

# ShadowFormer: Global Context Helps Shadow Removal

Lanqing Guo<sup>1</sup>, Siyu Huang<sup>2</sup>, Ding Liu<sup>3</sup>, Hao Cheng<sup>1</sup>, Bihan Wen<sup>1\*</sup>

<sup>1</sup>Nanyang Technological University, Singapore

<sup>2</sup>Harvard University, USA

<sup>3</sup>ByteDance Inc, USA

{lanqing001, hao006, bihan.wen}@ntu.edu.sg, huang@seas.harvard.edu, liuding@bytedance.com

## Abstract

Recent deep learning methods have achieved promising results in image shadow removal. However, most of the existing approaches focus on working locally within shadow and non-shadow regions, resulting in severe artifacts around the shadow boundaries as well as inconsistent illumination between shadow and non-shadow regions. It is still challenging for the deep shadow removal model to exploit the global contextual correlation between shadow and non-shadow regions. In this work, we first propose a Retinex-based shadow model, from which we derive a novel transformer-based network, dubbed *ShadowFormer*, to exploit non-shadow regions to help shadow region restoration. A multi-scale channel attention framework is employed to hierarchically capture the global information. Based on that, we propose a Shadow-Interaction Module (SIM) with Shadow-Interaction Attention (SIA) in the bottleneck stage to effectively model the context correlation between shadow and non-shadow regions. We conduct extensive experiments on three popular public datasets, including ISTD, ISTD+, and SRD, to evaluate the proposed method. Our method achieves state-of-the-art performance by using up to  $150\times$  fewer model parameters.<sup>1</sup>

## Introduction

Shadow is a ubiquitous phenomenon in capturing optical images under the condition of the light being partially or completely blocked. Shadow degrades the image quality by limiting both the human perception (Cucchiara et al. 2003; Nadimi and Bhanu 2004) as well as many subsequent vision tasks, *e.g.*, object detection, tracking, and semantic segmentation (Jung 2009; Sanin, Sanderson, and Lovell 2010; Zhang et al. 2018). Various solutions have been proposed for image shadow removal, including the classic approaches (Gryka, Terry, and Brostow 2015; Zhang, Zhang, and Xiao 2015; Xiao et al. 2013) that applied the physics-based illumination models. Such methods are highly limited in practice as the assumptions made in the illumination models are usually too restrictive for real-world shadow images. Recent methods (Qu et al. 2017; Hu et al. 2020) applied deep learning for image shadow removal, which have achieved remarkable performance thanks to highly flexible

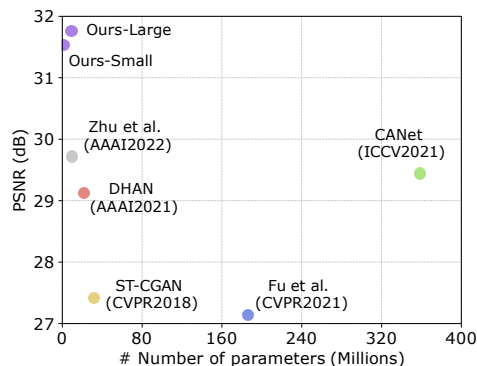


Figure 1: The PSNR performance, *v.s.*, the number of model parameters of shadow removal models on ISTD dataset.

deep models trained from large-scale training data. Comparing to many classic image restoration tasks, shadow removal is more challenging due to the following aspects: 1) The shadow patterns are arbitrary, diverse, and sometimes with highly-complex trace structures, posing challenges to supervised deep learning for achieving “trace-less” image recovery; 2) The shadow degradation is spatially non-uniform, leading to illumination and color inconsistency between the shadow and non-shadow regions.

Recent shadow removal algorithms (Le and Samaras 2019; Fu et al. 2021) attempted to tackle the first challenge by designing a separate refinement module to minimize the remaining shadow trace in the recovered images. Besides, DshadowNet (Qu et al. 2017) proposed to suppress the boundary trace artifacts by re-generating a more accurate shadow density matte. These methods, though mitigate the boundary trace artifacts at a certain level, adopt a sub-optimal deep restoration framework that incorporated multiple post-processing modules with huge computational overheads as shown in Figure 1. On the other hand, many deep shadow removal algorithms fail to preserve the illumination and color consistency in the recovered image (Qu et al. 2017; Le and Samaras 2019; Fu et al. 2021), by largely ignoring the global contextual correlation amongst the shadow and non-shadow regions. A more recent method (Chen et al.

\*Corresponding author: Bihan Wen.

<sup>1</sup><https://github.com/GuoLanqing/ShadowFormer>

2021b) employed an ad-hoc external patch matching step to exploit contextual information. However, such an approach requires the external collected image patch dataset for separately training the patch matching module and only selects the top K similar patches as a reference, in which the exploited contextual information is limited, and the computational cost is huge as shown in Figure 1.

In this work, we first extend the classic Retinex theory (Land 1977) to model the shadow degradation, thus deriving a physics-driven shadow removal process to exploit the information of non-shadow regions to help shadow region restoration. We then propose a lightweight transformer-based network to tackle the shadow removal challenges in an end-to-end way, named *ShadowFormer*. Figure 2 illustrates the proposed ShadowFormer network, in which we build the encoder-decoder framework via channel attention to efficiently stack the hierarchical information. In the bottleneck stage, we further introduce the Shadow-Interaction Module (SIM) with Shadow-Interaction Attention (SIA) to exploit the global contextual correlation between shadow and non-shadow regions across spatial and channel dimensions. By fusing the global correlation information from the bottleneck stage and low-level structure information from shallow stages, we address the challenge of color inconsistency and boundary trace in the restored shadow-free images. Experimental results show that the proposed ShadowFormer models can generate superior results consistently over the three widely-used shadow removal datasets by significantly outperforming the state-of-the-art methods using  $5\times$  to  $150\times$  fewer model parameters. The main contributions of this work are four-fold:

- We introduce a Retinex-based shadow model, which formulates the degradation of shadow and motivates us to exploit the global contexts for shadow removal.
- We propose a new single-stage transformer for shadow removal (ShadowFormer) based on the multi-scale channel attention framework.
- We propose a Shadow-Interaction Module with Shadow-Interaction Attention in the ShadowFormer to exploit the global contextual correlation between shadow and non-shadow regions.
- The comprehensive experimental results on the public ISTD, ISTD+, and SRD datasets show that the proposed ShadowFormer achieved a new state-of-the-art performance with a very lightweight network, *i.e.*, using up to  $150\times$  fewer model parameters.

