

ARGAN: Attentive Recurrent Generative Adversarial Network for Shadow Detection and Removal

Bin Ding¹, Chengjiang Long^{2*}, Ling Zhang³, Chunxia Xiao^{1*}

¹School of Computer Science, Wuhan University, Wuhan, Hubei, China

²Kitware Inc., Clifton Park, NY, USA

³Wuhan University of Science and Technology, Wuhan, Hubei, China

dingbin@whu.edu.cn, chengjiang.long@kitware.com, zhling@wust.edu.cn, cxxiao@whu.edu.cn

Abstract

In this paper we propose an attentive recurrent generative adversarial network (ARGAN) to detect and remove shadows in an image. The generator consists of multiple progressive steps. At each step a shadow attention detector is firstly exploited to generate an attention map which specifies shadow regions in the input image. Given the attention map, a negative residual by a shadow removal encoder will recover a shadow-lighter or even a shadow-free image. A discriminator is designed to classify whether the output image in the last progressive step is real or fake. Moreover, ARGAN is suitable to be trained with a semi-supervised strategy to make full use of sufficient unsupervised data. The experiments on four public datasets have demonstrated that our ARGAN is robust to detect both simple and complex shadows and to produce more realistic shadow removal results. It outperforms the state-of-the-art methods, especially in detail of recovering shadow areas.

1. Introduction

Shadow exists in most images and is formulated in interaction among the illumination, object materials and scene geometry. Obviously, the detected shadows can provide important clues for various applications of visual scene understanding, such as scene geometry depiction [41], camera location [21], object relighting [22], and scene illumination inference [60]. Meanwhile, the shadow removal is able to boost the performance of some computer vision and computer graphics tasks, such as object detection and tracking [37, 33, 36], object recognition [5, 17, 31, 32, 18], intrinsic image decomposition [29]. Therefore, it is desirable to develop an effective method of shadow detection and removal.

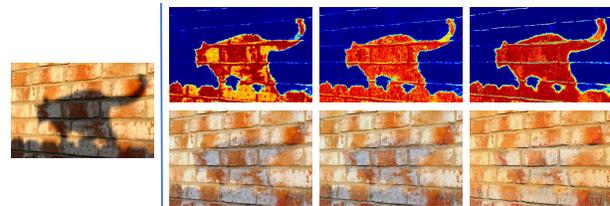


Figure 1. Given an input image (left) with shadow, our goal is to generate a more accurate attention map gradually indicating the detected shadow region (right-top) and recover a more realistic shadow-removal image (right-bottom) gradually with multiple progressive steps.

Previous works for shadow detection and removal can be mainly divided into two categories. One category is traditional methods [12, 61] based on some prior knowledge, such as consistent illumination in shadow regions. Its shortcoming is that the priors might dissatisfy some shadow images under a complex environment so that the performance of shadow removal result cannot be guaranteed. The other is deep learning methods [44, 15, 53, 26, 65] whose effectiveness rely heavily on supervised data to learn a robust model. Particularly, however, when the training data is insufficient, such deep learning methods often appear color distortion or other problems in the shadow removal results.

In this paper, we propose a novel attentive recurrent generative adversarial network (ARGAN) for shadow detection and removal. As illustrated in Figure 2, the generator involves multiple progressive step for shadow detection and removal in a coarse-to-fine fashion, and the discriminator is designed to classify whether the generated shadow-removal image at the last step from the generator is real or fake. At each progressive step in the generator, a shadow attention detector is used to generate an attention map. Then a shadow removal encoder is designed to combine the previous shadow-removal image and the current detected shadow attention map to obtain a negative residual [8] for recover-

*This work was co-supervised by Chengjiang Long and Chunxia Xiao.

ing the shadow-lighter or even shadow-free images.

The intuition behind multiple progressive steps in the generator is that it is much easier to detect and remove shadow gradually with a step-by-step approach so that we are able to handle shadows especially with complicated scenes. The detected shadow region and recovered shadow-lighter image from the previous step are the input of the present step. Therefore, we are able to explore a recurrent unit such as Long and Short Term Memory (LSTM) [14] to reserve the valuable and detailed information to make sure that the detected shadow regions are more and more accurate, and that the shadow-removal images are more and more realistic, as illustrated in Figure 1.

We shall emphasize that we adopt the adversarial training process [10] between the generator and the discriminator to generate a shadow-removal image. With the number of epochs increases, both models improve their functionalities so that it becomes harder and harder to distinguish a generated shadow-removal image from a real shadow-free image. Therefore, after a certain large number of training epochs, we can utilize the learned parameters in the generator to generate a shadow attention map and a shadow-removal image at each progressive step. The output from the last step is our final result.

Moreover, we apply a semi-supervised learning strategy [47] to make full use of sufficient unsupervised shadow images available online by modifying the original adversarial loss to cover both labeled data and unlabeled data. We can first use the generator and generate a shadow-removal image for any input images with shadow, and then just use the discriminator to discriminate whether the generated image is real or not. This treatment can improve the generalization ability and robustness of our ARGAN.

