

Leveraging Inpainting for Single-Image Shadow Removal

Xiaoguang Li^{1*}, Qing Guo^{2*†}, Rabab Abdelfattah¹, Di Lin³,
 Wei Feng³, Ivor Tsang², Song Wang¹

¹University of South Carolina, USA, ²IHPC and CFAR, Agency for Science, Technology and Research, Singapore, ³Tianjin University, China

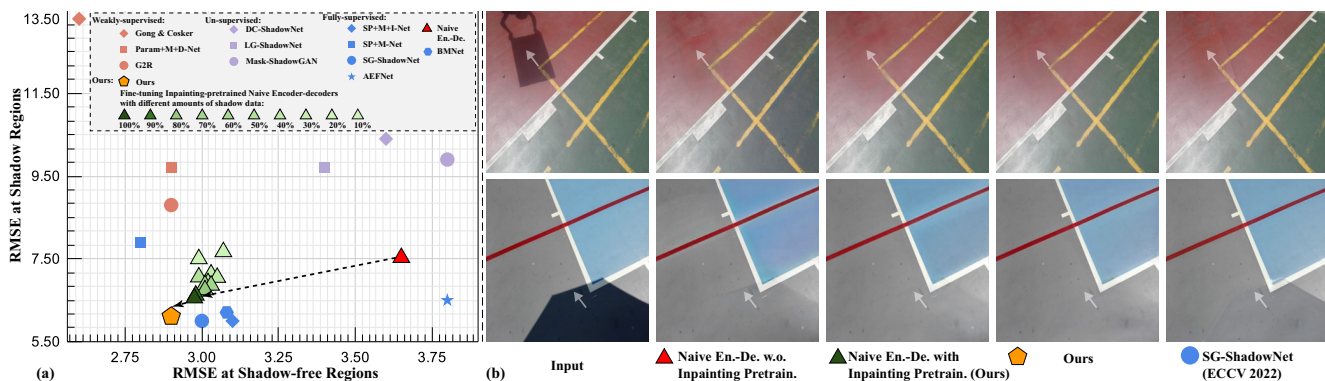


Figure 1: (a) Comparison with SOTA methods on ISTD+ dataset[24] at shadow (Y-axis) and shadow-free (X-axis) regions. We also present the results of a naive encoder-decoder network (Naive En.-De.) trained with different strategies: (1) training on inpainting data and fine-tuning on different amounts of shadow data; (2) only training on all shadow data. (b) Displays two examples of four methods and we use white arrows to highlight the main advantages.

Abstract

Fully-supervised shadow removal methods achieve the best restoration qualities on public datasets but still generate some shadow remnants. One of the reasons is the lack of large-scale shadow & shadow-free image pairs. Unsupervised methods can alleviate the issue but their restoration qualities are much lower than those of fully-supervised methods. In this work, we find that pretraining shadow removal networks on the image inpainting dataset can reduce the shadow remnants significantly: a naive encoder-decoder network gets competitive restoration quality w.r.t. the state-of-the-art methods via only 10% shadow & shadow-free image pairs. After analyzing networks with/without inpainting pretraining via the information stored in the weight (IIW), we find that inpainting pretraining improves restoration quality in non-shadow regions and enhances the generalization ability of networks significantly. Additionally, shadow removal fine-tuning enables networks to fill in the details of shadow regions. Inspired by these observations we formulate shadow removal as an adaptive fusion task that takes advantage of both shadow removal and image inpainting. Specifically, we develop an adaptive fusion network consisting of two

encoders, an adaptive fusion block, and a decoder. The two encoders are responsible for extracting the features from the shadow image and the shadow-masked image respectively. The adaptive fusion block is responsible for combining these features in an adaptive manner. Finally, the decoder converts the adaptive fused features to the desired shadow-free result. The extensive experiments show that our method empowered with inpainting outperforms all state-of-the-art methods. We have realized codes and models in <https://github.com/tsingqguo/inpaint4shadow>

1. Introduction

Shadows in images are formed when some objects block the light source, which not only reduces the image quality but also affects the subsequent intelligent tasks like visual object detection [34][9], objects tracking [36], face recognition [44], face landmark detection [8], etc. Single-image shadow removal is to map the shadow regions to their shadow-free counterparts, which can enhance the visual quality and benefit the intelligent tasks.

In recent years, with the development of advanced learning algorithms and deep architectures (i.e., GAN [12] and

CycleGAN [46]), deep learning-based shadow removal methods [25, 10, 37] have achieved significant progress and obtained top restoration qualities on public datasets. However, these methods may still lead to some shadow remnants. As the two cases shown in Fig. 1, even the state-of-the-art method SG-ShadowNet [37] generates obviously inconsistent colors across shadow boundaries. One of the reasons is the lack of large-scale shadow & shadow-free image pairs. The commonly-used ISTD+ dataset [24] only contains 1330 pairs of shadow and shadow-free images for training. To alleviate the requirements for the amounts of image pairs, researchers also develop weakly-supervised [11, 25, 31] and unsupervised shadow removal methods [22, 30, 20, 31]. However, all these methods still have a large gap to the fully-supervised methods (See the results in Fig. 1 (a)).

Meanwhile, several works [18, 40] have demonstrated that pretraining networks to predict masked patches from unmasked patches on a large-scale dataset can enhance the fully-supervised training on another small-scale dataset significantly. Actually, the task predicting masked patches is a special image inpainting task [27] that aims to fill missing pixels. Inspired by these works, we seek to leverage the image inpainting task for high-quality shadow removal. Specifically, given a deep network, we first train it on the image inpainting dataset where we can mask all clean images randomly and get large-scale clean & masked image pairs for free. Then, we finetune the network on the shadow removal dataset with limited shadow & shadow-free image pairs. As shown in Fig. 1, a naive encoder-decoder network pretrained on the inpainting dataset and fine-tuned on the 10% shadow data can achieve competitive restoration qualities (*i.e.* RMSE) at the shadow regions w.r.t. the state-of-the-art methods (*e.g.*, SP+M-Net [24]). Moreover, when we compare the networks with and without inpainting pretraining in the Fig. 1 (b), we see that inpainting pretraining can eliminate the shadow remnants effectively.