## Related Work

**Shadow Removal.** Classic shadow removal methods take advantage of prior information, *e.g.*, illumination (Zhang, Zhang, and Xiao 2015), gradient (Gryka, Terry, and Brostow 2015), and region (Guo, Dai, and Hoiem 2012), for recovering illumination in shadow regions. Recent learning-based methods (Zhang et al. 2021a; Guo et al. 2021a,b; Wang et al. 2021a; Yu et al. 2022; Wang et al. 2023) apply high-quality ground truth as guidance for image processing, *e.g.*, shadow removal. One group of works still reconstructed the shadow-free image under a physical illumination model and pre-

dicted an external shadow matte. For instance, Le *et al.* (Le and Samaras 2019) applied the physical linear transformation model to enhance shadow regions by image decomposition. However, most of them employed multiple networks to predict the matte or refine the wrongly amplified boundary, which suffers from huge computational overheads, and sub-optimal design.

Some existing works have tried to explore the context information for shadow removal. For instance, (Qu et al. 2017; Jin, Sharma, and Tan 2021; Guo et al. 2022) enlarged the receptive field of the model by fusing multi-level features, which exploited contextual semantic and appearance information. Cun *et al.* (Cun, Pun, and Shi 2020) utilized a series of dilated convolutions as the backbone to exploit the context features. After that, Chen *et al.* (Chen et al. 2021b) proposed to explore the potential contextual relationship by an ad-hoc external patch matting module, which only selects the top K similar patches as a reference. In contrast, our proposed ShadowFormer is an end-to-end method with only a unified module, in which we utilize the transformer building units to exploit the global contextual information.

**Vision Transformer.** Transformer-based models take advantage of the long-range dependencies within the context, which have gained improvements among many vision tasks (Fan et al. 2021; Yang et al. 2020). For low-level vision, the pioneering work, Chen *et al.* (Chen et al. 2021a) proposed a unified model with multi-heads and tails for various restoration tasks based on the vanilla transformer structure, which is built on a huge number of model parameters training on large-scale datasets. After that, Zhang *et al.* (Zhang et al. 2021b) proposed a structure transformer network STAR for image enhancement, which exploits long-range and short-range context information. Wang *et al.* (Wang et al. 2021b) designed a hierarchical U-shaped network via the swin transformer block for image restoration, including image denoising, deblurring, and deraining.

However, most of these methods focus on global corruption tasks, such as image denoising, image super-resolution, and image deblurring. Contextual similarity provides limited valid information since the context is also degraded. For the partial corruption problem, like shadow removal, the transformer-based network should be more powerful as there is little corrupted contextual information to guide image restoration. In this paper, we explore the effectiveness of the transformer block for shadow removal and introduce a novel end-to-end pipeline for the shadow removal problem.

## Retinex-based Shadow Model

A shadow region of an image  $\mathbf{I}_s$  is caused by partial or complete blocking. The classic model (Porter and Duff 1984) decomposes  $\mathbf{I}_s$  into the shadow and non-shadow regions:

$$\mathbf{I}_s = \mathbf{I}_m \odot \mathbf{I}_s + (1 - \mathbf{I}_m) \odot \mathbf{I}_{n_s}, \quad (1)$$

where  $\mathbf{I}_{n_s}$  and  $\mathbf{I}_m$  denote the non-shadow region and the mask indicating the shadow region, respectively.  $\odot$  denotes the element-wise multiplication. Shadow removal task is to recover the underlying shadow-free image  $\mathbf{I}_{s_f}$  from  $\mathbf{I}_s$  by adjusting the image illumination and color. It is highly related to low-light image enhancement task, in which the fa-

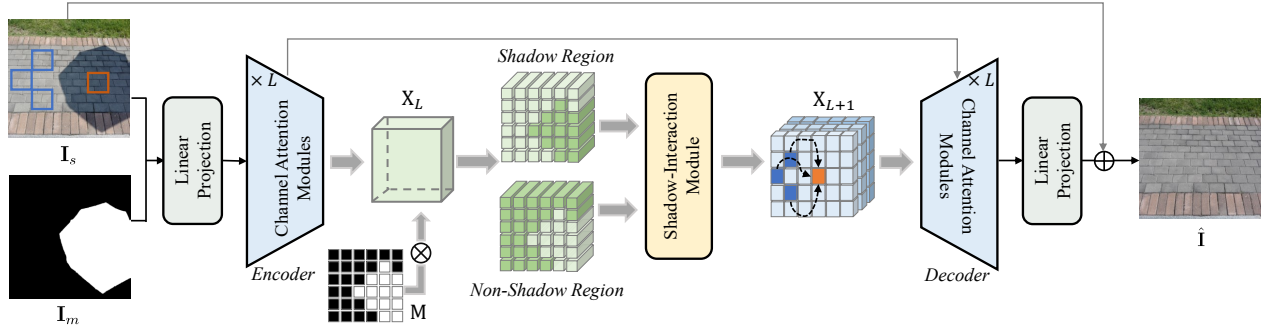


Figure 2: Overview of the ShadowFormer network. The channel attention transformer-based encoder and decoder are to extract hierarchical information from the input shadow image, and to reconstruct the shadow-free image, respectively, using a series of channel attention (CA) modules. In the bottleneck stage, we adopt a Shadow-Interaction Module (SIM) to exploit the context information across both spatial and channel dimensions from non-shadow region to help shadow region restoration.