Several aspects distinguish our work from the previous shadow detection and removal methods [12, 61, 44, 15, 53, 26, 65]. First of all, our proposed ARGAN adopts adversarial training process to optimize each shadow attention detector and each shadow removal encoder in the generator. Secondly, the generator involves multiple progressive steps for shadow detection and removal in a coarse-to-fine fashion so that it can handle shadows with complicated environment. Thirdly, a semi-supervised strategy by incorporating sufficient unsupervised shadow images available online is able to increase the robustness of our network. We evaluate our proposed ARGAN on the four public datasets and compare with the state-of-the-art methods on the performance of both shadow detection and shadow removal. The results clearly demonstrate the efficacy of our proposed model.

2. Related Work

The related work can be divided into four categories: *shadow detection methods*, *shadow removal methods*, *Generative Adversarial Network* and *attention mechanism*.

Shadow detection methods involve traditional methods, using user interactions [11, 61, 6] and hand-crafted features [25, 12, 51], and recent deep learning methods [23, 52, 40, 16, 26, 65] for automatic shadow detection. To specify, Khan *et al.* [23] detected the shadow by combining the boundary and region ConvNets in the CRF model. Vicente *et al.* [52] proposed a semantic-aware patch level CNN architecture for shadow detection. Nguyen *et al.* [40] detected shadow using conditional generative adversarial networks. Hu *et al.* [16] detected shadow by analyzing image context in a direction-aware manner. However, these methods only work well on image with simple shadow. They cannot detect accurate shadow with complex scenes.

Shadow removal methods based on gradient domain manipulation [39, 7], illumination [61, 56, 55, 23, 46], color transferring [45], accurate shadow matte [4, 11, 54] and depth information [57] have been exposed for a long time. Recently three deep learning methods have been proposed for shadow removal. One is Qu *et al.*'s multi-context embedding network [44] integrating high-level semantic context for shadow removal. One is Hu *et al.*'s [15] using direction-aware spatial context features for shadow detection and removal. Another is Wang *et al.*'s GAN-based method [53] which jointly learns shadow detection and shadow removal. Different from [53], our proposed ARGAN involves multiple progressive steps with attentive recurrent units in the generator to achieve better performance on shadow removal.

Generative Adversarial Network (GAN) [10] and its variants [60] have been proposed to deal with various image-to-image translation problems, such as image super-resolution [27], image inpainting [42], style transfer [28] and domain adaptation/transfer [30, 19, 50], raindrop removal [43], shadow detection and removal [53]. Unlike [43] which only progressively updates the attention map with the same input image for a one-step removal, our proposed ARGAN progressively detects shadow and removes shadow step by step in a coarse-to-fine fashion.

Attention mechanism [2] is designed to encode sequence data based on the assigned importance score of each element, which has attained significant improvement in various tasks in natural language processing [24, 49], speech recognition [3], computer vision [63], image captioning [59, 34, 9] and visual question answering (VQA) [35, 1]. Different from [63] which uses the progressive and recurrent idea to integrate multiple contextual information of multi-level features, our ARGAN progressively and recurrently updates the shadow attention maps and the shadow removal images in the generator in a coarse-to-fine fashion so that it can handle shadows with complicated environment.

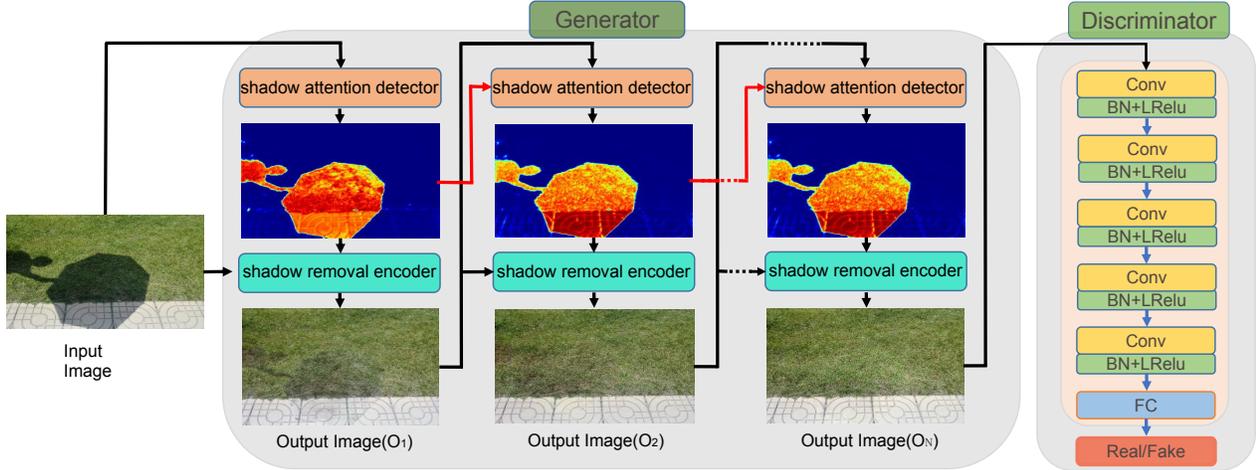


Figure 2. The framework overview of the proposed ARGAN consists of two components, *i.e.*, a generator and a discriminator. The generator consists of N shadow attention detectors and N shadow removal encoders. Each shadow attention detector is designed to generate the shadow attention map, and each shadow removal encoder is to produce shadow-lighter or even shadow-free image. The discriminator is formed by five convolutional layers and a fully connected layer to classify the output shadow-free image as real or fake.

