-training (*i.e.*, JFT-300M), ViT gets excellent results compared to state-of-the-art CNNs on image classification.

Since the introduction of ViT, many efforts have been made to reduce the quadratic computational cost of global self-attention for making Transformer more suitable for vision tasks. Some works [19, 62] focus on establishing a pyramid Transformer architecture similar to ConvNet-based structure. To overcome the quadratic complexity of original self-attention, self-attention is performed on local windows with the halo operation or window shift [43, 58] to help cross-window interaction, and get promising results. Rather than focusing on image classification, recent works [10, 14, 23, 70, 85] propose a brunch of Transformer-based backbones for more general high-level vision tasks.

Besides high-level discriminative tasks, there are also some Transformer-based works [25, 68, 86] for generative tasks. While there are enormous explorations in the vision area, introducing Transformer to low-level vision still lacks exploration. Early work [69] makes use of self-attention mechanism to learn texture for super-resolution. As for image restoration tasks, IPT [8] first applies standard Transformer blocks within a multi-task learning framework. However, IPT relies on pretraining on a large-scale synthesized dataset and multi-task learning. In contrast, we design a general U-shaped Transformer-based structure, which proves to be efficient and effective for image restoration.

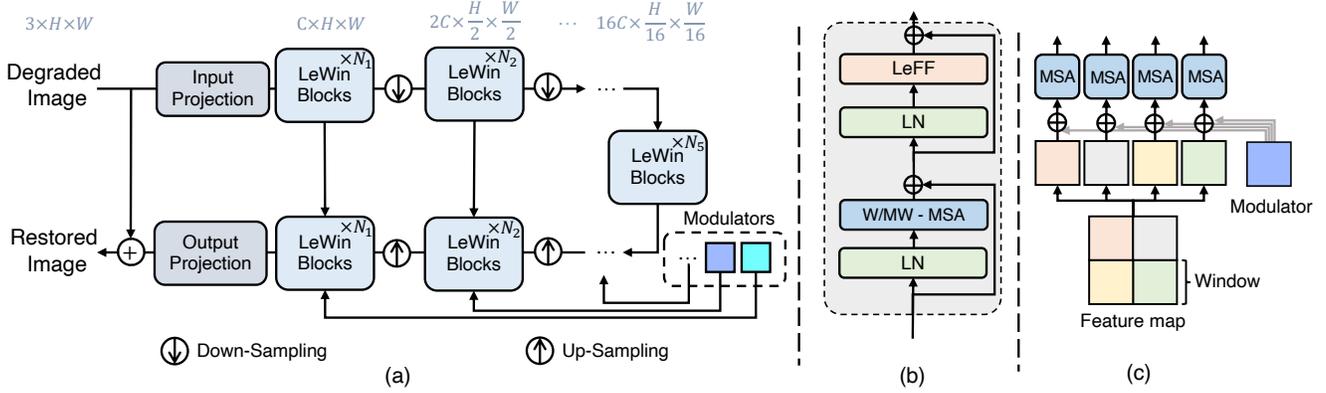


Figure 2. (a) Overview of Uformer. (b) LeWin Transformer block. (c) Illustration of how the modulators modulate the W-MSAs in each LeWin Transformer block which is named MW-MSA in (b).

3. Method

In this section, we first describe the overall pipeline and the hierarchical structure of Uformer for image restoration. Then, we provide the details of the *LeWin Transformer block* which is the basic component of Uformer. After that, we present the *multi-scale restoration modulator*.

3.1. Overall Pipeline

As shown in Figure 2(a), the overall structure of the proposed Uformer is a U-shaped hierarchical network with skip-connections between the encoder and the decoder. To be specific, given a degraded image $\mathbf{I} \in \mathbb{R}^{3 \times H \times W}$, Uformer firstly applies a 3×3 convolutional layer with LeakyReLU to extract low-level features $\mathbf{X}_0 \in \mathbb{R}^{C \times H \times W}$. Next, following the design of the U-shaped structures [24, 49], the feature maps \mathbf{X}_0 are passed through K encoder stages. Each stage contains a stack of the proposed LeWin Transformer blocks and one down-sampling layer. The LeWin Transformer block takes advantage of the self-attention mechanism for capturing long-range dependencies, and also cuts the computational cost due to the usage of self-attention through non-overlapping windows on the feature maps. In the down-sampling layer, we first reshape the flattened features into 2D spatial feature maps, and then down-sample the maps, double the channels using 4×4 convolution with stride 2. For example, given the input feature maps $\mathbf{X}_0 \in \mathbb{R}^{C \times H \times W}$, the l -th stage of the encoder produces the feature maps $\mathbf{X}_l \in \mathbb{R}^{2^l C \times \frac{H}{2^l} \times \frac{W}{2^l}}$.

Then, a bottleneck stage with a stack of LeWin Transformer blocks is added at the end of the encoder. In this stage, thanks to the hierarchical structure, the Transformer blocks capture longer (even global when the window size equals the feature map size) dependencies.

For feature reconstruction, the proposed decoder also contains K stages. Each consists of an up-sampling layer and a stack of LeWin Transformer blocks similar to the encoder.

We use 2×2 transposed convolution with stride 2 for the up-sampling. This layer reduces half of the feature channels and doubles the size of the feature maps. After that, the features input to the LeWin Transformer blocks are concatenation of the up-sampled features and the corresponding features from the encoder through skip-connection. Next, the LeWin Transformer blocks are utilized to learn to restore the image. After the K decoder stages, we reshape the flattened features to 2D feature maps and apply a 3×3 convolution layer to obtain a residual image $\mathbf{R} \in \mathbb{R}^{3 \times H \times W}$. Finally, the restored image is obtained by $\mathbf{I}' = \mathbf{I} + \mathbf{R}$. We train Uformer using the Charbonnier loss [7, 76]:

$$\ell(\mathbf{I}', \hat{\mathbf{I}}) = \sqrt{\|\mathbf{I}' - \hat{\mathbf{I}}\|^2 + \epsilon^2}, \quad (1)$$

where $\hat{\mathbf{I}}$ is the ground-truth image, and $\epsilon = 10^{-3}$ is a constant in all the experiments.