amous Retinex theory (Land 1977) assumes that an image  $\mathbf{I}$  can be decomposed into the illumination  $\mathbf{L}$  and reflectance  $\mathbf{R}$ , based on human color perception. We define the shadow-free image as  $\mathbf{I}_{sf} = \mathbf{L}_{sf} \odot \mathbf{R}$ . Similarly, for shadow removal, we propose a *Retinex-Based Shadow Model* as

$$\mathbf{I}_s = \mathbf{I}_m \odot \mathbf{L}_s \odot \mathbf{R} + (1 - \mathbf{I}_m) \odot \mathbf{L}_{ns} \odot \mathbf{R}. \quad (2)$$

Here,  $\mathbf{L}_s$ ,  $\mathbf{L}_{ns}$  and  $\mathbf{L}_{sf}$  denote the illumination of the shadow, non-shadow regions, and the shadow-free image, respectively. Our Model (2) indicates:

- The illumination degradation varies in the shadow and non-shadow regions<sup>2</sup>, while the underlying  $\mathbf{I}_{sf}$  contains a spatially consistent  $\mathbf{L}_{sf}$ . Most of the existing shadow removal methods ignored the illumination consistency between shadow and non-shadow regions, leading to the restored  $\mathbf{L}_{sf}$  that is spatially varying.
- Both the shadow and non-shadow regions capture the same underlying  $\mathbf{R}$ , leading to a strong global contextual correlation between the two regions. Such property has been generally ignored in existing methods, which limits their effectiveness in shadow removal tasks.

## ShadowFormer

Different from previous works (Liang et al. 2021; Wang et al. 2021b) which mainly focused on those image restoration tasks dealing with global corruption, shadow removal poses its unique problem with partial corruption, meaning that the non-shadow region plays a crucial role for shadow region restoration. Therefore, we propose two critical objectives for shadow removal, which are inspired by Retinex-Based Shadow Model (2): **First**, the receptive field of the model should be as large as possible to capture the global information; otherwise the local context might provide a wrong reference. **Second**, the illumination information from non-shadow region is an important prior for shadow region restoration, which preserves the illumination consistency between shadow and non-shadow regions.

<sup>2</sup>The illumination of the non-shadow region  $\mathbf{L}_{ns}$  may not exactly equal to  $\mathbf{L}_{sf}$ , as the shadow may affect the whole environmental illumination in practice.

To achieve these two objectives, we design a transformer-based network, dubbed *ShadowFormer*, for single-stage shadow removal as shown in Figure 2. We first adopt the channel attention (CA) (Hu, Shen, and Sun 2018) into the transformer blocks to build a multi-scale encoder-decoder pipeline as it captures the global information in an efficient way. Then we propose a Shadow-Interaction Module (SIM) with Shadow-Interaction Attention (SIA) for exploiting the global contextual information across both spatial and channel dimensions from non-shadow region to help shadow region restoration in the bottleneck stage.

## Overall Architecture

Given a shadow input  $\mathbf{I}_s \in \mathbb{R}^{3 \times H \times W}$  with the corresponding shadow mask  $\mathbf{I}_m \in \mathbb{R}^{H \times W}$ , we first apply one linear projection  $\text{LinearProj}(\cdot)$  to obtain the low-level feature embedding of input, denoted by  $\mathbf{X}_0 \in \mathbb{R}^{C \times H \times W}$ , where  $C$  is the embedding dimension. Then we feed the embedding  $\mathbf{X}_0$  into the CA transformer-based encoder and decoder, each consisting of  $L$  CA modules to stack multi-scale global features. Each CA module consists of two CA blocks, as well as a down-sampling layer in the encoder or an up-sampling layer in the decoder, as shown in Figure 3(a). The CA block sequentially squeezes the spatial information via CA and captures the long-range correlation via feed-forward MLP (Dosovitskiy et al. 2010) as follows:

$$\tilde{\mathbf{X}} = \text{CA}(\text{LN}(\mathbf{X})) + \mathbf{X}, \quad (3)$$

$$\hat{\mathbf{X}} = \text{GELU}(\text{MLP}(\text{LN}(\mathbf{X}))) + \tilde{\mathbf{X}}, \quad (4)$$

where  $\text{LN}(\cdot)$  denotes the layer normalization,  $\text{GELU}(\cdot)$  denotes the GELU activation layer, and  $\text{MLP}(\cdot)$  denotes multi-layer perceptron. After passing through  $L$  modules within the encoder, we receive the hierarchical features  $\{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_L\}$ , where  $\mathbf{X}_L \in \mathbb{R}^{2^L C \times \frac{H}{2^L} \times \frac{W}{2^L}}$ . We calculate the global contextual correlation via Shadow-Interaction Module (SIM) according to the pooled feature  $\mathbf{X}_L$  in the bottleneck stage (details refer to Section ). Next, the features input to each CA module of the decoder is the concatenation of the up-sampled features and the corresponding features from the encoder through skip-connection.

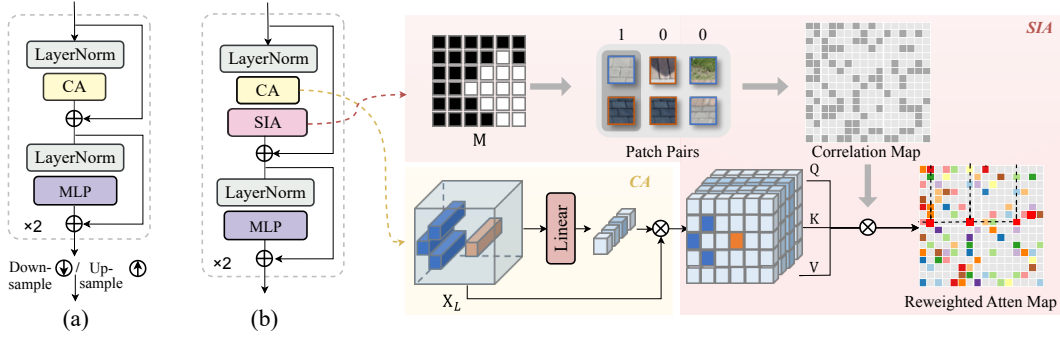


Figure 3: The detailed architectures of ShadowFormer model components. (a) Channel Attention (CA) Module in the encoder and decoder. (b) Shadow-Interaction Module (SIM), as well as an illustration of Shadow-Interaction Attention (SIA). Reweighting the attention map by the correlation map between shadow and non-shadow patches to emphasize the contextual correlation between shadow and non-shadow regions.