3. Approach

As illustrated in Figure 2, we present an attentive recurrent generative adversarial network (ARGAN) to explore the mapping relationship from shadow images to the corresponding shadow-free images. Like all the generative adversarial networks, our ARGAN contains two components, *i.e.*, a generator to produce a shadow-free image as real as possible, and a discriminator to classify whether the generated shadow-free image is indeed a real image or not.

At the generative stage, given an input shadow image I , we iteratively update the detected shadow region indicated with an attention map A_i by shadow attention detector G_{det}^i and output a shadow-lighter or even shadow-free image O_i by shadow removal encoder G_{rem}^i at i -th step by the following equations:

$$A_i = \begin{cases} G_{det}^i(I) & i = 1 \\ G_{det}^i(O_{i-1}, A_{i-1}) & i > 1 \end{cases} \quad (1)$$

$$O_i = \begin{cases} G_{rem}^i(I, A_i) & i = 1 \\ G_{rem}^i(O_{i-1}, A_i) & i > 1 \end{cases} \quad (2)$$

At the discriminative stage, we design a discriminator D to encode the final output shadow-free image O_N with semi-supervised strategy to handle the supervised data (the ground-truth shadow-free images F) and unsupervised data (F is missing) under the adversarial framework.

In the following subsections, we are going to discuss the generative network, the discriminative network, and the loss functions, as well as the implementation details.

3.1. Generative Network

Our generative network is composed of N progressive steps and each step has one shadow attention detector and

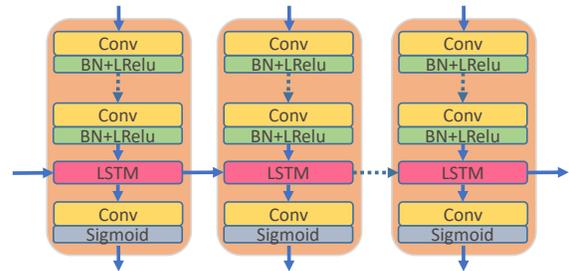


Figure 3. Each shadow attention detector consists of ten convolutional layers with 64 as output channel number, one LSTM layer and one convolutional layer. Note that every convolutional layer here is followed by a batch normalization and Leaky-ReLU active function. The stride is 1.

one shadow removal encoder.

Shadow attention detector. We incorporate attention mechanism to selectively choose what our network wants to observe, locate shadows of the input image and make the attention of shadow removal encoder focusing on the detected shadow regions. As shown in Figure 2, a recurrent unit through Long and Short Term Memory (LSTM) [14] can be formulated into our recurrent attention network, in which LSTM can make full use of the intermediate output of the previous step in the recurrent network and, as a prior, generate attention map which represents shadow region in the later steps.

As illustrated in Figure 3, the shadow attention detector in our recurrent network is designed as ten convolutional layers with batch normalization and Leaky ReLU activation function (Conv+BN+LRelu) to extract feature from the input image. A convolutional LSTM unit [58] is to estimate

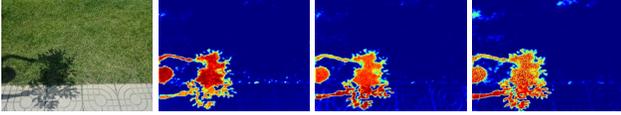


Figure 4. Attention maps generated by shadow attention detector at the first three progressive iterations. From left to right are input shadow image, region attention maps A_1 , A_2 and A_3 , respectively.

the shadow regions by combining information at the previous step and passing the estimated result to the next step and a convolutional layer to generate a shadow attention map, which is the shadow matte in Equation 1.

Note that all the N shadow attention detectors share the same architecture. Each output of attention map A_i is a matrix. Each value in such a matrix is in the range from 0 to 1, rather than a binary mask. The larger the value, the more attention at this area. It indicates that the area of large value is more likely to be marked as shadow region. As shown in Figure 4, the red area with greater attention value close to 1 is more likely to be shadow region, whereas the blue area where the attention value is close to 0 is more likely to be considered as non-shadow region. In this way, our attention map can effectively distinguish soft shadow and hard shadow in the image by giving different attention values.