3.2. LeWin Transformer Block

There are two main challenges to apply Transformer for image restoration. First, the standard Transformer architecture [15, 59] computes self-attention globally between all tokens, which contributes to the quadratic computation cost with respect to the number of tokens. It is unsuitable to apply global self-attention on high-resolution feature maps. Second, the local context information is essential for image restoration tasks since the neighborhood of a degraded pixel can be leveraged to restore its clean version, but previous works [35, 66] suggest that Transformer shows a limitation in capturing local dependencies.

To address the above mentioned two issues, we propose a **Locally-enhanced Window (LeWin) Transformer block**, as shown in Figure 2(b), which benefits from the self-attention in Transformer to capture long-range dependencies, and also involves the convolution operator into Transformer to capture useful local context. Specifically, given the features

at the $(l-1)$ -th block \mathbf{X}_{l-1} , we build the block with two core designs: (1) non-overlapping Window-based Multi-head Self-Attention (W-MSA) and (2) Locally-enhanced Feed-Forward Network (LeFF). The computation of a LeWin Transformer block is represented as:

$$\begin{aligned} \mathbf{X}'_l &= \text{W-MSA}(\text{LN}(\mathbf{X}_{l-1})) + \mathbf{X}_{l-1}, \\ \mathbf{X}_l &= \text{LeFF}(\text{LN}(\mathbf{X}'_l)) + \mathbf{X}'_l, \end{aligned} \quad (2)$$

where \mathbf{X}'_l and \mathbf{X}_l are the outputs of the W-MSA module and LeFF module, respectively. LN represents the layer normalization [5]. In the following, we elaborate W-MSA and LeFF separately.

Window-based Multi-head Self-Attention (W-MSA). Instead of using global self-attention like the vanilla Transformer, we perform the self-attention within non-overlapping local windows, which reduces the computational cost significantly. Given the 2D feature maps $\mathbf{X} \in \mathbb{R}^{C \times H \times W}$ with H and W being the height and width of the maps, we split \mathbf{X} into non-overlapping windows with the window size of $M \times M$, and then get the flattened and transposed features $\mathbf{X}^i \in \mathbb{R}^{M^2 \times C}$ from each window i . Next, we perform self-attention on the flattened features in each window.

Suppose the head number is k and the head dimension is $d_k = C/k$. Then computing the k -th head self-attention in the non-overlapping windows can be formulated as follows,

$$\begin{aligned} \mathbf{X} &= \{\mathbf{X}^1, \mathbf{X}^2, \dots, \mathbf{X}^N\}, \quad N = HW/M^2, \\ \mathbf{Y}_k^i &= \text{Attention}(\mathbf{X}^i \mathbf{W}_k^Q, \mathbf{X}^i \mathbf{W}_k^K, \mathbf{X}^i \mathbf{W}_k^V), \quad i = 1, \dots, N, \\ \hat{\mathbf{X}}_k &= \{\mathbf{Y}_k^1, \mathbf{Y}_k^2, \dots, \mathbf{Y}_k^N\}, \end{aligned} \quad (3)$$

where $\mathbf{W}_k^Q, \mathbf{W}_k^K, \mathbf{W}_k^V \in \mathbb{R}^{C \times d_k}$ represent the projection matrices of the queries, keys, and values for the k -th head, respectively. $\hat{\mathbf{X}}_k$ is the output of the k -th head. Then the outputs for all heads $\{1, 2, \dots, k\}$ are concatenated and then linearly projected to get the final result. Inspired by previous works [43, 51], we also apply the relative position encoding into the attention module, so the attention calculation can be formulated as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{SoftMax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}} + \mathbf{B}\right)\mathbf{V}, \quad (4)$$

where \mathbf{B} is the relative position bias, whose values are taken from $\hat{\mathbf{B}} \in \mathbb{R}^{(2M-1) \times (2M-1)}$ with learnable parameters [43, 51].

Window-based self-attention can significantly reduce the computational cost compared with global self-attention. Given the feature maps $\mathbf{X} \in \mathbb{R}^{C \times H \times W}$, the computational complexity drops from $O(H^2W^2C)$ to $O(\frac{HW}{M^2}M^4C) = O(M^2HWC)$. Since we design Uformer as a hierarchical architecture, our window-based self-attention at low resolution feature maps works on larger receptive fields and is sufficient to learn long-range dependencies. We also try the

shifted-window strategy [43] in the even LeWin Transformer block of each stage in our framework, which gives only slightly better results.

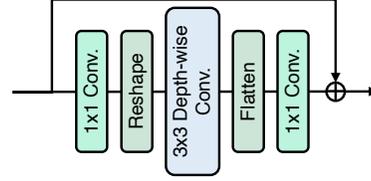


Figure 3. Locally-enhanced feed-forward network.

Locally-enhanced Feed-Forward Network (LeFF). As pointed out by previous works [66, 72], the Feed-Forward Network (FFN) in the standard Transformer suffers limited capability to leverage local context. Actually, neighboring pixels are crucial references for image restoration [6, 22]. To overcome this issue, we add a depth-wise convolutional block to the FFN in our Transformer-based structure following the recent works [35, 50, 72]. As shown in Figure 3, we first apply a linear projection layer to each token to increase its feature dimension. Next, we reshape the tokens to 2D feature maps, and use a 3×3 depth-wise convolution to capture local information. Then we flatten back the features to tokens and shrink the channels via another linear layer to match the dimension of the input channels. We use GELU [18] as the activation function after each linear/convolution layer.

3.3. Multi-Scale Restoration Modulator

Different types of image degradation (e.g. blur, noise, rain, etc.) have their own distinctive perturbed patterns to be handled or restored. To further boost the capability of Uformer for approaching various perturbations, we propose a light-weight multi-scale restoration modulator to calibrate the features and encourage more details recovered.

As shown in Figure 2(a) and 2(c), the multi-scale restoration modulator applies multiple modulators in the Uformer decoder. Specially in each LeWin Transformer block, a modulator is formulated as a learnable tensor with a shape of $M \times M \times C$, where M is the window size and C is the channel dimension of current feature map. Each modulator is simply served as a shared bias term that is added into all non-overlapping windows before self-attention module. Due to this light-weight addition operation and window-sized shape, the multi-scale restoration modulator introduces marginal extra parameters and computational cost.