We further analyze the features of the networks with/without inpainting pretraining and see that inpainting pretraining not only improves restoration quality in non-shadow regions but also enhances the generalization ability of networks. This improvement in generalization ability can be proved by using the PAC-Bayes theorem to measure the amount of information stored in the network's weights (IIW) [38]. Furthermore, fine-tuning the networks for shadow removal enables them to fill in the details of shadow regions, thereby improving the overall restoration quality of the images. To utilize the respective advantages of shadow removal and image inpainting, we formulate shadow removal as an adaptive fusion task of shadow removal and image inpainting where make the shadow removal and image inpainting focus on the shadow and non-shadow regions respectively. To address this task, we propose an adaptive fusion network which consisting of two encoders, an adaptive fusion block, and a

decoder. The two encoders are responsible for extracting the features from the shadow image and the shadow-masked image respectively. The adaptive fusion block is responsible for combining the extracted features in an adaptive manner and the decoder converts the adaptively fused features into the desired shadow-free result. Our final method outperforms all state-of-the-art methods on the public datasets (See Fig. 1(a)) and reduces the shadow remnants significantly.

2. Related Work

2.1. General Shadow Removal Methods

To reconstruct the shadow-free image from its shadow counterpart, previous methods [7][6][17][41][39][13][32] focus more on exploiting the priors information. [17] utilize illumination difference to distinguish the shadow and shadow-free regions, then recover the shadow-free image by a lighting model. [7] [6] rebuild the shadow-free image based on the gradient information. [13] proposes to learn a patch mapping function to perform shadow removal.

Recently, deep learning-based methods have achieved excellent performance in shadow removal. [24] decomposes the shadow image as its shadow-free counterpart and a set of shadow parameters and uses illumination transformation to remove the shadow. [47] claims that the networks of shadow removal and generation can mutually promote, then proposes a unified framework to perform shadow removal and generation together. [2][37] design a feature transformation network to transform the contextual information from non-shadow regions to shadow regions. [10] reformulates shadow removal as an exposure fusion problem, utilizing the deep neuron network to predict the parameters of exposure. Although achieving a remarkable result, the above supervised-based methods require a large amount of paired shadow & shadow-free images to train the network. To alleviate this limitation, the unsupervised methods [22][30][20][31] employ the generative adversarial network (GAN) to train the network with large numbers of unpaired shadow & shadow-free images. However, the restoration results of unsupervised methods are not satisfied. In this work, we formulate shadow removal as an adaptive fusion task and propose a novel adaptive fusion network for shadow removal.

2.2. Predictive Filtering Technique

The predictive filtering technique has been widely used for image processing [1][33][3][14][42][28]. Different from the traditional convolutional neural network (CNN) which shares the same kernel along the spacial level for each convolution layer, the predictive filtering technique predicts unique kernel for each location by using a deep neural network which makes it can utilize the neighbors' information explicitly. [16] use the predictive filtering on image level for deraining. [15] employ it to improve the performance of

the generative adversarial network (GAN) based inpainting methods. [27] extend the image level filtering into the feature level and design an end-to-end image inpainting pipeline that can restore the images with large missing regions. In this work, we exploit the mechanism of the predictive filtering technique to predict the adaptive weights in an element-wise manner which can be used to take information adaptively from the extracted features.

3. Leveraging Inpainting-pretrained Networks for Shadow Removal

In this section, we first train a naive encoder-decoder network through an inpainting dataset (See Sec. 3.1). Then, we train the same network with paired shadow & shadow-free images (See Sec. 3.2) based on the inpainting-pretrained network and randomly initialized network, respectively. We conduct extensive empirical study and analysis in Sec. 3.3.

3.1. Training Network for Image Inpainting

Given a corrupted image I where some regions are cropped and indicated by a mask M , an image inpainting network is to fill the missing pixels and generate an image that is desired to be identical to the ground truth image (*i.e.*, I^*). To train the network, we obtain the corrupted image I by cropping the ground truth image I^* according to the mask, and then the image pair (I, I^*) forms a training example.

Here, we use the following setups for training an image inpainting network: ❶ Network architecture. Instead of using advanced inpainting networks (*e.g.*, MISF [27] and CICM [5]), we employ a naive encoder-decoder network that contains 11 convolution layers for the encoder and 3 transpose convolution layers for the decoder. By doing this, we can avoid the effects of some advanced designs on the observations. ❷ Dataset and masks. We train the encoder-decoder network on the *Places2 dataset*[45] that is widely used in the inpainting field. We randomly sample masks from a third-party mask dataset [29] and use them to crop the clean images. ❸ Loss functions. We follow the designs of [27] and use the GAN loss, perceptual loss, style loss, and L_1 loss for training. With the above setups, we can train a naive encoder-decoder network for image inpainting.

3.2. Training Network for Shadow Removal

Given a shadow image I and a shadow mask M , a shadow removal network is to map the shadow regions indicated by the mask to their shadow-free counterparts and is desired to produce a clean image I^* . In contrast to the image inpainting task, we can hardly collect numerous paired images (*i.e.*, (I, I^*)) for training the shadow removal network since the shadow image cannot be naively simulated via the ground truth image I^* and the two images should be captured under the same scene and the same environmental factors. Even

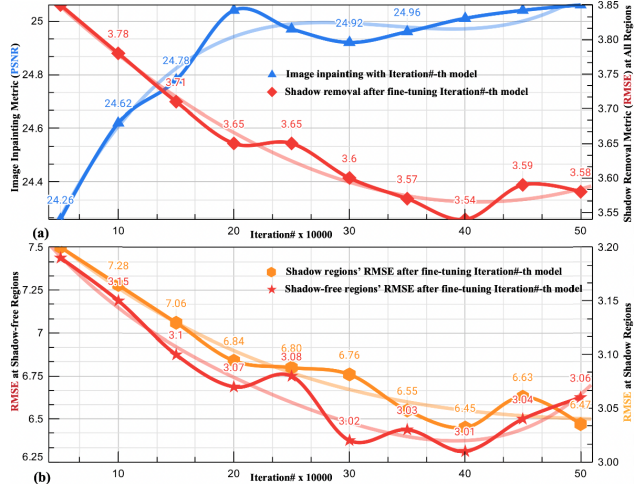


Figure 2: We first train the naive encoder-decoder network on the Places2 dataset and save the networks every 5×10^5 iterations, and then fine-tune each model with the shadow removal dataset (*i.e.*, ISTD+[24]). The blue line in (a) presents the inpainting results (*i.e.*, PSNR) of different networks on the Places2 dataset, and the red line shows shadow removal results (*i.e.*, root mean square error (RMSE)) of fine-tuned networks on both shadow and shadow-free regions. (b) presents the shadow removal results (*i.e.*, RMSE) on the shadow regions and shadow-free regions, respectively.

the widely-used ISTD+ dataset[24] only contains 1,330 pairs of shadow and shadow-free images for training.