### Shadow-Interaction Module

Since shadow removal is a partial corruption task, the existing local attention mechanisms (Liu et al. 2021; Wang et al. 2021b) would be highly limited for shadow removal since the regions inside the window may all be corrupted. To this end, we propose a novel Shadow-Interaction Module (SIM) to exploit the global attention information across both spatial and channel dimensions.

Given a feature map  $\mathbf{Y} \in \mathbb{R}^{\hat{C} \times \hat{H} \times \hat{W}}$  normalized by a LayerNorm (LN) layer, the SIM first employs CA to re-weight the channels. To reduce the computational cost, we then split the re-weighted feature map into a sequence of non-overlapping windows  $\mathbf{X} \in \mathbb{R}^{N \times (P^2 \times \hat{C})}$ , where  $P \times P$  is the size of each window, and  $N = \hat{H}\hat{W}/P^2$  is the number of windows. Here the corresponding receptive field of one window  $(P \times 2^L)^2$  is very large while mapping to the spatial domain. After that, we perform a Shadow-Interaction Attention (SIA) on the flattened features in each window to capture the global contextual information. Figure 3(b) illustrates the detailed architecture of SIM, which consists of two blocks, each can be represented as follows:

$$\tilde{\mathbf{X}} = \text{SIA}(\text{CA}(\text{LN}(\mathbf{X})), \Sigma) + \mathbf{X}, \quad (5)$$

$$\hat{\mathbf{X}} = \text{GELU}(\text{MLP}(\text{LN}(\mathbf{X}))) + \tilde{\mathbf{X}}, \quad (6)$$

where  $\Sigma$  denotes the patch-wise correlation map.

**Shadow-Interaction Attention.** One location vector  $\mathbf{X}_L^{ij} \in \mathbb{R}^{2^L C \times 1 \times 1}$  of the pooled feature map  $\mathbf{X}_L$  can correspond to one patch in the input shadow image, where  $i, j$  denote the spatial indexes of the feature map as well as its corresponding mask. In the meanwhile, we apply max-pooling to the shadow mask  $\mathbf{I}_m$  into the same spatial dimension of  $\mathbf{X}_L$ , denoted as  $\mathbf{M}$ . According to the shadow mask, the patches from shadow image can be categorized into shadow patches, and non-shadow patches, denoted as 1 and 0, respectively. Intuitively, we can exploit the patch-wise correlation map  $\Sigma$  between non-shadow and shadow regions as follows:

$$\Sigma^{ij} = \mathbf{M}^i \oplus \mathbf{M}^j \quad \forall i, j, \quad (7)$$

where  $\oplus$  denotes exclusive OR operation. Based on that, we propose a Shadow-Interaction Attention (SIA) to reweight

the attention map to emphasize the similarity between non-shadow region and shadow region as follows:

$$\text{SIA}(\mathbf{X}, \Sigma) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{d}\right)\mathbf{V}[\sigma\Sigma + (1-\sigma)\mathbf{1}], \quad (8)$$

where  $\sigma \in (0, 1)$  adjusts the weight of shadow-shadow and nonshadow-nonshadow pairs and  $d$  is the scaling parameter.  $\mathbf{Q}$ ,  $\mathbf{K}$ , and  $\mathbf{V}$  represent the projected queries, keys, and values of input feature map  $\mathbf{X}$ .

Finally, we apply a linear projection to obtain a residual image  $\mathbf{I}_r$ . The final output is obtained by  $\hat{\mathbf{I}} = \mathbf{I}_s + \mathbf{I}_r$ . Different from other shadow removal methods (Fu et al. 2021; Chen et al. 2021b; Zhu et al. 2022) are trained under a hybrid loss function, we only use one  $\ell_1$  loss to constrain the pixel-wise consistency:

$$\mathcal{L}(\mathbf{I}_{gt}, \hat{\mathbf{I}}) = \|\mathbf{I}_{gt} - \hat{\mathbf{I}}\|, \quad (9)$$

where the  $\hat{\mathbf{I}}$  is the output image and  $\mathbf{I}_{gt}$  is the ground truth shadow-free image.

## Experiments

### Implementation Details

**Networks.** The proposed ShadowFormer is implemented using PyTorch. Following (Vaswani et al. 2017), we train our model using AdamW optimizer (Loshchilov and Hutter 2017) with the momentum as (0.9, 0.999) and the weight decay as 0.02. The initial learning rate is  $2e^{-4}$ , then gradually reduces to  $1e^{-6}$  with the cosine annealing (Loshchilov and Hutter 2016). We set the  $\sigma = 0.2$  in our experiments. (We can achieve comparable results with  $\sigma = 0, 0.1, 0.3$ .) We propose two variants of ShadowFormer, denoted as Ours-Large and Ours-Small. Ours-Large model adopts a four-scale encoder-decoder structure ( $L = 3$ ), while Ours-Small uses a three-scale encoder-decoder structure ( $L = 2$ ). We set the first feature embedding dimension as  $C = 32$  and  $C = 24$ , for Ours-Large and Ours-Small, respectively. More experimental settings can be found in **supplementary**.