We prove the effectiveness of the multi-scale restoration modulator on two typical image restoration tasks: image deblurring and image denoising. The visualization comparisons are presented in Figure 4. We observe that adding the multi-scale restoration modulator makes more motion blur/noising patterns removed and yields a much cleaner image. These results show that our multi-scale restoration modulator truly helps to recover restoration details with little

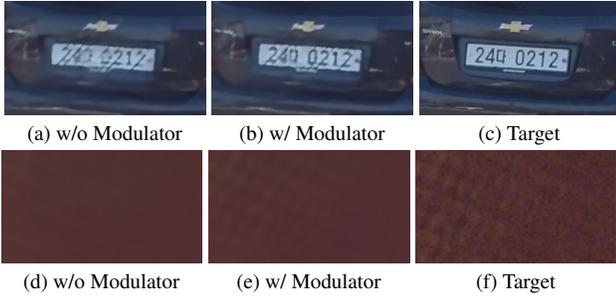


Figure 4. Effect of the multi-scale restoration modulator on image deblurring (top samples from GoPro [45]) and denoising (bottom samples from SIDD [1]). Compared with (a), Uformer w/ Modulator (b) can remove much more blur and recover the numbers accurately. Compared with (d), the image restored by Uformer w/ Modulator (e) is closer to the target with more details.

computation cost. One possible explanation is that adding modulators at each stage of the decoder enables a flexible adjustment of the feature maps that boosts the performance for restoring details. This is consistent with the previous work StyleGAN [26] using a multi-scale noise term adding to the convolution features, which realizes stochastic variation for generating photo-realistic images.

4. Experiments

In this section, we first discuss the experimental setup. After that, we verify the effectiveness and efficiency of Uformer on various image restoration tasks on eight datasets. Finally, we perform comprehensive ablation studies to evaluate each component of our proposed Uformer.

4.1. Experimental Setup

Basic settings. Following the common training strategy of Transformer [59], we train our framework using the AdamW optimizer [44] with the momentum terms of (0.9, 0.999) and the weight decay of 0.02. We randomly augment the training samples using the horizontal flipping and rotate the images by 90° , 180° , or 270° . We use the cosine decay strategy to decrease the learning rate to $1e-6$ with the initial learning rate $2e-4$. We set the window size to 8×8 in all LeWin Transformer blocks. The number of Uformer encoder/decoder stages K equals 4 by default. And the dimension of each head in Transformer block d_k equals C . More dataset-specific experimental settings can be found in the supplementary materials.

Architecture variants. For a concise description, we introduce three Uformer variants in our experiments, Uformer-T (Tiny), Uformer-S (Small), and Uformer-B (Base) by setting different Transformer feature channels C and the numbers of the Transformer blocks in each encoder and decoder stages. The details are listed as follows:

- Uformer-T: $C = 16$, depths of Encoder = {2, 2, 2, 2},

- Uformer-S: $C = 32$, depths of Encoder = {2, 2, 2, 2},
- Uformer-B: $C = 32$, depths of Encoder = {1, 2, 8, 8},

and the depths of Decoder are mirrored depths of Encoder.

Evaluation metrics. We adopt the commonly-used PSNR and SSIM [64] metrics to evaluate the restoration performance. These metrics are calculated in the RGB color space except for deraining where we evaluate the PSNR and SSIM on the Y channel in the YCbCr color space, following the previous work [60].

Method	SIDD [1]		DND [46]	
	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
BM3D [13]	25.65	0.685	34.51	0.851
RIDNet [4]	38.71	0.914	39.26	0.953
VDN [73]	39.28	0.909	39.38	0.952
DANet [74]	39.47	0.918	39.59	<u>0.955</u>
CycleISP [75]	39.52	0.957	39.56	0.956
MIRNet [76]	39.72	<u>0.959</u>	39.88	0.956
MPRNet [77]	39.71	0.958	39.80	0.954
NBNet [9]	<u>39.75</u>	<u>0.959</u>	<u>39.89</u>	<u>0.955</u>
Uformer-B	39.89	0.960	40.04	0.956

Table 1. Denoising results on the SIDD [1] and DND [46] datasets.

4.2. Real Noise Removal

Table 1 reports the results of real noise removal on the SIDD [1] and DND [46] datasets. We compare Uformer with 8 state-of-the-art denoising methods, including the feature-based BM3D [13] and seven learning-based methods: RIDNet [4], VDN [73], CycleISP [75], NBNet [9], DANet [74], MIRNet [76], and MPRNet [77]. Our Uformer-B achieves 39.89 dB on PSNR, surpassing all the other methods by at least 0.14 dB. As for the DND dataset, we follow the common evaluation strategy and test our model trained on SIDD via the online server testing. Uformer outperforms the previous state-of-the-art method NBNet [9] by 0.09 dB. To verify whether the gains benefit from more computation cost, we present the results of PSNR vs. computational cost in Figure 1. We notice that our Uformer-T can achieve a better performance than most models but with the least computation cost, which demonstrates the efficiency and effectiveness of Uformer. We also show the qualitative results on the SIDD and DND datasets in Figure 5, in which Uformer can not only successfully remove the noise but also keep the texture details.

Further, we give quantitatively and qualitatively comparisons with IPT [8] for denoising in Table 3 and Figure 6. We observe that Uformer outperforms IPT significantly on SIDD and DND with much fewer GMACs and parameters, while not requiring large scale pre-training.

4.3. Motion Blur Removal

For motion blur removal, Uformer also shows state-of-the-art performance. We follow the previous method [77] to