Here, we use the following setups for training a shadow removal network: ❶ Network architecture. We employ the same encoder-decoder network in Sec. 3.1. ❷ Dataset and masks. We use the ISTD+ dataset[24] for network training and testing where the shadow mask for each image is given. ❸ Loss functions. We only use L_1 loss for training. Then, we set up two training strategies: *First*, we randomly initialize the network’s weights and train the network on the ISTD+ dataset[24]. *Second*, we use the inpainting-pretrained network as the initialization and fine-tuning the network via the ISTD+ dataset[24]. Based on the above setups, we can study how inpainting affects the shadow removal network training.

3.3. Empirical Study and Analysis

Fine-tuning inpainting-pretrained networks for shadow removal. Following the setups in Sec. 3.1, we first train an encoder-decoder network for image inpainting and store the intermediate models every 5×10^5 iterations. Then, we fine-tune these intermediate models on the shadow removal dataset (See Sec. 3.2). We show the inpainting results (*i.e.*, PSNR) of the intermediate models and the corresponding shadow removal performance (*i.e.*, RMSE) after fine-tuning in Fig. 1(a) and Fig. 2. We see that: ❶ All shadow removal networks (*i.e.*, green triangles in Fig. 1(a)) that are fine-tuned on the inpainting-pretrained models achieve much lower RMSEs at both shadow and shadow-free regions than the network trained on the randomly initialized model (*i.e.*, red triangle in Fig. 1(a)). This demonstrates that inpaint-

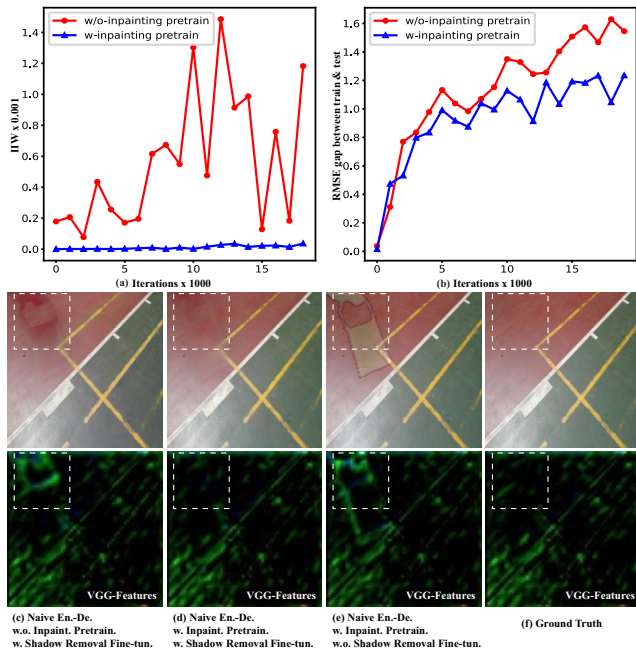


Figure 3: (a) Analysis with/without inpainting pretraining via IIW. (b) RMSE gap between the training and testing datasets. (c)-(f) Visualized comparison with/without inpainting pretraining and shadow removal fine-tuning.

ing pretraining indeed helps shadow removal task get much higher restoration quality in both shadow and shadow-free regions. ② As the iteration becomes larger, the inpainting capability of intermediate models increases (*i.e.*, the PSNR on the inpainting dataset increases) and the corresponding fine-tuned shadow removal networks get lower RMSEs at the shadow and shadow-free regions. This observation demonstrates that the shadow removal performance is related to the inpainting capability of pretrained models, directly.

Fine-tuning inpainting-pretrained networks with different amounts of shadow & shadow-free image pairs. With an adequately inpainting-pretrained network, we use different amounts of shadow & shadow-free image pairs that are randomly selected from the ISTD+ [24] to fine-tune the network, respectively. As shown in Fig. 1 (a), we see that: ① With only 10% ISTD+ [24], the fine-tuned encoder-decoder network achieves lower RMSE at shadow-free regions than the state-of-the-art method AEFNet [10] (*i.e.*, 3.07 vs. 3.80) and lower RMSE on shadow regions than the SP+M-Net method (*i.e.*, 7.67 vs. 7.9). Both methods are trained with all examples in ISTD+ dataset [24]. ② Inpainting-pretrained networks are fine-tuned with different amounts of image pairs and have similar RMSEs at both shadow and shadow-free regions. ③ The network fine-tuned with 100% image pairs outperforms three recent works including SG-ShadowNet [37], SP+M+I-Net [26], and AEFNet [10] at the shadow-free regions, which demonstrates that inpainting pretraining benefits the shadow-free preservation.

Analysis with/without inpainting pretraining via in-

formation stored in the weight (IIW) [38]. To further understand the benefits of using inpainting pretraining for shadow removal, we fine-tune a naive encoder-decoder network both with and without inpainting pretraining on the ISTD+ dataset and evaluate the network’s performance every 1000 iterations. To assess the network’s performance, we monitor two key metrics. First, we use the PAC-Bayes theorem to measure the amount of information stored in the network’s weights (IIW) [38], which is a promising indicator of a deep neural network’s generalization ability. As shown in Fig. 3 (a), with inpainting pretraining, the network’s generalization capability improved significantly. In addition to the IIW, we also calculated the Root Mean Squared Error (RMSE) gap, which is the difference in RMSE values obtained by evaluating the network on the training and testing datasets. As shown in Fig. 3 (b), the RMSE gap was smaller with inpainting pretraining than without it which can prove that inpainting pretraining can decrease the degree of overfitting of the network.

Visualized comparison with/without inpainting pretraining and shadow removal fine-tuning.

We also analyze the performance of with/without inpainting pretraining in both image and feature level as shown in Fig. 3 (c)-(f). Comparing Fig. 3 (c) with (d), we see that the inpainting pretrained network can suppress the shadow patterns (*i.e.*, the white dashed rectangles) at both the semantic level and image level effectively. Comparing Fig. 3 (d) and (e), we see that without the shadow removal fine-tuning, the network cannot recover the colors properly and the shadow patterns become more obvious. We have similar observations at the feature level.