We observe that the attention module may not focus on all shadow regions of the image initially. The attention will constantly and gradually cover to the target region in the subsequent recurrent iterations. Therefore, it gradually shifts the attention to focus on all shadow regions. Figure 4 shows one example of attention maps generated at the first three progressive steps. The detected shadow region marked in red are more and more accurate as the step number increases.

Shadow removal encoder. In order to remove shadow in an image, we follow the idea of negative residual [8] and design the shadow removal encoder, which incorporates both the image with shadow, and the detected shadow regions attention to generate a negative residual for recovering a shadow-lighter or shadow-free image.

As illustrated in Figure 5, the encoder firstly uses eight Conv+BN+LRelu to extract feature from the image. Then it takes eight deconvolutional layers with batch normalization and Leaky ReLU activation function (Deconv+BN+LRelu) to generate image with feature data of a particular distribution. Skip connection [13] is applied between convolutional layers and deconvolutional layers because it is able to not only increase the number of channels in the network, but also preserve the context information of front layer.

Following the last deconvolutional layer, 2 Conv+BN+LRelu are applied to extract the feature map and a convolutional layer with sigmoid activation to convert the feature map into a corresponding map with 3

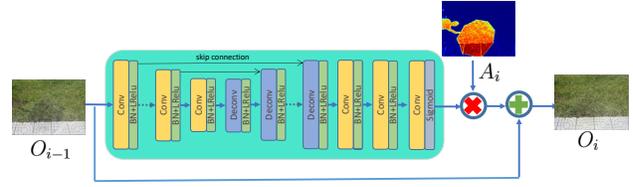


Figure 5. The architecture of the shadow removal encoder. It consists of 8 Conv+BN+LRelu and 8 Deconv+BN+LRelu. The skip connections are linked between convolutional layers and deconvolutional layers. After that, 2 Conv+BN+LRelu, 1 Conv and 1 Sigmoid layer generates a corresponding map with 3 channels and the same size of the input O_{i-1} . A product operation is applied between the sigmoid output and the detected attention map A_i . Then a negative residual is obtained to recover a shadow-lighter or even shadow-free image O_i .

Layer	Output	Kernel	Stride	Layer	Output	Kernel	Stride
Conv	64	3×3	2	Deconv	512	4×4	2
Conv	128	3×3	2	Deconv	512	4×4	2
Conv	256	3×3	2	Deconv	512	4×4	2
Conv	512	3×3	2	Deconv	256	4×4	2
Conv	512	3×3	2	Deconv	128	4×4	2
Conv	512	3×3	2	Deconv	64	4×4	2
Conv	512	3×3	2	Deconv	3	4×4	2
Conv	512	3×3	2	Conv	3	3×3	1
Deconv	512	4×4	2	Conv	3	3×3	1

Table 1. The architecture for shadow removal encoder. Conv means the convolutional layer, Deconv means the deconvolutional layer. Output channels denotes for the amount of output channels in current layer. Kernel means the convolutional kernel size. Stride denotes the moving step size of convolutional kernel.

channels and the same size of the input O_{i-1} . The parameters for all the convolutional layers and deconvolutional layers are summarized in Table 1. Finally, we convert the sigmoid output to the residual information with the detected attention map A_i by a product operation to get a negative residual for recovering a shadow-lighter or even a shadow-free image O_i from the input image O_{i-1} .

As we can observe in Figure 6, as the progressive step number increases, the shadow in the output image becomes lighter and lighter, and the last output image O_3 is almost shadow-free.

3.2. Discriminative Network

Discriminator is designed as a binary classifier to predict whether the final output image O_N from the generator is real or fake. It is worth mentioning that both generator and discriminator constantly improve their ability. Finally, they achieve a balanced state that the image produced by generator seems to be a real shadow-free image, which is consistent with our expectation that our generator model can produce a realistic shadow-free image so that the discriminator may consider it as a real shadow-free image.

To make it simple, we design our discriminator with five Conv+BN+LRelu and one fully connected layer, as shown



Figure 6. Output images generated by shadow removal encoder at the first three progressive iterations. From left to right are input shadow image, output images O_1 , O_2 and O_3 , respectively.

in Figure 2. The output channel numbers for all these six layer are 64, 128, 256, 512 and 1, respectively. Note that the kernel size for convolutional layers is 4×4 , and the stride length is 2. The last fully connected layer outputs the actual probability value for the input image. The fake image and the real image are distinguished by calculating the cross entropy loss between them.

It worths paying attention to that some shadow scenes may be missed in the shadow datasets. To solve this problem, inspired by [47], we apply semi-supervised learning strategy to our network. We use shadow image without ground-truth as unsupervised data in the training process. For each training process, we also input an unsupervised data to the generator and generate a shadow-free image. The discriminator discriminates whether the generated image is real or not. The semi-supervised strategy can improve the generalization ability of our network and make our generator models more robust.

In addition, to make our ARGAN more stable, we use the latest spectral normalization [38] method to stabilize the training process of discriminator network, because spectral normalization is a simple and effective standardized method for limiting the optimization process of the discriminator in GAN, and it can make the whole Generator model better.