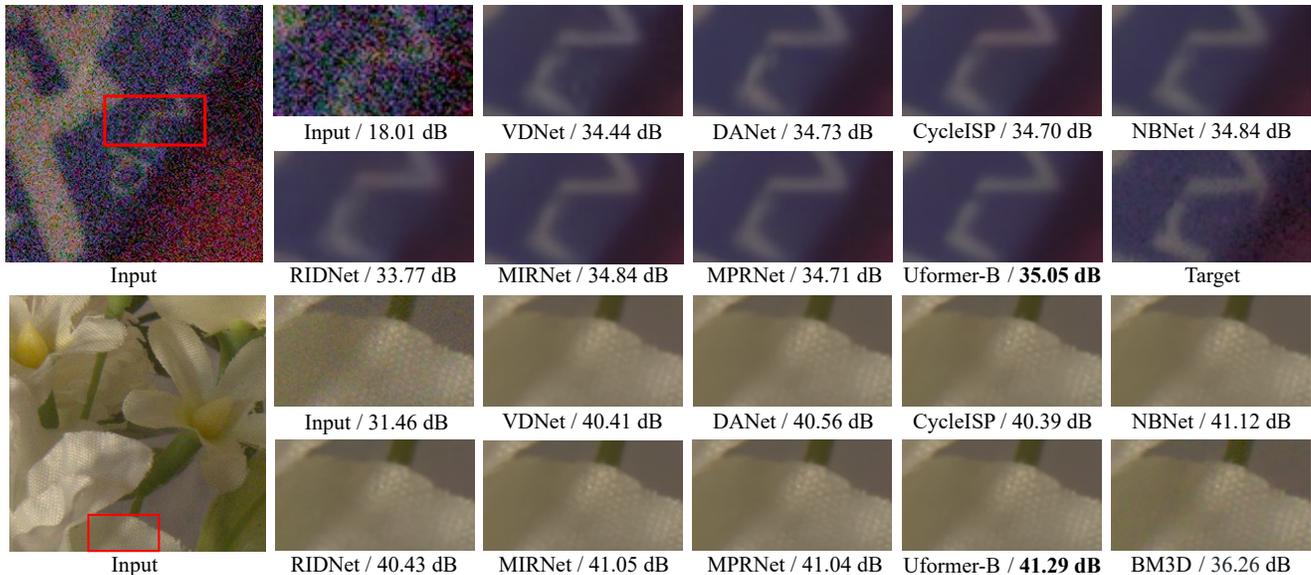


Figure 5. Visual comparisons with state-of-the-art methods on real noise removal. The top sample comes from SIDD [1] while the bottom one is from DND [46].

Method	GoPro [45]		HIDE [52]		RealBlur-R [48]		RealBlur-J [48]	
	PSNR \uparrow	SSIM \uparrow						
Nah <i>et al.</i> [45]	29.08	0.914	25.73	0.874	32.51	0.841	27.87	0.827
DeblurGAN [28]	28.70	0.858	24.51	0.871	33.79	0.903	27.97	0.834
Xu <i>et al.</i> [67]	21.00	0.741	-	-	34.46	0.937	27.14	0.830
DeblurGAN-v2 [29]	29.55	0.934	26.61	0.875	35.26	0.944	28.70	0.866
DBGAN [80]	31.10	0.942	28.94	0.915	-	-	-	-
SPAIR [47]	32.06	0.953	30.29	0.931	-	-	<u>28.81</u>	<u>0.875</u>
\dagger Zhang <i>et al.</i> [79]	29.19	0.931	-	-	35.48	0.947	27.80	0.847
\dagger SRN [56]	30.26	0.934	28.36	0.915	35.66	0.947	28.56	0.867
\dagger DMPHN [78]	31.20	0.940	29.09	0.924	35.70	0.948	28.42	0.860
\dagger MPRNet [77]	32.66	0.959	30.96	0.939	35.99	0.952	28.70	0.873
Uformer-B	33.06	0.967	<u>30.90</u>	0.953	36.19	0.956	29.09	0.886

Table 2. Results on motion deblurring. Following previous works [28, 29, 77], our Uformer is only trained on the GoPro dataset [45]. Then we apply our GoPro trained model directly on the HIDE dataset [52] and the RealBlur dataset [48] to evaluate the generalization on real scenes. \dagger denotes recurrent/multi-stage designs for better performance.

	GMACs	# Param	SIDD [1]		DND [46]	
			PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
IPT [8]	\sim 380G	115.31M	39.10	0.954	39.62	0.952
Uformer-B	89.46G	50.88M	39.89	0.960	40.04	0.956

Table 3. Comparisons of IPT [8] and Uformer for denoising.



Figure 6. Visual comparisons of IPT [8] and Uformer for denoising.

train Uformer on the GoPro dataset and test it on the four datasets: two synthesized datasets (HIDE [52] and the test set of GoPro [45]), and two real-world datasets (RealBlur-R

and -J [48]). We compare Uformer with ten state-of-the-art methods: Nah *et al.* [45], DeblurGAN [28], Xu *et al.* [67], DeblurGAN-v2 [29], DBGAN [80], SPAIR [47], Zhang *et al.* [79], SRN [56], DMPHN [78], and MPRNet [77]. The results are reported in Table 2. For synthetic deblurring, Uformer gets significant better performance on GoPro than previous state-of-the-art method [77] and shows a comparable result on the HIDE dataset. As for real-world deblurring, the causes of blur are complicated so the task is usually more challenging. Our Uformer outperforms other methods by at least 0.23 dB and 0.36 dB on RealBlur-R and RealBlur-J, respectively, showing a strong generalization ability. Besides,

we show some visual results in Figure 7. Compared with other methods, the images restored by Uformer are more clear and closer to ground truth.

4.4. Defocus Blur Removal

We perform defocus blur removal on the DPD dataset [3]. Table 4 and Figure 8 report the quantitative and qualitative results, respectively. Uformer achieves a better performance (1.04 dB, 1.15 dB, 1.44 dB, and 1.87 dB) over previous state-of-the-art methods KPAC [54], DPDNet [3], JNB [53], and DMENet [30], respectively. From the visualization results, we observe that the images recovered by Uformer are sharper and closer to the ground-truth images.

	DMENet [30]	JNB [53]	DPDNet [3]	KPAC [54]	Uformer-B
PSNR \uparrow	23.41	23.84	25.13	25.24	26.28
SSIM \uparrow	0.714	0.715	0.786	0.842	0.891

Table 4. Results on the DPD dataset [3] for defocus blur removal.

4.5. Real Rain Removal

We conduct the deraining experiments on SPAD [61] and compare with 6 deraining methods: GMM [34], RESCAN [33], SPANet [61], JORDER-E [71], RCDNet [60], and SPAIR [47]. As shown in Table 5, Uformer presents a significantly better performance, achieving 3.74 dB improvement over the previous best work [47]. This indicates the strong capability of Uformer for deraining on this real derain dataset. We also provide the visual results in Figure 8 where Uformer can remove the rain more successfully while introducing fewer artifacts.