4. Methodology

In Sec. 3, we demonstrate the significant benefits of inpainting pretraining for shadow removal. Building upon these insights, in Sec. 4, we explicitly integrate inpainting into shadow removal to fully harness its potential capabilities and formulate the shadow removal as a special adaptive fusion task, as detailed in Sec. 4.1. To this end, we propose an adaptive fusion network in Sec. 4.2 to realize the task in a unified architecture and detail the implementations in Sec. 4.3.

4.1. Problem Formulation

Given a shadow image I and a binary mask M , we first crop the shadow regions and produce a shadow-masked image $\tilde{I} = I \odot (1 - M)$. Then we aim to reconstruct the shadow-free counterpart based on the shadow image I and shadow-masked image \tilde{I} . We denote the whole process as an *adaptive fusion* task. Intuitively, such a task inherits the advantages of image inpainting and could fill the missing semantic information with similar color distributions as the shadow-free regions (See discussions in Sec. 3.3). Besides,

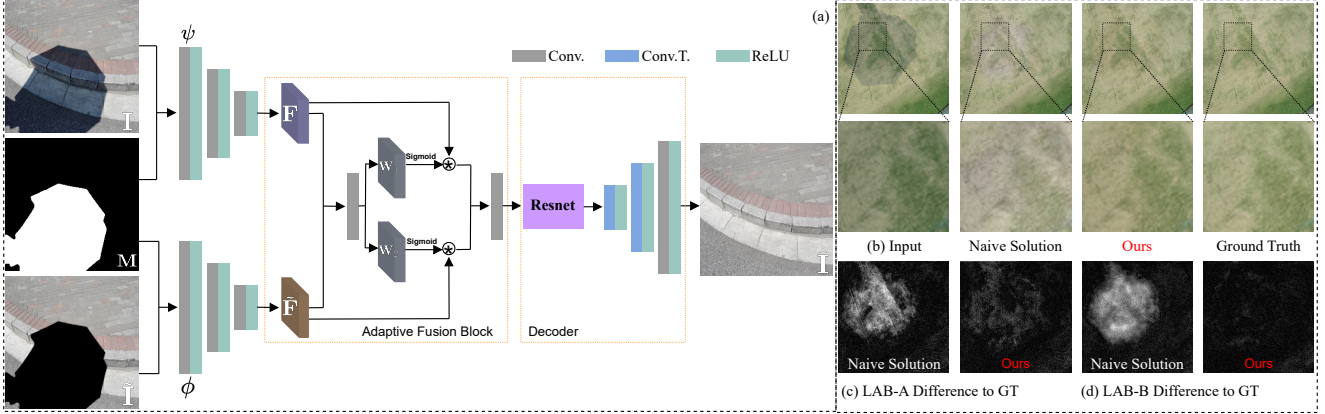


Figure 4: (a) Displays the whole framework of the proposed adaptive fusion network. (b) Shows examples of using a simple solution that input the concatenated shadow image and shadow-masked image into a naive encoder-decoder and the proposed method to handle the shadow image, respectively. We also display the enlarged regions in the second row. We further calculate the color difference to the ground truth of the A channel (*i.e.*, (c)) and B channel (*i.e.*, (d)) in the LAB color space.

the task could recover missing details in the shadow regions since the whole information is contained in the raw shadow image. However, the implementation of this task is not straightforward.

A simple solution is to concatenate the shadow-masked image and the shadow image and feed them into the same encoder-decoder network used in the Sec. 3.1. We can also perform inpainting pretraining and shadow removal finetuning to that network. We show an example in Fig. 4 (b), (c), and (d). Clearly, the result of the naive solution still has an obvious color shifting to the ground truth (See the LAB-A/B differences in Fig. 4 (c) and (d)), which demonstrates that such a naive solution cannot properly take the advantages of both tasks and the color inconsistency still remains.

4.2. Adaptive Fusion Network

Instead of concatenating the two images, we propose to handle the shadow image and shadow-masked image via two encoders which are responsible for extracting the feature from the shadow image and shadow-masked image respectively. We show the whole framework in Fig. 4(a). we denote the encoder receiving the shadow-masked image as $\phi(\cdot)$. we can represent the feature $\tilde{F} \in \mathbb{R}^{H_l \times W_l \times C_l}$ extracted from shadow-masked image as

$$\tilde{F} = \phi_l(\dots(\dots\phi_2(\phi_1(\tilde{I}, M))), \quad (1)$$

where $\phi_l(\cdot)$ is the l th layer in the encoder. Meanwhile, we can use another encoder (*i.e.*, $\psi(\cdot)$) to get the feature of the shadow image

$$F = \psi_l(\dots(\dots\psi_2(\psi_1(I, M))). \quad (2)$$

Instead of passing the shadow-masked feature (*i.e.*, \tilde{F}) and shadow feature (*i.e.* F) to the decoder directly, we design an adaptive fusion block to fuse \tilde{F} and F in an adaptive manner. Specifically, we input \tilde{F} and F into a convolution




Figure 5: Visualization of the shadow image, shadow-masked image, and the corresponding weights W_1 and W_2 .

layer to predict the adaptive weights followed by a sigmoid operation which can be represented as

$$[W_1, W_2] = \text{Sigmoid}(\text{Conv}_{\text{weight}}(\tilde{F}, F)), \quad (3)$$

where ‘Conv_{weight}’ denotes a convolution layer to predict the weights W_1 and W_2 . With the predicted weights, we can obtain the adaptive fused feature \hat{F} by

$$\hat{F} = \text{Conv}_{\text{fusion}}(\tilde{F} \otimes W_1 + F \otimes W_2), \quad (4)$$

where ‘ \otimes ’ is the element-wise multiplication operation and ‘Conv_{fusion}’ denotes a convolution layer to fuse the weighted features. When we conduct the average pooling along the channels of W_1 and W_2 respectively and illustrate the result, as shown in Fig. 5, we can observe that: in the shadow regions, the adaptive fused feature \hat{F} take more information from F which is extracted from the shadow image; in contrast, in the non-shadow regions, it takes more information from \tilde{F} which is extracted from the shadow-masked image. Finally, the adaptive fused feature \hat{F} is passed into the decoder to get the de-shadowed result.