3.3. Loss functions

The loss function we use to optimize our ARGAN comes from the shadow attention detector, shadow removal encoder, and discriminator. The total loss L_{total} can be formulated as:

$$L_{total} = L_{det} + L_{rem} + L_{adv}, \quad (3)$$

where the corresponding loss components are described as following.

Shadow attention detector loss L_{det} in each shadow attention detector is defined as the mean square error (MSE) between shadow matte M (which is obtained by comparing the ground-truth shadow-free image and its corresponding shadow image) and the output attention map. In our generator, we apply the shadow attention detector for N iterations, and the loss function is expressed as:

$$L_{det} = \sum_{i=1}^N \beta_i V_{MSE}(A_i, M), \quad (4)$$

where β_i is the weight of the MSE loss at i -th iteration, $\beta_i = 0.7^{N-i+1}$, and $V_{MSE}(A_i, M)$ is the mean square error between M and A_i .

Shadow removal encoder loss L_{rem} contains accuracy loss and perceptual loss [20]. We define the loss function as:

$$L_{rem} = \sum_{i=1}^N L_{mse}(O_i, F) + \sum_{i=1}^N L_{per}(O_i, F), \quad (5)$$

where O_i is the shadow-lighter or even shadow-free image generated by shadow removal encoder, F is the corresponding ground-truth shadow-free image, $L_{mse}(O_i, F)$ is accuracy loss, and $L_{per}(O_i, F)$ is perceptual loss.

$L_{mse}(O_i, F)$ is used to measure the difference between the ground-truth image and shadow-free image generated by shadow removal encoder at the i -th iteration. The smaller the value of MSE, the more accurate the shadow removal encoder. The accuracy loss function is defined as:

$$L_{mse}(O_i, F) = \beta_i V_{MSE}(O_i, F). \quad (6)$$

$L_{per}(O_i, F)$ is used to calculate global difference between the ground-truth image and the shadow removal result. We extract image features using the pre-trained VGG16 model [48] on the ImageNet dataset. The loss function is defined as :

$$L_{per}(O_i, F) = V_{MSE}(VGG(O_i), VGG(F)), \quad (7)$$

where $VGG(O_i)$ and $VGG(F)$ are the feature of image O_i and F extracted from VGG16 model.

Adversarial loss L_{adv} with supervised learning is expressed as:

$$L_{adv} = \mathbb{E}_{(I, F)} [\log(D(y)) + \log(1 - D(G(I)))], \quad (8)$$

and with semi-supervised learning, it is defined as

$$L_{adv} = \lambda \mathbb{E}_{(I, F)} [\log(D(y)) + \log(1 - D(G(I)))] + (1 - \lambda) \mathbb{E}_{(\hat{I})} (\log(1 - D(G(\hat{I}))), \quad (9)$$

where \hat{I} is an unsupervised data. G is the generator. The output of the discriminator D represents the probability of input image is the real image. λ is a weighting hyper-parameter and the expectation value is over a supervised training set (I, F) of input-output pairs.

3.4. Implementation Details

Our proposed ARGAN is implemented in Tensorflow on a computer with Intel(R) Xeon(R) Silver 4114 CPU @ 2.20GHz 192G RAM NVIDIA GeForce GTX 1080Ti. In our experiments, the input size of image is 256×256 . We set $N = 3$ and $\lambda = 0.7$. The minibatch size of 4. The initial learning rate is set as 0.0002. We use Momentum Optimizer to optimize our generator and use Adam Optimizer for the discriminator. We alternatively train the generative network and the discriminative network for 100,000 epochs.

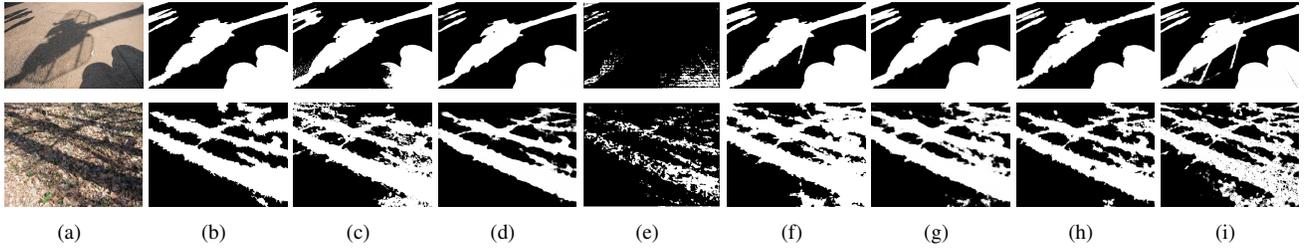


Figure 7. Shadow detection results comparisons. (a) is input images. (b) is results of Guo [12]. (c) is results of Zhang [61]. (d) is results of DSC [15]. (e) is results of ST-CGAN [53]. (f) is results of A+D Net [26]. (g) is results of BDRAR [65]. (h) is ground-truth. (i) is our ARGAN’s results.