**Datasets.** We work with three benchmark datasets for the various shadow removal experiments: (1) ISTD (Wang, Li,

	Method	Params	Shadow Region (S)			Non-Shadow Region (NS)			All Image (ALL)		
			PSNR $\uparrow$	SSIM $\uparrow$	RMSE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	RMSE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	RMSE $\downarrow$
256 $\times$ 256	Input Image	-	22.40	0.936	32.10	27.32	0.976	7.09	20.56	0.893	10.88
	Guo <i>et al.</i> (Guo, Dai, and Hoiem 2012)	-	27.76	0.964	18.65	26.44	0.975	7.76	23.08	0.919	9.26
	MaskShadow-GAN (Hu et al. 2019)	13.8M	-	-	12.67	-	-	6.68	-	-	7.41
	ST-CGAN (Wang, Li, and Yang 2018)	31.8M	33.74	0.981	9.99	29.51	0.958	6.05	27.44	0.929	6.65
	DSC (Hu et al. 2020)	22.3M	34.64	0.984	8.72	31.26	0.969	5.04	29.00	0.944	5.59
	DHAN (Cun, Pun, and Shi 2020)	21.8M	35.53	0.988	7.49	31.05	0.971	5.30	29.11	0.954	5.66
	Fu <i>et al.</i> (Fu et al. 2021)	186.5M	34.71	0.975	7.91	28.61	0.880	5.51	27.19	0.945	5.88
	Zhu <i>et al.</i> (Zhu et al. 2022)	10.1M	36.95	0.987	8.29	31.54	0.978	4.55	29.85	0.960	5.09
	Ours-Small	2.4M	<b>37.99</b>	<b>0.990</b>	<b>6.16</b>	<b>33.89</b>	<b>0.980</b>	<b>3.90</b>	<b>31.81</b>	<b>0.967</b>	<b>4.27</b>
	Ours-Large	9.3M	<b>38.19</b>	<b>0.991</b>	<b>5.96</b>	<b>34.32</b>	<b>0.981</b>	<b>3.72</b>	<b>32.21</b>	<b>0.968</b>	<b>4.09</b>
Original	Input Image	-	22.34	0.935	33.23	26.45	0.947	7.25	20.33	0.874	11.35
	ARGAN (Ding et al. 2019)	-	-	-	9.21	-	-	6.27	-	-	6.63
	DHAN (Cun, Pun, and Shi 2020)	21.8M	34.79	0.983	8.13	29.54	0.941	5.94	27.88	0.921	6.29
	CANet (Chen et al. 2021b)	358.2M	-	-	8.86	-	-	6.07	-	-	6.15
	Ours-Small	2.4M	<b>36.85</b>	<b>0.985</b>	<b>6.93</b>	<b>31.88</b>	<b>0.952</b>	<b>4.59</b>	<b>30.16</b>	<b>0.934</b>	<b>4.96</b>
	Ours-Large	9.3M	<b>37.03</b>	<b>0.985</b>	<b>6.76</b>	<b>32.20</b>	<b>0.953</b>	<b>4.44</b>	<b>30.47</b>	<b>0.935</b>	<b>4.79</b>

Table 1: The quantitative results of shadow removal using our models and recent methods on ISTD (Wang, Li, and Yang 2018) datasets. We put “-” to denote those models or results that are not available.

Method	Shadow Region (S)			Non-Shadow Region (NS)			All Image (ALL)		
	PSNR $\uparrow$	SSIM $\uparrow$	RMSE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	RMSE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	RMSE $\downarrow$
Input Image	18.96	0.871	36.69	31.47	0.975	4.83	18.19	0.830	14.05
Guo <i>et al.</i> (Guo, Dai, and Hoiem 2012)	-	-	29.89	-	-	6.47	-	-	12.60
DSC (Hu et al. 2020)	30.65	0.960	8.62	31.94	0.965	4.41	27.76	0.903	5.71
DHAN (Cun, Pun, and Shi 2020)	33.67	0.978	8.94	34.79	0.979	4.80	30.51	0.949	5.67
Fu <i>et al.</i> (Fu et al. 2021)	32.26	0.966	9.55	31.87	0.945	5.74	28.40	0.893	6.50
Zhu <i>et al.</i> (Zhu et al. 2022)	34.94	0.980	7.44	35.85	0.982	3.74	31.72	0.952	4.79
Ours-Small	<b>36.13</b>	<b>0.988</b>	<b>6.05</b>	<b>35.95</b>	<b>0.986</b>	<b>3.55</b>	<b>32.38</b>	<b>0.955</b>	<b>4.09</b>
Ours-Large	<b>36.91</b>	<b>0.989</b>	<b>5.90</b>	<b>36.22</b>	<b>0.989</b>	<b>3.44</b>	<b>32.90</b>	<b>0.958</b>	<b>4.04</b>

Table 2: The quantitative results of shadow removal using our models and recent methods on SRD (Qu et al. 2017) datasets.

and Yang 2018) dataset includes 1330 training and 540 testing triplets (shadow images, masks, and shadow-free images). (2) Adjusted ISTD (ISTD+) dataset (Le and Samaras 2019) reduces the illumination inconsistency between the shadow and shadow-free image of ISTD by the image processing algorithm, which has the same number of triplets with ISTD. (3) SRD (Qu et al. 2017) dataset consists of 2680 training and 408 testing pairs of shadow and shadow-free images without the ground truth shadow masks. We use the predicted masks that are provided by DHAN (Cun, Pun, and Shi 2020) for training and testing.

**Evaluation metrics.** Following the previous works (Wang, Li, and Yang 2018; Guo, Dai, and Hoiem 2012; Qu et al. 2017; Le and Samaras 2019; Cun, Pun, and Shi 2020; Fu et al. 2021), we utilize the root mean square error (RMSE) in the LAB color space as the **quantitative** evaluation metric of the shadow removal results, comparing to the ground truth shadow-free images. Besides, we also adopt the Peak Signal-to-Noise Ratio (PSNR) and the structural similarity (SSIM) to measure the performance of various methods in the RGB color space. For the PSNR and SSIM metrics, higher values represent better results.