	GMM [34]	RESCAN [33]	SPANet [61]	JORDER-E [71]	RCDNet [60]	SPAIR [47]	Uformer-B
PSNR \uparrow	34.30	38.11	40.24	40.78	41.47	44.10	47.84
SSIM \uparrow	0.9428	0.9707	0.9811	0.9811	0.9834	0.9872	0.9925

Table 5. Results on the SPAD dataset [61] for real rain removal.

4.6. Ablation Study

In this section, we analyze the effect of each component of Uformer in detail. The evaluations are conducted on image denoising (SIDD [1]), deblurring (GoPro [45], RealBlur [48]), and deraining (SPAD [61]) using different variants. The ablation results are reported in Tables 6–8. **Transformer vs. convolution.** We replace all the LeWin Transformer blocks in Uformer with the convolution-based ResBlocks [9], resulting in the so-called "UNet", while keeping all others unchanged. Similar to the Uformer variants, we design UNet-T/-S/-B:

- UNet-T: $C = 32$, depths of Encoder = {2, 2, 2, 2},
- UNet-S: $C = 48$, depths of Encoder = {2, 2, 2, 2},

	GMACs	# Param	PSNR \uparrow
UNet-T	15.49G	9.50M	39.62
UNet-S	34.76G	21.38M	39.65
UNet-B	86.97G	53.58M	39.71
ViT	8.83G	14.86M	38.51
Uformer-T	12.00G	5.23M	39.66
Uformer-S	43.86G	20.63M	39.77
Uformer-B	89.46G	50.88M	39.89

Table 6. Comparison of different network architectures for denoising on the SIDD dataset [1].

	W-MSA	FFN	GMACs	# Param	PSNR \uparrow
Uformer-S (SIDD [1])	-	-	43.00G	20.47M	39.74
	\checkmark	-	43.64G	20.59M	39.72
	-	\checkmark	43.86G	20.63M	39.77
Uformer-B (RealBlur-R/-J [48])	-	-	88.31G	50.45M	36.15/28.99
	\checkmark	\checkmark	90.31G	51.20M	36.19/28.85
	-	\checkmark	89.46G	50.88M	36.22/29.06

Table 7. Effect of enhancing locality in different modules.

- UNet-B: $C = 76$, depths of Encoder = {2, 2, 2, 2},

and the depths of Decoder are mirrored depths of Encoder.

Table 6 reports the comparison results. We observe that Uformer-T achieves 39.66 dB and outperforms UNet-T by 0.04 dB with fewer parameters and less computation. Uformer-S achieves 39.77 dB and outperforms UNet-S by 0.12 dB with fewer parameters and a slightly higher computation cost. And Uformer-B achieves 39.89 dB which outperforms UNet-B by 0.18 dB. This study indicates the effectiveness of the proposed LeWin Transformer block, compared with the original convolutional block.

Hierarchical structure vs. single scale. We further build a ViT-based architecture which only contains a single scale of the feature maps for image denoising. This architecture employs a head of two convolution layers for extracting features from the input image and also a tail of two convolution layers for the output. 12 standard Transformer blocks are used between the head and the tail. We train the ViT with the hidden dimension of 256 on patch size 16×16 . The results are presented in Table 6. We observe that the vanilla ViT structure gets an unsatisfactory result compared with UNet, while our Uformer significantly outperforms both the ViT-based and UNet architectures, which demonstrates the effectiveness of hierarchical structure for image restoration.

Where to enhance locality? Table 7 compares the results of no locality enhancement and enhancing locality in the self-attention calculation [66] or the feed-forward network based on Uformer-S and Uformer-B. We observe that introducing locality into the feed-forward network yields 0.03 dB (SIDD), 0.07 dB (RealBlur-R)/0.07 dB (RealBlur-J) over the baseline (no locality enhancement), while introducing locality into the self-attention yields -0.02 dB (SIDD). Further, we combine introducing locality into the feed-

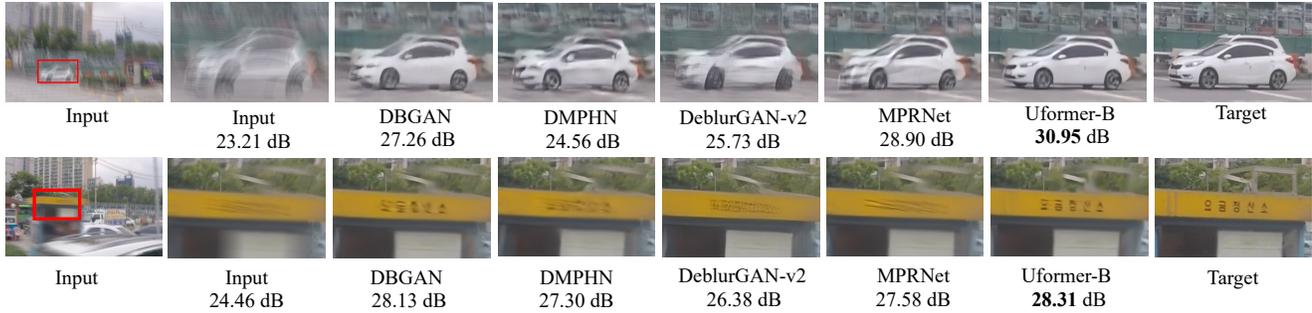


Figure 7. Visual comparisons with state-of-the-art methods on the GoPro dataset [45] for motion blur removal.

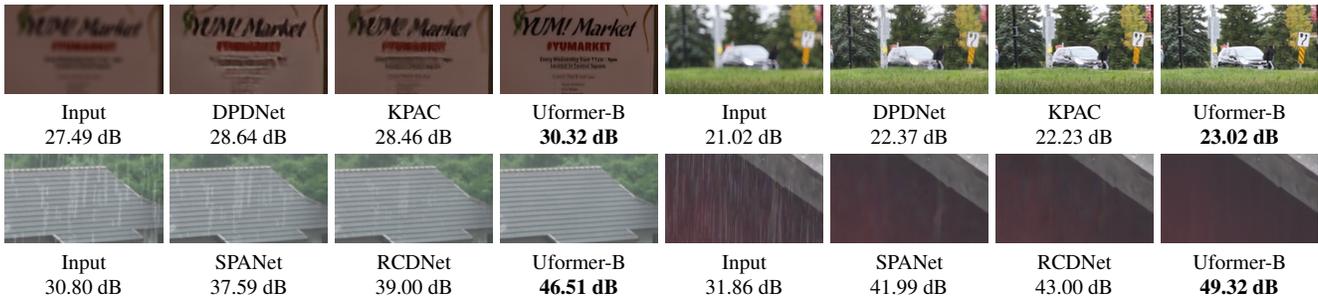


Figure 8. Top row: Visual comparisons with state-of-the-art methods on the DPD dataset [3] for defocus blur removal. Bottom row: Visual comparisons with state-of-the-art methods on the SPAD dataset [61] for real rain removal.

forward network and introducing into the self-attention. The results on RealBlur-R/-J also drop from 36.22 dB/29.06 dB to 36.19 dB/28.85 dB, indicating that compared to involving locality into self-attention, introducing locality into the feed-forward network is more suitable for image restoration tasks.