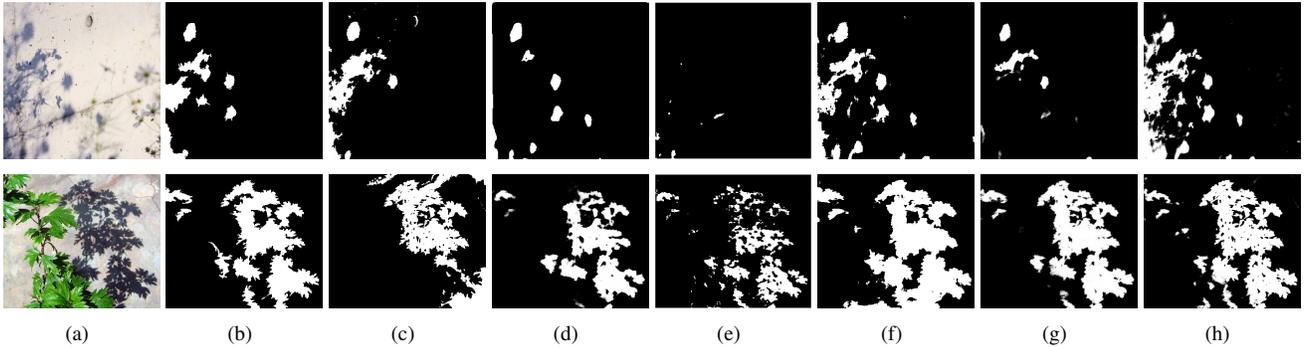


Figure 8. Shadow detection results comparisons. (a) is input images. (b) is results of Guo [12]. (c) is results of Zhang [61]. (d) is results of DSC [15]. (e) is results of ST-CGAN [53]. (f) is results of A+D Net [26]. (g) is results of BDRAR [65]. (h) is our ARGAN’s results.

4. Experiments

To verify the effectiveness of our proposed ARGAN+SS, we compare our method with several state-of-the-art shadow detection and removal methods on four datasets: SBU dataset [52], UCF dataset [64], SRD dataset [44], and ISTD dataset [53].

4.1. Datasets and metrics

SBU dataset [52] contains 4727 pairs of shadow and shadow mask image. UCF dataset [64] contains 110 images with corresponding shadow masks. Both datasets have no shadow-free images so that we can only evaluate the detection performance on them. SRD dataset [44] has 408 pairs of shadow and shadow-free images publicly available but, without the ground-truth shadow masks, can only be used for the evaluation of shadow removal. ISTD dataset [53] contains 1870 image triplets of shadow image, shadow mask, and shadow-free image. We can use this dataset to train our ARGAN and evaluate the performance on both shadow detection and shadow removal. We use 1330 triplets of shadow image, shadow mask and shadow-free images of ISTD [53] train dataset for training as supervised data, and the remaining 540 triplets for evaluation.

For shadow detection, we employ Balance Error Rate (BER) [40] between the ground-truth mask and the predicted shadow matte to evaluate the shadow detection performance. For shadow removal, we utilize the root mean square error (RMSE) in LAB color space between the recovered shadow removal result and the ground-truth image.

4.2. Performance comparison on shadow detection

We compare our shadow detection results with some state-of-the-art shadow detection methods including two traditional methods, *i.e.*, Guo [12] and Zhang [61], and four recent deep learning methods, *i.e.*, ST-CGAN [53], DSC [15], A+D Net [26] and BDRAR [65]. To further verify the effectiveness of the LSTM layer in the shadow attention detector, we remove all the LSTM layers and get a variant network which we call “AGAN”. For the fair comparison, we train all the competing models together with our proposed ARGAN on the same training data in supervised learning, and evaluate the shadow detection performance on the SBU, UCF and ISTD datasets. We also collect 1330 images of a wide variety of scenes from online and take them as unsupervised data to train our model in semi-supervised learning. We denote this method as “ARGAN+SS” where “SS” represents semi-supervised.

The results are summarized in Table 2. As we can observe, (1) among all the competing methods, our ARGAN works the best BER on all the three datasets, which strongly demonstrates that our proposed ARGAN is able to detect accurate shadow regions; (2) without LSTM layers, AGAN performs much worse than ARGAN, which indirectly verify the effectiveness of the LSTM layers in our shadow attention detector; (3) with semi-supervised learning, ARGAN+SS further improve the performance from ARGAN, which strongly proves the robustness of our proposed model.