### Comparison with State-of-the-Art Methods

We compare the proposed method with the popular or state-of-the-art (SOTA) shadow removal algorithms, including one traditional method, *i.e.*, Guo *et al.* (Guo, Dai, and Hoiem 2012), and several deep learning-based methods, *i.e.*,

Method	Shadow		Non-Shadow		All	
	PSNR $\uparrow$	RMSE $\downarrow$	PSNR $\uparrow$	RMSE $\downarrow$	PSNR $\uparrow$	RMSE $\downarrow$
Input Image	20.83	40.2	37.46	2.6	20.46	8.5
Param-Net	-	9.7	-	3.0	-	4.0
SP+M-Net	37.59	5.9	36.02	3.0	32.94	3.5
DHAN	32.92	11.2	27.15	7.1	25.66	7.8
Fu <i>et al.</i>	36.04	6.6	31.16	3.8	29.45	4.2
Ours-Small	<b>39.53</b>	<b>5.4</b>	<b>38.67</b>	<b>2.4</b>	<b>35.42</b>	<b>2.8</b>
Ours-Large	<b>39.67</b>	<b>5.2</b>	<b>38.82</b>	<b>2.3</b>	<b>35.46</b>	<b>2.8</b>

Table 3: The quantitative results of shadow removal using our models and recent methods on ISTD+ dataset.

MaskShadow-GAN (Hu et al. 2019), ST-CGAN (Wang, Li, and Yang 2018), DSC (Hu et al. 2020), ARGAN (Ding et al. 2019), DHAN (Cun, Pun, and Shi 2020), SP+M-Net (Le and Samaras 2019), Fu *et al.* (Fu et al. 2021), CANet (Chen et al. 2021b), and Zhu *et al.* (Zhu et al. 2022). As there are different experiment settings (in training and testing) that were adopted by previous shadow removal works, for a fair comparison, we conduct experiments on two major settings over ISTD dataset following two most recent works, *i.e.*, Zhu *et al.* (Zhu et al. 2022) and DHAN (Cun, Pun, and Shi 2020): (1) The results are resized into a resolution of 256  $\times$  256 for evaluation. (2) The original image resolutions are kept in both training and testing stages. For SRD and ISTD+ datasets, we only employ the setting (1) following the setting of most competing methods.






Figure 4: Examples of shadow removal results on the ISTD (Wang, Li, and Yang 2018) dataset and SRD (Qu et al. 2017). The input shadow image (a), the estimated results of DSC (Hu et al. 2020) (b), Fu *et al.* (Fu et al. 2021) (c), DHAN (Cun, Pun, and Shi 2020) (d), Ours-Small (e), Ours-Large (f), and the ground truth (g), respectively. Zoom in to see the details.

**Quantitative measure.** Tables 1, 2&3 show the quantitative results on the testing sets over ISTD, ISTD+, and SRD, respectively. It is clear that our methods outperform all competing methods by large margins in shadow area, non-shadow area and the whole image over all of the three datasets. Some methods would even destroy the non-shadow region after shadow removal as shown in the result of non-shadow area (N) over ISTD+ dataset in Table 3. With the merits of the multi-scale CA transformer architecture, ShadowFormer can effectively exploit the hierarchical global contextual information towards trace-less results. The proposed SIM can exploit the contextual correlation between shadow and non-shadow regions to correct the illumination of shadow region, preserving the illumination consistency. Besides, the most recent SOTA methods, *i.e.*, Fu *et al.* (Fu et al. 2021) and CANet (Chen et al. 2021b), employed very deep and large backbones, *e.g.*, U-Net256 (Ronneberger, Fischer, and Brox 2015) and ResNeXt (Xie et al. 2017). Compared with them, our methods only use less than 0.5% (Ours-Small) or 2.5% (Ours-Large) of the number of parameters of those competing methods. With far fewer parameters, ShadowFormer provides more superior shadow removal performance on the test sets than all the recently proposed models, achieving the new SOTA results over three widely-used benchmarks.

**Qualitative measure.** To further demonstrate the advantage of ShadowFormer against other competing methods, Figure 4 presents the visual examples of the shadow removal results on ISTD and SRD, respectively. More visual examples

can be found in the **supplementary**. Note that the images from the ISTD dataset have high context similarity and the scene is relatively simple. In these samples of ISTD dataset, previous works usually produce illumination inconsistencies and wrongly-enhanced shadow boundaries. The DSC (Hu et al. 2020) and DHAN (Cun, Pun, and Shi 2020) methods would wrongly enlighten some regions with insufficient lightness, *e.g.*, the sky region in third row in Figure 4, leading to many ghosts. In addition, almost all competing methods cannot preserve the illumination consistencies between shadow and non-shadow regions, which seriously destroy the image structure and patterns, thus causing sharp boundaries as shown in the first and second rows in Figure 4. However, with the merits of proposed SIM, it is clear that our methods can successfully enhance the shadow region with the help of non-shadow region. On the other hand, the image structures on SRD dataset are more complicated and always with diverse colors. In these samples of SRD dataset, previous works lost the ability to enhance the illumination of the background and suppress the boundary artifacts in a complicated and colorful region, *e.g.*, the blue poster in the fourth row in Figure 4.

### Ablation Study

We implement and evaluate a series of variants of ShadowFormer over the ISTD dataset, to thoroughly investigate the impact of these components we proposed in this work. Table 4 shows the evaluation results and Figure 5 demonstrates the visual examples on different variants.




Figure 5: Visual examples of the results, the zoom-in regions, and the zoom-in error maps for ablation study, including shadow input, and results of four ablation experiments corresponding to the No. in Table 4.

	CA	SA	SIA	Shadow		Non-Shadow		All	
				PSNR $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$
①		✓	✓	36.61	0.986	32.00	0.950	30.04	0.934
②	✓			36.49	0.984	30.95	0.952	29.41	0.933
③	✓	✓		36.48	0.985	32.00	0.952	30.16	0.934
④	✓		✓	<b>37.03</b>	<b>0.985</b>	<b>32.20</b>	<b>0.953</b>	<b>30.47</b>	<b>0.935</b>

Table 4: Quantitative evaluation results on ISTD dataset over Ours-Large model against its variants without the CA, without the SIA, and without spatial attention (SA).