4.3. Implementation Details

Network architectures. The encoders ϕ and ψ employ the same architecture. Specifically, the encoders consist of three convolution layers and each convolution layer is followed by a ReLU activation. The extracted features by ϕ and ψ are passed into the convolution layer Conv_{weight} to predict the weights followed by a sigmoid operation. The predicted

weights will be used to adaptive take information from the features (*i.e.* \mathbf{F} and $\bar{\mathbf{F}}$) extracted from the shadow image and shadow-masked image. $\text{Conv}_{\text{weight}}$ use a kernel of size 3×3 . Then the weighted feature is fused by another convolution layer $\text{Conv}_{\text{fusion}}$ with a kernel of size 3×3 . Finally, the fused feature is passed into the decoder. The decoder consists of eight Resnet blocks and three transposition convolution layers. Theoretically, we can add the adaptive fusion operation to all layers, which would lead to high computational complexity. Here, we equip the adaptive fusion operation to the last layer features of ϕ and ψ in order to save computation.

Loss functions. When training on the inpainting dataset, we follow [27] and use GAN loss, perceptual loss, Style loss, and L_1 loss, which guides the network to fill in missing semantic information and details. When fine-tuning the network on the shadow dataset, we only use L_1 loss to optimize since this stage mainly focuses on restoration fidelity.

Training details. We first pretrain the network 450,000 iterations with batch size 8 on the Places2 dataset, then fine-tune it with 250,000 iterations under the same batch size on the shadow removal datasets. Following [10], the input images are resized to 256×256 . We use Adam as the optimizer to optimize the network with a learning rate of 0.00005 and all the experiments are conducted on the Linux server with two NVIDIA Tesla V100 GPUs.

5. Experiments

5.1. Setups

Datasets. We first train our network on the image inpainting dataset (*i.e.*, Places2 challenge dataset [45]) and then fine-tune it on the ISTD+[24] and SRD[35] datasets, respectively. Finally, we evaluate the trained networks on the ISTD+[24] and SRD datasets[35]. The Places2 dataset contains more than millions of images under over 365 scenes. We use the irregular masks in [29]. The ISTD+ dataset[24] is constituted of 1330 triplets for training and 540 triplets for testing. We follow the previous method and use ground truth masks during the training. For the evaluation, we follow [10] using Ostu’s algorithm to calculate the difference between the shadow and shadow-free images to get the masks. The SRD dataset[35] contains 2680 paired shadow and shadow-free images for training and 408 paired shadow and shadow-free images for testing without providing ground truth masks. Therefore, during the training and testing, we use the detected masks provided by DHAN[4].

Metrics. To prove the effectiveness of our method, we follow the previous method [10] to calculate the root mean square error (RMSE) between the reconstructed shadow-free images and the corresponding ground truth in the LAB color space. Besides, we also use the peak signal-to-noise ratio (PSNR[23]), structural similarity index (SSIM), and the learned perceptual image patch similarity (LPIPS[43])

to measure the quality of recovered shadow-free images. To ensure a fair comparison, we reevaluate all baseline methods on the same machine using various metrics.

Baselines. We compare our proposed method with ten state-of-the-art shadow removal methods which includes SP+M-Net[24], DSC[19], DHAN[4], Param+M+D-Net[25], LG-shadowNet[30], DC-ShadowNet[22], G2R-ShadowNet[31], Fu et al.[10], BMNet[47], and SG-ShadowNet[37].

5.2. Comparison Results

Quantitative comparison. We first compare our method with the competitors on the ISTD+ dataset[24]. As shown in Table 1, we can observe that: ① for the whole image, our method outperforms all the competitors over each metric *i.e.* RMSE, PSNR, SSIM, and LPIPS. Specifically, our method decreases RMSE 5.48% and LPIPS 4.41% as well as increases PSNR 5.70% and SSIM 0.58% compared with BMNet[47]. ② for the shadow region our method also outperforms other competitors over all the metrics. Specifically, compared with LG-shadowNet[30] our method decreases RMSE 38.87% as well as increases PSNR 17.79% and SSIM 0.90% respectively. ③ for the non-shadow region, our method still gets the best result. Specifically, compared with Fu et al.[10] our method decreases RMSE 24.17%, increases PSNR 20.19%, and increases SSIM 11.62% respectively.

Besides, we further compare with the competitors on the SRD dataset[35]. As shown in Table 1, we see that: for the shadow region, non-shadow region, and the whole image, our method outperforms all baselines over the four metrics, *i.e.*, RMSE, PSNR, SSIM, and LPIPS. Compared with DC-ShadowNet[22], for the shadow region, our method decreases RMSE 24.80% as well as increases PSNR 7.22% and SSIM 0.91%. For the non-shadow region, our method decreases RMSE 19.08% as well as increases PSNR 7.59% and SSIM 1.39%. For the whole image, our method decreases RMSE 21.68% and LPIPS 30.71% as well as increases PSNR 7.87% and SSIM 3.07%.

Qualitative comparison. We visualize four samples from ISTD+[24] and SRD datasets[35] in Fig. 6. We can observe that ① our method can generate more natural and realistic color in the shadow regions of the restored result. As shown in Case2 and Case4, the generated floor in the shadow regions by our method is almost identical to the ground truth. Compared with other methods *i.e.* BMNet[47] and DSC[19], we can find obvious color inconsistencies between the shadow regions and non-shadow regions in their generated result. ② Although without any post process, the generated result by our method is more smooth along the shadow boundary than other methods generated. As shown in Case1, the boundary in our restored result is almost unidentified by human eyes. However, we can observe the

Table 1: Comparison results on ISTD+[24] and SRD[35] datasets.