To further explain the outperforming of our approach, we

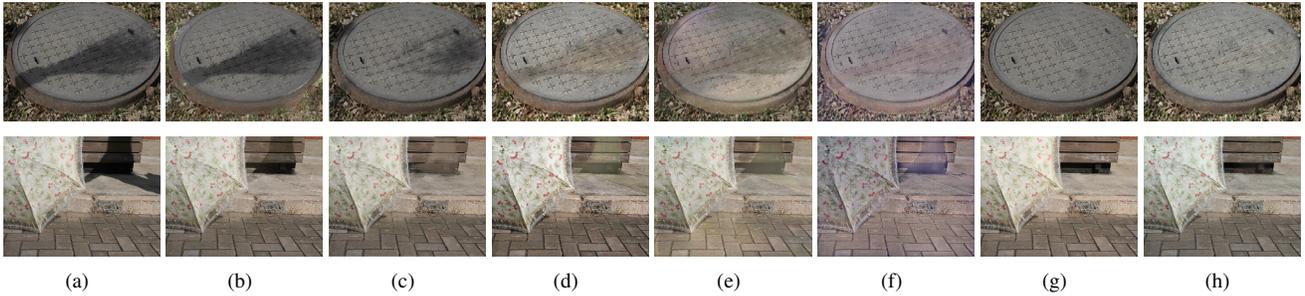


Figure 9. Shadow removal results. (a) is the input images. (b) is results of Guo [12]. (c) is the results of Zhang [61]. (d) is results of DshadowNet [44]. (e) is results of DSC [15]. (f) is results of ST-CGAN [53]. (g) is ground-truth. (h) is our ARGAN’s results.

visualize some results in Figure 7. As we can see, (1) traditional methods of Guo [12] and Zhang [61] are not able to effectively detect slender shadows in the image; (2) among all deep learning methods, comparing with ST-CGAN [53], DSC [15], A+D Net [26] and BDRAR [65], our proposed ARGAN is able to detect more accurate shadow regions and even more close to our human observation.

Figure 8 presents two more shadow images with more complex scenes. Apparently, our ARGAN achieves the best performance on shadow detection. This can be explained by the fact that the shadow attention detectors with recurrent units LSTMs keep updating the detection results gradually from coarse to fine in multiple progressive steps.

	Year	SBU	UCF	ISTD
Guo	2011	25.03	28.32	27.16
Zhang	2015	7.13	9.21	8.56
DSC	2018	5.31	8.73	2.40
ST-CGAN	2018	13.56	17.69	3.84
A+D Net	2018	7.67	11.05	2.97
BDRAR	2018	6.61	9.45	2.20
AGAN	2019	7.24	8.67	4.23
ARGAN	2019	3.09	3.76	2.01
ARGAN+SS	2019	2.56	3.03	1.75

Table 2. Quantitative comparison results on shadow detection with BER metric. The best and second best results are marked in red and blue colors, respectively.

4.3. Performance comparison on shadow removal

We compare our proposed ARGAN with the state-of-the-art methods including the traditional methods, *i.e.*, Guo [12] and Zhang [61] and the recent deep learning methods, *i.e.*, DshadowNet [44], DSC [15], and ST-CGAN [53]. We also compare with two variants of our model, *i.e.*, AGAN and ARGAN+SS. We evaluate the performance of shadow removal on SRD dataset and ISTD dataset.

The results are summarized in Table 3. As we can see, (1) our proposed ARGAN achieves the best RMSE in both shadow regions and the whole image on the two datasets, which suggests that ARGAN is promising to removal shadows and recover more realistic shadow-free images; (2) without LSTM layers, AGAN cannot recover the shadow-free images qualitatively as well as ARGAN. It can

be explained by the fact that LSTM layers affect the detected shadow attention map which affects the quality of final recovered shadow-free images; (3) ARGAN+SS constantly improves the performance for ARGAN no matter in shadow regions, non-shadow regions, or even the whole images. Again, this clearly demonstrates that our proposed ARGAN is good at utilizing sufficient unsupervised shadow images in a semi-supervised learning to improve the quality of the generator and guarantee the performance of both shadow detection and shadow-free image recovery.

We continue to analyze the performance comparison with visualization in Figure 9. In these input images, some areas in non-shadow regions are dark in color. As we can observe, Guo [12], Zhang [61], DshadowNet [44], ST-CGAN [53], DSC [15], sometimes will consider the areas with dark color as shadow regions, and recover illumination in these dark areas. However, the illumination enhancement changes the color of non-shadow regions, which is not desirable. In contrast, our proposed ARGAN takes full account of color information in the whole image and can produce more natural and realistic shadow removal results.

To further verify the robustness of our ARGAN, we choose 7 images with shadows in complex scenes from online and apply all the competing methods to produce shadow-free images, as shown in Figure 10. Obviously, our ARGAN is robust to deal with shadows with complex scenes. The recovered illumination in original shadow regions in the recovered shadow-free image is consistent with surrounding environment and the texture details in shadow regions are well preserved.

4.4. Discussion

To further explore how the value of N affects the final performance, we take experiments $N = 1, 2, 3$ and 4 to generate a shadow-free images by our ARGAN. On ISTD, the BER values for shadow detection are 2.22, 2.08, 2.01 and 2.01, respectively; and the RMSE values for shadow removal on the whole images are 7.35, 6.97, 6.68 and 6.67, respectively. We observe that $N = 3$ is a good trade-off between performance and complexity, as shown in Figure 11.