**The effectiveness of CA transformer:** We remove the Channel Attention (CA) in encoder and decoder, and replacing with the same number of vanilla window-based spatial attention (SA) (Wang et al. 2021b), denoted by ①. It turns out that the SA within small window in the shallow layers is useless, leading to shadow and non-shadow regions inconsistency of restored image, as shown in Table 4. **Comparison between SA and SIA:** We also conduct the experiment to verify the effectiveness of the Shadow-Interaction Attention (SIA). Specifically, we first remove the whole SIA in the bottleneck stage, denoted by ②. As shown in ② of Figure 5, the result of the model with only channel attention has obvious boundary artifacts and illumination inconsistency between the shadow and non-shadow regions since the channel attention cannot preserve the spatial consistency. Besides, we compare the SIA with vanilla SA, where we do not use the correlation map between shadow/non-shadow patches to re-weight the attention map, but still preserve the vanilla SA, denoted by ③. The vanilla SA might over-explore the correlation between shadow and shadow regions, leading to a wrong guidance for restoration. In contrast, the SIA in the bottleneck stage contributes to both shadow and non-shadow regions restoration, achieving trace-less results. It exploits the self-similarity between each spatial location and refines the contour of input shadow mask. From the comparison in ④ of Figure 5, the result of the model with SIA is trace-less and the artifacts has been suppressed to some extension. More ablation studies are included in **supplementary**.

## Network Analysis

**The activation area and effect of ShadowFormer.** To verify our motivation that the employed Shadow-Interaction Module (SIM) can exploit the contextual information between shadow and non-shadow regions, where the non-corruption area can help shadow area enhancement. We illustrate the attention maps for some key patches within




Figure 6: Visualization of correlation maps in SIM for **three** (red, purple, and blue) key points within shadow region between shadow and non-shadow regions.

Method	Real Data			Synthetic Data		
	S	N	A	S	N	A
SP+M-Net	8.5	3.6	4.4	11.3(+2.8)	3.6(+0)	4.9(+0.5)
DHAN	7.4	4.8	5.2	9.7(+2.3)	4.0(-0.8)	4.9(-0.3)
Ours-Large	<b>5.2</b>	<b>2.3</b>	<b>2.8</b>	<b>6.1(+0.9)</b>	<b>2.3(+0)</b>	<b>2.9(+0.1)</b>

Table 5: Quantitative evaluation results (RMSE $\downarrow$ ) on ISTD+ dataset over the SOTA methods and our method training with real data or synthetic data (where S, N, A represent the shadow area, non-shadow area and all image, respectively).

shadow region in Figure 6. We can see that the corresponding regions, *e.g.*, glass, tiles, and stones, can successfully find the content related non-corruption references.

**The robustness to synthetic data.** Obtaining a large-scale, diverse, and accurate dataset has still been a big challenge, and it limits the generalization and performance of the learned models on shadow images with unseen shapes/scenes. Recently, many researchers (Sidorov 2019; Inoue and Yamasaki 2020) tried to propose the shadow synthetic algorithm to tackle the data limitation problem. In order to evaluate the robustness of our model, we apply one synthetic shadow dataset (Inoue and Yamasaki 2020) as training set, then evaluate on real ISTD+ testing set. For a fair comparison, we also employ the same training set for the competing methods. As shown in Table 5, the performance of previous methods would be significantly affected when training with synthetic data especially for the shadow region. For example, the RMSE performance of shadow area for SP+M-Net (Le and Samaras 2019) drops from 8.5 to 11.3, whereas our method only drops 0.9. With merits of shadow and non-shadow correlation exploiting, the proposed ShadowFormer has better generalizability.

## Conclusion

In this work, we propose a Retinex-based shadow model, which motivates us to exploit the non-shadow region to help shadow region restoration. We introduce a novel single-stage transformer for shadow removal, called ShadowFormer, which is built on a channel attention encoder-decoder architecture, and a Shadow-Interaction Module in the bottleneck stage. We further introduce a Shadow-Interaction Attention to exploit the contextual information of images, such that the non-shadow region with mild degradation would guide the shadow restoration to improve illumination and color consistency. We show that ShadowFormer can effectively alleviate the contour artifacts towards trace-less image reconstruction, outperforming competing methods over all datasets, with far fewer amount of model parameters.

## Acknowledgment

This work was carried out at Rapid-Rich Object Search (ROSE) Lab, Nanyang Technological University (NTU). This research is supported in part by the MOE AcRF Tier 1 (RG61/22), Start-Up Grant and ASPIRE League Seed Fund.