	GoPro [45]	SIDD [1]	SPAD [61]
	Uformer-T	Uformer-B	Uformer-B
Modulator	- ✓	- ✓	- ✓
PSNR ↑	29.11 29.57	39.86 39.89	47.43 47.84

Table 8. Effect of the multi-scale restoration modulator.

Effect of the multi-scale restoration modulator. In Table 8, to verify the effect of the modulator, we conduct experiments on GoPro for image deblurring, SIDD for image denoising, and SPAD for deraining. For deblurring, we observe that w/ modulator can bring a performance improvement of 0.46 dB, which reveals the effectiveness of the modulator for deblurring. We also compare the results of Uformer-B with/without the modulator on SIDD and SPAD, and the comparisons indicate that the proposed modulator introduces 0.03 dB improvement (SIDD)/0.41 dB improvement (SPAD). In Figure 4, we have provided visual comparisons of Uformer w/ and wo/ the modulator. This study validates the proposed modulator can bring extra ability of restoring more details. More comparisons about the modula-

tor can be found in supplementary materials.

5. Discussion and Conclusion

In this paper, we have presented an alternative architecture Uformer for image restoration tasks by introducing the Transformer block. In contrast to existing ConvNet-based structures, our Uformer builds upon the main component LeWin Transformer block, which can not only handle local context but also capture long-range dependencies efficiently. To handle various image restoration degradation and enhance restoration quality, we propose a learnable multi-scale restoration modulator inserted into the Uformer decoder. Extensive experiments demonstrate that Uformer achieves state-of-the-art performance on several tasks, including denoising, motion deblurring, defocus deblurring, and deraining. Uformer also surpasses the UNet family by a large margin with less computation cost and fewer model parameters.

Acknowledgement. This work was supported in part by the National Natural Science Foundation of China under Contract 61836011 and 62021001, and in part by the Youth Innovation Promotion Association CAS under Grant 2018497. It was also supported by the GPU cluster built by MCC Lab of Information Science and Technology Institution, USTC. We also thank MindSpore¹ for the partial support of this work, which is a new deep learning computing framework.

¹<https://www.mindspore.cn/>

References

- [1] Abdelrahman Abdelhamed, Stephen Lin, and Michael S. Brown. A High-Quality Denoising Dataset for Smartphone Cameras. In *CVPR*, 2018. 1, 2, 5, 6, 7, 8
- [2] Abdelrahman Abdelhamed, Radu Timofte, and Michael S. Brown. NTIRE 2019 Challenge on Real Image Denoising: Methods and Results. In *CVPR Workshop*, 2019. 2
- [3] Abdullah Abuolaim and Michael S Brown. Defocus Deblurring Using Dual-Pixel Data. In *ECCV*. Springer, 2020. 2, 7, 8
- [4] Saeed Anwar and Nick Barnes. Real Image Denoising with Feature Attention. In *ICCV*, 2019. 5
- [5] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer Normalization. *arXiv preprint arXiv:1607.06450*, 2016. 4
- [6] A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. In *CVPR*, 2005. 4
- [7] P. Charbonnier, L. Blanc-Feraud, G. Aubert, and M. Barlaud. Two Deterministic Half-Quadratic Regularization Algorithms for Computed Imaging. In *ICIP*, 1994. 3
- [8] Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. Pre-Trained Image Processing Transformer. In *CVPR*, 2021. 2, 5, 6
- [9] Shen Cheng, Yuzhi Wang, Haibin Huang, Donghao Liu, Haoqiang Fan, and Shuaicheng Liu. NBNNet: Noise Basis Learning for Image Denoising with Subspace Projection. In *CVPR*, 2021. 1, 2, 5, 7
- [10] Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting the Design of Spatial Attention in Vision Transformers. In *NeurIPS*, 2021. 2
- [11] Xiaodong Cun and Chi-Man Pun. Improving the harmony of the composite image by spatial-separated attention module. *TIP*, 29:4759–4771, 2020. 2
- [12] Xiaodong Cun, Chi-Man Pun, and Cheng Shi. Towards ghost-free shadow removal via dual hierarchical aggregation network and shadow matting GAN. In *AAAI*, 2020. 2
- [13] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image Denoising by Sparse 3-D Transform-Domain Collaborative Filtering. *TIP*, 16(8):2080–2095, 2007. 5
- [14] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. CSWin Transformer: A General Vision Transformer Backbone with Cross-Shaped Windows. *arXiv preprint arXiv:2107.00652*, 2021. 2
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*, 2020. 2, 3
- [16] Shuhang Gu, Yawei Li, Luc Van Gool, and Radu Timofte. Self-Guided Network for Fast Image Denoising. In *ICCV*, 2019. 2
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016. 2
- [18] Dan Hendrycks and Kevin Gimpel. Gaussian Error Linear Units (GELUs). *arXiv preprint arXiv:1606.08415*, 2016. 4
- [19] Byeongho Heo, Sangdoon Yun, Dongyoon Han, Sanghyuk Chun, Junsuk Choe, and Seong Joon Oh. Rethinking Spatial Dimensions of Vision Transformers. In *ICCV*, 2021. 2
- [20] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-Excitation Networks. In *CVPR*, 2018. 2
- [21] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q Weinberger. Densely Connected Convolutional Networks. In *CVPR*, 2017. 2
- [22] Tao Huang, Songjiang Li, Xu Jia, Huchuan Lu, and Jianzhuang Liu. Neighbor2Neighbor: Self-Supervised Denoising from Single Noisy Images. In *CVPR*, 2021. 4
- [23] Zilong Huang, Youcheng Ben, Guozhong Luo, Pei Cheng, Gang Yu, and Bin Fu. Shuffle Transformer: Rethinking Spatial Shuffle for Vision Transformer. *arXiv preprint arXiv:2106.03650*, 2021. 2
- [24] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-Image Translation with Conditional Adversarial Networks. In *CVPR*, 2017. 3
- [25] Yifan Jiang, Shiyu Chang, and Zhangyang Wang. TransGAN: Two Pure Transformers Can Make One Strong GAN, and That Can Scale Up. *arXiv preprint arXiv:2102.07074*, 2021. 2
- [26] Tero Karras, Samuli Laine, and Timo Aila. A Style-Based Generator Architecture for Generative Adversarial Networks. In *CVPR*, 2019. 5
- [27] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *NeurIPS*, 2012. 2
- [28] Orest Kupyn, Volodymyr Budzan, Mykola Mykhailych, Dmytro Mishkin, and Jiri Matas. DeblurGAN: Blind Motion Deblurring Using Conditional Adversarial Networks. *ArXiv e-prints*, 2017. 2, 6
- [29] Orest Kupyn, Tetiana Martyniuk, Junru Wu, and Zhangyang Wang. DeblurGAN-v2: Deblurring (Orders-of-Magnitude) Faster and Better. In *ICCV*, 2019. 2, 6
- [30] Junyong Lee, Sungkil Lee, Sunghyun Cho, and Seungyong Lee. Deep Defocus Map Estimation using Domain Adaptation. In *CVPR*, 2019. 7
- [31] Guanbin Li, Xiang He, Wei Zhang, Huiyou Chang, Le Dong, and Liang Lin. Non-locally Enhanced Encoder-Decoder Network for Single Image De-raining. In *ACMMM*, 2018. 1
- [32] Siyuan Li, Iago Breno Araujo, Wenqi Ren, Zhangyang Wang, Eric K Tokuda, Roberto Hirata Junior, Roberto Cesar-Junior, Jiawan Zhang, Xiaojie Guo, and Xiaochun Cao. Single image deraining: A comprehensive benchmark analysis. In *CVPR*, 2019. 2
- [33] Xia Li, Jianlong Wu, Zhouchen Lin, Hong Liu, and Hongbin Zha. Recurrent Squeeze-and-Excitation Context Aggregation Net for Single Image Deraining. In *ECCV*, 2018. 2, 7
- [34] Yu Li, Robby T. Tan, Xiaojie Guo, Jiangbo Lu, and Michael S. Brown. Rain Streak Removal Using Layer Priors. In *CVPR*, 2016. 7