Datasets	Method	All				Shadow			Non-Shadow		
		RMSE↓	PSNR↑	SSIM↑	LPIPS↓	RMSE↓	PSNR↑	SSIM↑	RMSE↓	PSNR↑	SSIM↑
ISTD+	SP+M-Net[24]	3.610	32.33	0.9479	0.0716	7.205	36.16	0.9871	2.913	35.84	0.9723
	Param+M+D-Net[25]	4.045	30.12	0.9420	0.0759	9.714	33.59	0.9850	2.935	34.33	0.9723
	SynShadow[21]	4.000	-	-	-	6.900	-	-	3.400	-	-
	Fu et al.[10]	4.278	29.43	0.8404	0.1673	6.583	36.41	0.9769	3.827	31.01	0.8755
	LG-ShadowNet[30]	4.402	29.20	0.9335	0.0920	9.709	32.65	0.9806	3.363	33.36	0.9683
	DC-ShadowNet[22]	4.781	28.76	0.9219	0.1112	10.434	32.20	0.9758	3.674	33.21	0.9630
	G2R-ShadowNet[31]	3.970	30.49	0.9330	0.0868	8.872	34.01	0.9770	3.010	34.62	0.9707
	BMNet[47]	3.595	32.30	0.9551	0.0567	6.189	37.30	0.9899	3.087	35.06	0.9738
	SG-ShadowNet[37]	3.531	32.41	0.9524	0.0594	6.019	37.41	0.9893	3.044	34.95	0.9725
	Ours	3.398	34.14	0.9606	0.0542	5.935	38.46	0.9894	2.902	37.27	0.9772
SRD	DSC[19]	5.704	29.01	0.9044	0.1145	8.828	34.20	0.9702	4.509	31.85	0.9555
	DHAN[4]	4.666	30.67	0.9278	0.0792	7.771	37.05	0.9818	3.486	32.98	0.9591
	SynShadow[21]	5.200	-	-	-	10.90	-	-	3.600	-	-
	Fu et al.[10]	6.269	27.90	0.8430	0.1820	8.927	36.13	0.9742	5.259	29.43	0.8888
	DC-ShadowNet[22]	4.893	30.75	0.9118	0.1084	8.103	36.68	0.9759	3.674	33.10	0.9540
	BMNet[47]	4.240	31.88	0.9376	0.0817	6.982	37.41	0.9816	3.198	35.09	0.9676
	SG-ShadowNet[37]	4.297	31.31	0.9273	0.0835	7.564	36.55	0.9807	3.056	34.23	0.9611
	Ours	3.832	33.17	0.9398	0.0751	6.094	39.33	0.9848	2.973	35.61	0.9673




Figure 6: Visualization results on ISTD+ dataset[24] (first two rows) and SRD dataset[35] (last two rows)

boundary clearly in the restored result by other methods. Also in Case3, our restored result is almost identical to the ground truth. But we can find a clear ghost along the shadow boundary in restored results by other methods.

5.3. Ablation Study

To prove the effectiveness of each part of our method, we conduct the following experiments for ablations: **Exp1**, we remove the adaptive fusion block and combine the extracted features $\hat{\mathbf{F}}$ and \mathbf{F} directly (See the first row of Table 2). We pretrain this variant on the image inpainting dataset and fine-tune that on the shadow dataset. **Exp2**, we concatenate all the inputs and input that into a naive encoder-decoder (See the second row of Table 2). We also perform the pretraining and fine-tuning like exp1. **Exp3**, we pretrain a naive encoder-decode network on the image inpainting dataset and fine-tune it on the shadow dataset like Exp1 and Exp2. Mean-

while, we also train another naive encoder-decode directly on the shadow dataset (See the third and fourth row of Table 2). **Exp4**, we train an inpainting model based on the naive encoder-decoder architecture, then input the shadow-masked image into the inpainting model to perform the shadow removal. As shown in Table 2, we have the following observations: ❶ without the adaptive fusion the performance decreases a lot, *e.g.* the RMSE decreases 14.33% and 2.65% in the shadow and non-shadow regions respectively. ❷ Both solutions of concatenating all the inputs and performing the shadow removal by an inpainting model get a worse result in the shadow and non-shadow regions compared with our method. Specifically, concatenating all the input decreases the RMSE 12.15% in the shadow regions and 1.93% in the non-shadow regions. Performing the shadow removal by an inpainting model decreases the RMSE 59.46% in the shadow regions and 32.48% in the non-shadow regions.

Table 2: Ablation study on ISTD+ dataset[24].

Method	All				Shadow			Non-Shadow		
	RMSE↓	PSNR↑	SSIM↑	LPIPS↓	RMSE↓	PSNR↑	SSIM↑	RMSE↓	PSNR↑	SSIM↑
w/o adaptive fusion	3.627	33.57	0.9573	0.0603	6.928	37.68	0.9884	2.981	36.87	0.9756
Concatenate inputs	3.581	33.58	0.9573	0.0628	6.756	37.58	0.9879	2.959	37.03	0.9763
Encoder-decode w pretraining	3.527	33.59	0.9589	0.0603	6.377	37.66	0.9886	2.969	36.96	0.9769
Encoder-decode w/o pretraining	4.216	32.75	0.9545	0.0661	7.428	36.89	0.9877	3.588	36.09	0.9744
De-shadow by inpainting	5.992	28.62	0.8706	0.1420	14.64	30.98	0.9227	4.298	34.58	0.9672
Ours	3.398	34.14	0.9606	0.0542	5.935	38.46	0.9894	2.902	37.27	0.9772

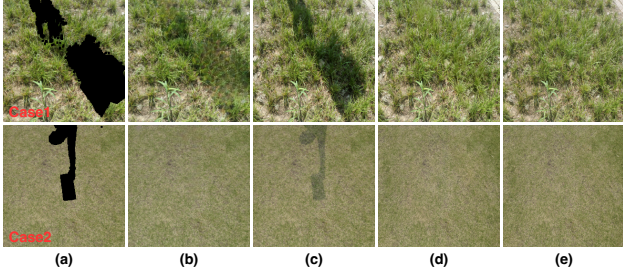


Figure 7: (a) and (b) are cropped images and restored results of [27]. (c) and (d) are shadow images and restored results of our method. (e) is the corresponding ground truth.

5.4. Relationship to SOTA Image Inpainting

For the single-image inpainting task, the input is a corrupted image as shown in Fig. 7 (a) which can be got by multiplying the original image with a binary mask. The masked regions are generated totally based on the unmasked background. For the shadow removal task, the input is a shadow image as shown in Fig. 7 (c) and the shadow regions is indicated by a binary mask. When the binary mask is small or thin *i.e.* (a) and (c) in Case2, the restored result of inpainting is indistinguishable compared with the restored result of shadow removal (see (b) and (d) in Case2). However, when the binary mask is large *i.e.* (a) and (c) in Case1, the texture information of the restored result based on inpainting is worse than the shadow removal-based counterpart. This is because the shadow regions *i.e.* (c) can provide more detailed texture information than the masked regions *i.e.* (a). But from (b) in Case1, we can still find that the restored color in the masked regions is similar to the background. This phenomenon proves that the inpainting method has the capability to transfer the color from the background to the masked regions even when the masked regions are large. Based on the above observations, we can find the potential mutual promotion between image inpainting and shadow removal *i.e.* image inpainting and shadow removal can provide color and textual information respectively.