## References

- Chen, H.; Wang, Y.; Guo, T.; Xu, C.; Deng, Y.; Liu, Z.; Ma, S.; Xu, C.; Xu, C.; and Gao, W. 2021a. Pre-trained image processing transformer. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 12299–12310.
- Chen, Z.; Long, C.; Zhang, L.; and Xiao, C. 2021b. CANet: A Context-Aware Network for Shadow Removal. In *Proc. IEEE Int'l Conf. Computer Vision*, 4743–4752.
- Cucchiara, R.; Grana, C.; Piccardi, M.; and Prati, A. 2003. Detecting moving objects, ghosts, and shadows in video streams. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 25(10): 1337–1342.
- Cun, X.; Pun, C.-M.; and Shi, C. 2020. Towards Ghost-Free Shadow Removal via Dual Hierarchical Aggregation Network and Shadow Matting GAN. In *Proc. AAAI Conf. on Artificial Intelligence*, 10680–10687.
- Ding, B.; Long, C.; Zhang, L.; and Xiao, C. 2019. Argan: Attentive recurrent generative adversarial network for shadow detection and removal. In *Proc. IEEE Int'l Conf. Computer Vision*, 10213–10222.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2010. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Fan, H.; Xiong, B.; Mangalam, K.; Li, Y.; Yan, Z.; Malik, J.; and Feichtenhofer, C. 2021. Multiscale vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6824–6835.
- Fu, L.; Zhou, C.; Guo, Q.; Juefei-Xu, F.; Yu, H.; Feng, W.; Liu, Y.; and Wang, S. 2021. Auto-exposure fusion for single-image shadow removal. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 10571–10580.
- Gryka, M.; Terry, M.; and Brostow, G. J. 2015. Learning to remove soft shadows. *ACM Transactions on Graphics*, 34(5): 1–15.
- Guo, L.; Huang, S.; Liu, H.; and Wen, B. 2021a. FINO: Flow-based Joint Image and Noise Model. *arXiv preprint arXiv:2111.06031*.
- Guo, L.; Wan, R.; Su, G.-M.; Kot, A. C.; and Wen, B. 2021b. Multi-Scale Feature Guided Low-Light Image Enhancement. In *2021 IEEE International Conference on Image Processing (ICIP)*, 554–558. IEEE.
- Guo, L.; Wang, C.; Yang, W.; Huang, S.; Wang, Y.; Pfister, H.; and Wen, B. 2022. ShadowDiffusion: When Degradation Prior Meets Diffusion Model for Shadow Removal. *arXiv preprint arXiv:2212.04711*.
- Guo, R.; Dai, Q.; and Hoiem, D. 2012. Paired regions for shadow detection and removal. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 35(12): 2956–2967.
- 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.
- Hu, X.; Fu, C.-W.; Zhu, L.; Qin, J.; and Heng, P.-A. 2020. Direction-Aware Spatial Context Features for Shadow Detection and Removal. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 42(11): 2795–2808.
- Hu, X.; Jiang, Y.; Fu, C.-W.; and Heng, P.-A. 2019. Mask-ShadowGAN: Learning to remove shadows from unpaired data. In *Proc. IEEE Int'l Conf. Computer Vision*, 2472–2481.
- Inoue, N.; and Yamasaki, T. 2020. Learning from synthetic shadows for shadow detection and removal. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(11): 4187–4197.
- Jin, Y.; Sharma, A.; and Tan, R. T. 2021. DC-ShadowNet: Single-Image Hard and Soft Shadow Removal Using Unsupervised Domain-Classifer Guided Network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 5027–5036.
- Jung, C. R. 2009. Efficient background subtraction and shadow removal for monochromatic video sequences. *IEEE Trans. on Multimedia*, 11(3): 571–577.
- Land, E. H. 1977. The retinex theory of color vision. *Scientific american*, 237(6): 108–129.
- Le, H.; and Samaras, D. 2019. Shadow removal via shadow image decomposition. In *Proc. IEEE Int'l Conf. Computer Vision*, 8578–8587.
- 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.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*.
- Loshchilov, I.; and Hutter, F. 2016. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Nadimi, S.; and Bhanu, B. 2004. Physical models for moving shadow and object detection in video. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 26(8): 1079–1087.
- Porter, T.; and Duff, T. 1984. Compositing digital images. In *Proceedings of the 11th annual conference on Computer graphics and interactive techniques*, 253–259.
- Qu, L.; Tian, J.; He, S.; Tang, Y.; and Lau, R. W. 2017. Deshadownet: A multi-context embedding deep network for shadow removal. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 4067–4075.
- 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.



- Sanin, A.; Sanderson, C.; and Lovell, B. C. 2010. Improved shadow removal for robust person tracking in surveillance scenarios. In *International Conference on Pattern Recognition*, 141–144. IEEE.
- Sidorov, O. 2019. Conditional gans for multi-illuminant color constancy: Revolution or yet another approach? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 0–0.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. 5998–6008.
- Wang, J.; Li, X.; and Yang, J. 2018. Stacked conditional generative adversarial networks for jointly learning shadow detection and shadow removal. In *Proc. IEEE Int’l Conf. Computer Vision and Pattern Recognition*, 1788–1797.
- Wang, Y.; Wan, R.; Yang, W.; Li, H.; Chau, L.-P.; and Kot, A. C. 2021a. Low-Light Image Enhancement with Normalizing Flow. *arXiv preprint arXiv:2109.05923*.
- Wang, Y.; Yu, Y.; Yang, W.; Guo, L.; Chau, L.-P.; Kot, A.; and Wen, B. 2023. Raw image reconstruction with learned compact metadata. *arXiv preprint arXiv:2302.12995*.
- Wang, Z.; Cun, X.; Bao, J.; and Liu, J. 2021b. Uformer: A general u-shaped transformer for image restoration. *arXiv preprint arXiv:2106.03106*.
- Xiao, C.; She, R.; Xiao, D.; and Ma, K.-L. 2013. Fast shadow removal using adaptive multi-scale illumination transfer. In *Computer Graphics Forum*, volume 32, 207–218. Wiley Online Library.
- Xie, S.; Girshick, R.; Dollár, P.; Tu, Z.; and He, K. 2017. Aggregated residual transformations for deep neural networks. In *Proc. IEEE Int’l Conf. Computer Vision and Pattern Recognition*, 1492–1500.
- Yang, Z.; Wang, Y.; Chen, X.; Liu, J.; and Qiao, Y. 2020. Context-Transformer: Tackling Object Confusion for Few-Shot Detection. In *AAAI*, 12653–12660.
- Yu, Y.; Yang, W.; Tan, Y.-P.; and Kot, A. C. 2022. Towards robust rain removal against adversarial attacks: A comprehensive benchmark analysis and beyond. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6013–6022.
- Zhang, L.; Zhang, Q.; and Xiao, C. 2015. Shadow remover: Image shadow removal based on illumination recovering optimization. *IEEE Trans. on Image Processing*, 24(11): 4623–4636.
- Zhang, R.; Guo, L.; Huang, S.; and Wen, B. 2021a. Relie: Deep reinforcement learning for customized low-light image enhancement. In *Proceedings of the 29th ACM International Conference on Multimedia*, 2429–2437.
- Zhang, W.; Zhao, X.; Morvan, J.-M.; and Chen, L. 2018. Improving shadow suppression for illumination robust face recognition. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 41(3): 611–624.
- Zhang, Z.; Jiang, Y.; Jiang, J.; Wang, X.; Luo, P.; and Gu, J. 2021b. STAR: A Structure-Aware Lightweight Transformer for Real-Time Image Enhancement. In *Proc. IEEE Int’l Conf. Computer Vision*, 4106–4115.
- Zhu, Y.; Xiao, Z.; Fang, Y.; Fu, X.; Xiong, Z.; and Zha, Z.-J. 2022. Efficient Model-Driven Network for Shadow Removal.