- [35] Yawei Li, Kai Zhang, Jiezhong Cao, Radu Timofte, and Luc Van Gool. LocalViT: Bringing Locality to Vision Transformers. *arXiv preprint arXiv:2104.05707*, 2021. 1, 3, 4
- [36] Jingyun Liang, Jiezhong Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. SwinIR: Image Restoration Using Swin Transformer. In *ICCV Workshops*, 2021. 2
- [37] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced Deep Residual Networks for Single Image Super-Resolution. In *CVPR Workshop*, 2017. 2
- [38] Bolin Liu, Xiao Shu, and Xiaolin Wu. Demoiréing of Camera-Captured Screen Images Using Deep Convolutional Neural Network. *arXiv preprint arXiv:1804.03809*, 2018. 2
- [39] Ding Liu, Bihan Wen, Yuchen Fan, Chen Change Loy, and Thomas S Huang. Non-Local Recurrent Network for Image Restoration. In *NeurIPS*, 2018. 1
- [40] Lin Liu, Xu Jia, Jianzhuang Liu, and Qi Tian. Joint Demosaicing and Denoising With Self Guidance. In *CVPR*, 2020. 2
- [41] Lin Liu, Jianzhuang Liu, Shanxin Yuan, Gregory Slabaugh, Ales Leonardis, Wengang Zhou, and Qi Tian. Wavelet-Based Dual-Branch Network for Image Demoiréing. In *ECCV*, 2020. 2
- [42] Xing Liu, Masanori Suganuma, Zhun Sun, and Takayuki Okatani. Dual Residual Networks Leveraging the Potential of Paired Operations for Image Restoration. In *CVPR*, 2019. 2
- [43] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. *arXiv preprint arXiv:2103.14030*, 2021. 2, 4
- [44] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. *arXiv preprint arXiv:1711.05101*, 2017. 5
- [45] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep Multi-Scale Convolutional Neural Network for Dynamic Scene Deblurring. In *CVPR*, 2017. 2, 5, 6, 7, 8
- [46] Tobias Plotz and Stefan Roth. Benchmarking Denoising Algorithms with Real Photographs. In *CVPR*, 2017. 2, 5, 6
- [47] Kuldeep Purohit, Maitreya Suin, AN Rajagopalan, and Vishnu Naresh Boddeti. Spatially-Adaptive Image Restoration using Distortion-Guided Networks. In *ICCV*, 2021. 1, 2, 6, 7
- [48] Jaesung Rim, Haeyun Lee, Jucheol Won, and Sunghyun Cho. Real-World Blur Dataset for Learning and Benchmarking Deblurring Algorithms. In *ECCV*, 2020. 2, 6, 7
- [49] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *MICCAI*. Springer, 2015. 1, 3
- [50] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobileNetV2: Inverted Residuals and Linear Bottlenecks. In *CVPR*, 2018. 4
- [51] Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. Self-Attention with Relative Position Representations. *arXiv preprint arXiv:1803.02155*, 2018. 4
- [52] Ziyi Shen, Wenguan Wang, Jianbing Shen, Haibin Ling, Tingfa Xu, and Ling Shao. Human-Aware Motion Deblurring. In *ICCV*, 2019. 2, 6
- [53] Jianping Shi, Li Xu, and Jiaya Jia. Just Noticeable Defocus Blur Detection and Estimation. In *CVPR*, 2015. 7
- [54] Hyeonseok Son, Junyong Lee, Sunghyun Cho, and Seungyong Lee. Single Image Defocus Deblurring Using Kernel-Sharing Parallel Atrous Convolutions. In *ICCV*, 2021. 2, 7
- [55] Yujing Sun, Yizhou Yu, and Wenping Wang. Moiré photo restoration using multiresolution convolutional neural networks. *TIP*, 27(8):4160–4172, 2018. 2
- [56] Xin Tao, Hongyun Gao, Xiaoyong Shen, Jue Wang, and Jiaya Jia. Scale-recurrent Network for Deep Image Deblurring. In *CVPR*, 2018. 6
- [57] Chunwei Tian, Yong Xu, Lunke Fei, and Ke Yan. Deep Learning for Image Denoising: A Survey. In *International Conference on Genetic and Evolutionary Computing*, 2018. 2
- [58] Ashish Vaswani, Prajit Ramachandran, Aravind Srinivas, Niki Parmar, Blake Hechtman, and Jonathon Shlens. Scaling Local Self-Attention for Parameter Efficient Visual Backbones. In *CVPR*, 2021. 2
- [59] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. In *NeurIPS*, 2017. 2, 3, 5
- [60] Hong Wang, Qi Xie, Qian Zhao, and Deyu Meng. A Model-Driven Deep Neural Network for Single Image Rain Removal. In *CVPR*, 2020. 5, 7
- [61] Tianyu Wang, Xin Yang, Ke Xu, Shaozhe Chen, Qiang Zhang, and Rynson WH Lau. Spatial Attentive Single-Image Deraining with a High Quality Real Rain Dataset. In *CVPR*, 2019. 2, 7, 8
- [62] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid Vision Transformer: A Versatile Backbone for Dense Prediction Without Convolutions. In *ICCV*, 2021. 2
- [63] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local Neural Networks. In *CVPR*, 2018. 2
- [64] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image Quality Assessment: From Error Visibility to Structural Similarity. *TIP*, 13(4):600–612, 2004. 5
- [65] Zhihao Wang, Jian Chen, and Steven CH Hoi. Deep Learning for Image Super-resolution: A Survey. *TPAMI*, 2020. 2
- [66] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. CVT: Introducing convolutions to vision transformers. *arXiv preprint arXiv:2103.15808*, 2021. 3, 4, 7
- [67] Li Xu, Shicheng Zheng, and Jiaya Jia. Unnatural L0 Sparse Representation for Natural Image Deblurring. In *CVPR*, 2013. 6
- [68] Rui Xu, Xiangyu Xu, Kai Chen, Bolei Zhou, and Chen Change Loy. STransGAN: An Empirical Study on Transformer in GANs. *arXiv preprint arXiv:2110.13107*, 2021. 2
- [69] Fuzhi Yang, Huan Yang, Jianlong Fu, Hongtao Lu, and Baining Guo. Learning Texture Transformer Network for Image Super-Resolution. In *CVPR*, 2020. 2
- [70] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Xiyang Dai, Bin Xiao, Lu Yuan, and Jianfeng Gao. Focal Self-attention for Local-Global Interactions in Vision Transformers. *arXiv preprint arXiv:2107.00641*, 2021. 2