5.5. Effectiveness of the inpainting branch and the shadow removal branch

Based on our network architecture (Fig. 4), we zero out the outputs of the first shadow removal encoder while preserving the second encoder and the decoder for the inpainting task. Remarkably, without any retraining or fine-tuning, the inpainting branch successfully fills in missing pixels (See Fig. 8(c) vs.(d)) with consistent and smooth colors across the affected regions. Similarly, we can remove the inpainting




Figure 8: From left to right: shadow image, restored result by shadow removal branch, masked image, restored result by inpainting branch, restored result by two branches.

Table 3: Pretraining our network on three datasets with 10,000 images randomly sampled from each dataset. For the last row, we use full Place2 for pretraining.

Method	All				Shadow	Non-Shadow
	RMSE↓	PSNR↑	SSIM↑	LPIPS↓	RMSE↓	RMSE↓
w/o pretrain	3.860	33.30	0.9555	0.0599	7.144	3.263
CelabA (10 ⁴ images)	3.557	33.56	0.9586	0.0559	6.343	3.012
Paris-Street-View (10 ⁴ images)	3.511	33.64	0.9586	0.0551	6.286	2.967
Places2 (10 ⁴ images)	3.476	33.76	0.9601	0.0550	6.232	2.937
Places2 (1.8 × 10⁶ images)	3.398	34.14	0.9606	0.0542	5.935	2.902

branch to focus solely on shadow removal. As depicted in Fig. 8(a) vs. (b), the colors in the shadow region are recovered and resemble those of shadow-free regions in (a), albeit with some noticeable boundary effects.

5.6. Effects of inpainting datasets.

We pretrain our network on three inpainting datasets (See Table 3), respectively, and fine-tune on ISTD+ dataset. In Table 3, we see that: ❶ With different datasets, inpainting pretraining yields significant improvements in shadow removal performance. ❷ Pretraining with a larger dataset leads to better results.

6. Conclusions

In this work, we found that pretraining a network on the image inpainting dataset and fine-tuning it on the paired shadow & shadow-free images can enhance the network’s capability for shadow removal significantly. We conducted extensive experiments to analyze and compare the networks with or without inpainting pretraining. Inspired by these studies, we formulated shadow removal as an adaptive fusion task and propose a novel adaptive fusion network for shadow removal. The extensive experiments show that our method empowered with inpainting outperforms all state-of-the-art methods.

Acknowledgement. This work is supported by the A*STAR Centre for Frontier AI Research and the National Research Foundation, Singapore, and DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG2-GC-2023-008).

References

- [1] Steve Bako, Thijs Vogels, Brian McWilliams, Mark Meyer, Jan Novák, Alex Harvill, Pradeep Sen, Tony Derose, and Fabrice Rousselle. Kernel-predicting convolutional networks for denoising monte carlo renderings. *ACM Trans. Graph.*, 36(4):97–1, 2017. [2](#)
- [2] Zipei Chen, Chengjiang Long, Ling Zhang, and Chunxia Xiao. Canet: A context-aware network for shadow removal. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4743–4752, 2021. [2](#)
- [3] Yupeng Cheng, Felix Juefei-Xu, Qing Guo, Huazhu Fu, Xiaofei Xie, Shang-Wei Lin, Weisi Lin, and Yang Liu. Adversarial exposure attack on diabetic retinopathy imagery. *arXiv preprint arXiv:2009.09231*, 2020. [2](#)
- [4] Xiaodong Cun, Chi-Man Pun, and Cheng Shi. Towards ghost-free shadow removal via dual hierarchical aggregation network and shadow matting gan. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10680–10687, 2020. [6](#), [7](#)
- [5] Tingliang Feng, Wei Feng, Weiqi Li, and Di Lin. Cross-image context for single image inpainting. *Advances in Neural Information Processing Systems*, 35:1474–1487, 2022. [3](#)
- [6] Graham D Finlayson, Mark S Drew, and Cheng Lu. Entropy minimization for shadow removal. *International Journal of Computer Vision*, 85(1):35–57, 2009. [2](#)
- [7] Graham D Finlayson, Steven D Hordley, Cheng Lu, and Mark S Drew. On the removal of shadows from images. *IEEE transactions on pattern analysis and machine intelligence*, 28(1):59–68, 2005. [2](#)
- [8] Lan Fu, Qing Guo, Felix Juefei-Xu, Hongkai Yu, Wei Feng, Yang Liu, and Song Wang. Benchmarking shadow removal for facial landmark detection and beyond. *arXiv preprint arXiv:2111.13790*, 2021. [1](#)
- [9] Lan Fu, Hongkai Yu, Xiaoguang Li, Craig P Przybyla, and Song Wang. Deep learning for object detection in materials-science images: A tutorial. *IEEE Signal Processing Magazine*, 39(1):78–88, 2021. [1](#)
- [10] Lan Fu, Changqing Zhou, Qing Guo, Felix Juefei-Xu, Hongkai Yu, Wei Feng, Yang Liu, and Song Wang. Auto-exposure fusion for single-image shadow removal. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10571–10580, 2021. [2](#), [4](#), [6](#), [7](#)
- [11] Han Gong and Darren Cosker. Interactive removal and ground truth for difficult shadow scenes. *JOSA A*, 33(9):1798–1811, 2016. [2](#)
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. [1](#)
- [13] Maciej Gryka, Michael Terry, and Gabriel J Brostow. Learning to remove soft shadows. *ACM Transactions on Graphics (TOG)*, 34(5):1–15, 2015. [2](#)
- [14] Qing Guo, Felix Juefei-Xu, Xiaofei Xie, Lei Ma, Jian Wang, Bing Yu, Wei Feng, and Yang Liu. Watch out! motion is blurring the vision of your deep neural networks. *Advances in Neural Information Processing Systems*, 33, 2020. [2](#)
- [15] Qing Guo, Xiaoguang Li, Felix Juefei-Xu, Hongkai Yu, Yang Liu, and Song Wang. Jpgnet: Joint predictive filtering and generative network for image inpainting. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 386–394, 2021. [2](#)
- [16] Qing Guo, Jingyang Sun, Felix Juefei-Xu, Lei Ma, Xiaofei Xie, Wei Feng, Yang Liu, and Jianjun Zhao. Efficientderain: Learning pixel-wise dilation filtering for high-efficiency single-image deraining. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1487–1495, 2021. [2](#)
- [17] Ruiqi Guo, Qieyun Dai, and Derek Hoiem. Paired regions for shadow detection and removal. *IEEE transactions on pattern analysis and machine intelligence*, 35(12):2956–2967, 2012. [2](#)
- [18] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. [2](#)
- [19] Xiaowei Hu, Chi-Wing Fu, Lei Zhu, Jing Qin, and Pheng-Ann Heng. Direction-aware spatial context features for shadow detection and removal. *IEEE TPAMI*, 42(11):2795–2808, 2019. [6](#), [7](#)
- [20] Xiaowei Hu, Yitong Jiang, Chi-Wing Fu, and Pheng-Ann Heng. Mask-shadowgan: Learning to remove shadows from unpaired data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2472–2481, 2019. [2](#)
- [21] Naoto Inoue and Toshihiko Yamasaki. Learning from synthetic shadows for shadow detection and removal. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(11):4187–4197, 2020. [7](#)
- [22] Yeying Jin, Aashish Sharma, and Robby T Tan. Dc-shadownet: Single-image hard and soft shadow removal using unsupervised domain-classifier guided network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5027–5036, 2021. [2](#), [6](#), [7](#)
- [23] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. [6](#)
- [24] Hieu Le and Dimitris Samaras. Shadow removal via shadow image decomposition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8578–8587, 2019. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [25] Hieu Le and Dimitris Samaras. From shadow segmentation to shadow removal. In *European Conference on Computer Vision*, pages 264–281. Springer, 2020. [2](#), [6](#), [7](#)
- [26] Hieu Le and Dimitris Samaras. Physics-based shadow image decomposition for shadow removal. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (01):1–1, 2021. [4](#)
- [27] Xiaoguang Li, Qing Guo, Di Lin, Ping Li, Wei Feng, and Song Wang. Misf: Multi-level interactive siamese filtering for high-fidelity image inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1869–1878, 2022. [2](#), [3](#), [6](#), [8](#)