- [71] Wenhan Yang, Robby T. Tan, Jiashi Feng, Zongming Guo, Shuicheng Yan, and Jiaying Liu. Joint Rain Detection and Removal from a Single Image with Contextualized Deep Networks. *TPAMI*, 42(6):1377–1393, 2020. [7](#)
- [72] Kun Yuan, Shaopeng Guo, Ziwei Liu, Aojun Zhou, Fengwei Yu, and Wei Wu. Incorporating Convolution Designs into Visual Transformers. *arXiv preprint arXiv:2103.11816*, 2021. [1](#), [4](#)
- [73] Zongsheng Yue, Hongwei Yong, Qian Zhao, Deyu Meng, and Lei Zhang. Variational Denoising Network: Toward Blind Noise Modeling and Removal. In *NeurIPS*, 2019. [5](#)
- [74] Zongsheng Yue, Qian Zhao, Lei Zhang, and Deyu Meng. Dual Adversarial Network: Toward Real-world Noise Removal and Noise Generation. In *ECCV*, 2020. [1](#), [2](#), [5](#)
- [75] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. CycleISP: Real image restoration via improved data synthesis. In *CVPR*, 2020. [5](#)
- [76] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Learning Enriched Features for Real Image Restoration and Enhancement. In *ECCV*, 2020. [1](#), [2](#), [3](#), [5](#)
- [77] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-Stage Progressive Image Restoration. In *CVPR*, 2021. [1](#), [2](#), [5](#), [6](#)
- [78] Hongguang Zhang, Yuchao Dai, Hongdong Li, and Piotr Koniusz. Deep Stacked Hierarchical Multi-patch Network for Image Deblurring. In *CVPR*, 2019. [2](#), [6](#)
- [79] Jiawei Zhang, Jinshan Pan, Jimmy Ren, Yibing Song, Linchao Bao, Rynson WH Lau, and Ming-Hsuan Yang. Dynamic Scene Deblurring Using Spatially Variant Recurrent Neural Networks. In *CVPR*, 2018. [6](#)
- [80] Kaihao Zhang, Wenhan Luo, Yiran Zhong, Lin Ma, Bjorn Stenger, Wei Liu, and Hongdong Li. Deblurring by Realistic Blurring. In *CVPR*, 2020. [6](#)
- [81] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising. *TIP*, 26(7):3142–3155, 2017. [2](#)
- [82] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image Super-Resolution Using Very Deep Residual Channel Attention Networks. In *ECCV*, 2018. [2](#)
- [83] Yulun Zhang, Kunpeng Li, Kai Li, Bineng Zhong, and Yun Fu. Residual Non-local Attention Networks for Image Restoration. In *ICLR*, 2019. [1](#), [2](#)
- [84] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual Dense Network for Image Super-Resolution. In *CVPR*, 2018. [2](#)
- [85] Zizhao Zhang, Han Zhang, Long Zhao, Ting Chen, and Tomas Pfister. Aggregating Nested Transformers. In *arXiv preprint arXiv:2105.12723*, 2021. [2](#)
- [86] Long Zhao, Zizhao Zhang, Ting Chen, Dimitris N Metaxas, and Han Zhang. Improved Transformer for High-Resolution GANs. *arXiv preprint arXiv:2106.07631*, 2021. [2](#)