- [28] Di Lin, Xin WANG, Jia Shen, Renjie Zhang, Ruonan Liu, Miaohui Wang, Wuyuan Xie, Qing Guo, and Ping Li. Generative status estimation and information decoupling for image rain removal. *Advances in Neural Information Processing Systems*, 35:4612–4625, 2022. [2](#)
- [29] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *European Conference on Computer Vision (ECCV)*, pages 85–100, 2018. [3](#), [6](#)
- [30] Zhihao Liu, Hui Yin, Yang Mi, Mengyang Pu, and Song Wang. Shadow removal by a lightness-guided network with training on unpaired data. *IEEE Transactions on Image Processing*, 30:1853–1865, 2021. [2](#), [6](#), [7](#)
- [31] Zhihao Liu, Hui Yin, Xinyi Wu, Zhenyao Wu, Yang Mi, and Song Wang. From shadow generation to shadow removal. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4927–4936, 2021. [2](#), [6](#), [7](#)
- [32] Li-Qian Ma, Jue Wang, Eli Shechtman, Kalyan Sunkavalli, and Shi-Min Hu. Appearance harmonization for single image shadow removal. In *Computer Graphics Forum*, volume 35, pages 189–197. Wiley Online Library, 2016. [2](#)
- [33] Ben Mildenhall, Jonathan T Barron, Jiawen Chen, Dillon Sharlet, Ren Ng, and Robert Carroll. Burst denoising with kernel prediction networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2502–2510, 2018. [2](#)
- [34] Sohail Nadimi and Bir Bhanu. Physical models for moving shadow and object detection in video. *IEEE transactions on pattern analysis and machine intelligence*, 26(8):1079–1087, 2004. [1](#)
- [35] Liangqiong Qu, Jiandong Tian, Shengfeng He, Yandong Tang, and Rynson WH Lau. Deshadownet: A multi-context embedding deep network for shadow removal. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4067–4075, 2017. [6](#), [7](#)
- [36] Andres Sanin, Conrad Sanderson, and Brian C Lovell. Improved shadow removal for robust person tracking in surveillance scenarios. In *ICPR*, pages 141–144, 2010. [1](#)
- [37] Jin Wan, Hui Yin, Zhenyao Wu, Xinyi Wu, Yanting Liu, and Song Wang. Style-guided shadow removal. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. [2](#), [4](#), [6](#), [7](#)
- [38] Zifeng Wang, Shao-Lun Huang, Ercan E Kuruoglu, Jimeng Sun, Xi Chen, and Yefeng Zheng. Pac-bayes information bottleneck. *arXiv preprint arXiv:2109.14509*, 2021. [2](#), [4](#)
- [39] Chunxia Xiao, Ruiyun She, Donglin Xiao, and Kwan-Liu Ma. Fast shadow removal using adaptive multi-scale illumination transfer. In *Computer Graphics Forum*, volume 32, pages 207–218. Wiley Online Library, 2013. [2](#)
- [40] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9653–9663, 2022. [2](#)
- [41] Qingxiong Yang, Kar-Han Tan, and Narendra Ahuja. Shadow removal using bilateral filtering. *IEEE Transactions on Image processing*, 21(10):4361–4368, 2012. [2](#)
- [42] Liming Zhai, Felix Juefei-Xu, Qing Guo, Xiaofei Xie, Lei Ma, Wei Feng, Shengchao Qin, and Yang Liu. It’s raining cats or dogs? adversarial rain attack on dnn perception. *arXiv preprint arXiv:2009.09205*, 2020. [2](#)
- [43] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–595, 2018. [6](#)
- [44] Wuming Zhang, Xi Zhao, Jean-Marie Morvan, and Liming Chen. Improving shadow suppression for illumination robust face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 41(3):611–624, 2018. [1](#)
- [45] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6):1452–1464, 2017. [3](#), [6](#)
- [46] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. [2](#)
- [47] Yurui Zhu, Jie Huang, Xueyang Fu, Feng Zhao, Qibin Sun, and Zheng-Jun Zha. Bijective mapping network for shadow removal. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5627–5636, 2022. [2](#), [6](#), [7](#)