In addition, we also visualize the shadow-removal results with AGAN, ARGAN and ARGAN+SS in Figure 12. We

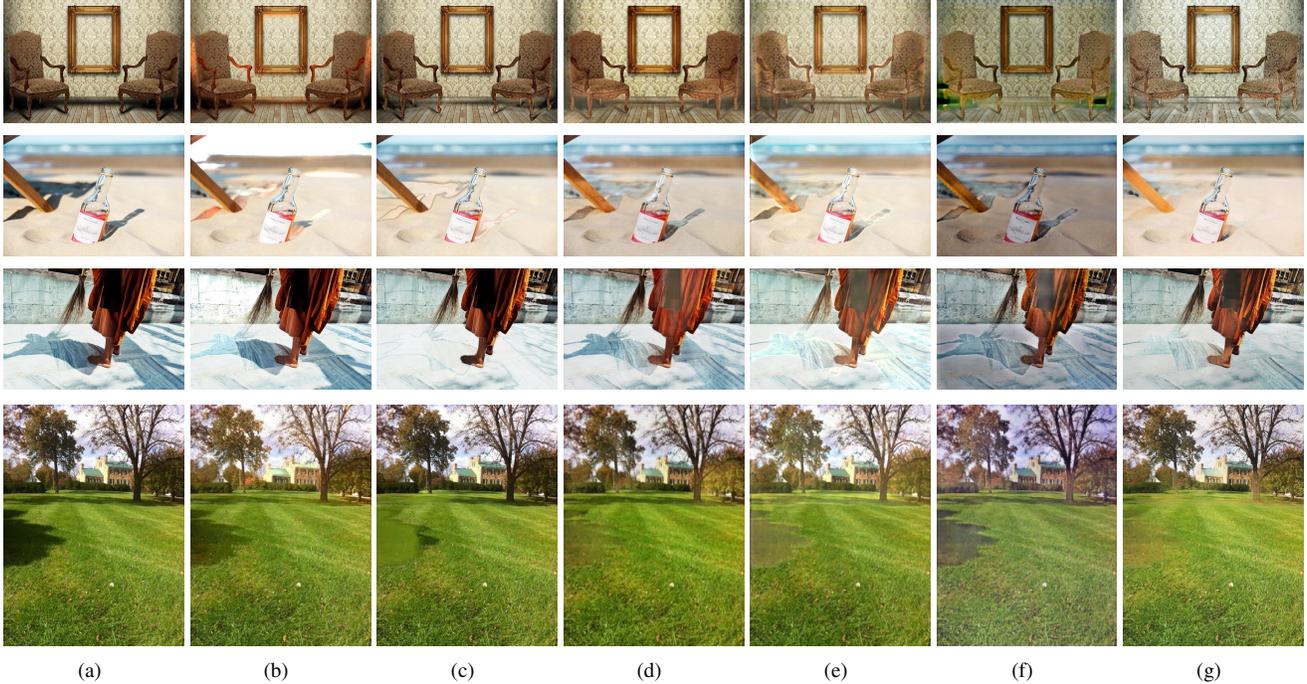


Figure 10. Shadow removal results. (a) is the input images. (b) is results of Guo [12]. (c) is the results of Zhang [61]. (d) is results of DeshadowNet [44]. (e) is results of DSC [15]. (f) is results of ST-CGAN [53]. (g) is our ARGAN's results.

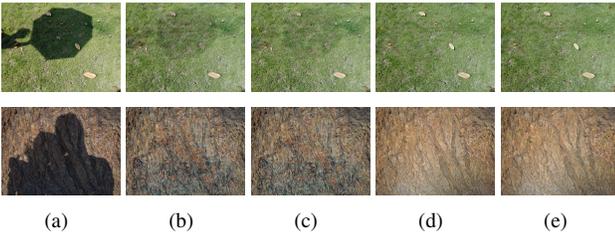


Figure 11. The shadow removal results of recurrent N in a different value. (a) is the input image. (b), (c), (d) and (e) are the results of N taking 1, 2, 3, 4 respectively.



Figure 12. The visualization results of ablation analysis. (a) is the input images. (b) is shadow removal results without using LSTM. (c) is shadow removal results without using unsupervised data. (d) is our final shadow removal results.

can observe that LSTM layers really matter for shadow-free image recovery, and with semi-supervised strategy our ARGAN+SS is powerful to handle shadow images with complex scenes.

5. Conclusions

In this paper, we have proposed a robust attentive recurrent generative adversarial network for shadow detection and removal. The generator both to generate shadow attention maps and to recover the shadow-removal images

	SRD			ISTD		
	S	N	A	S	N	A
Guo	29.89	6.47	12.60	18.95	7.46	9.3
Zhang	9.56	6.9	7.24	9.77	7.12	8.16
DeshadowNet	17.96	6.53	8.47	12.76	7.19	7.83
ST-CGAN	18.64	6.37	8.23	10.31	6.92	7.46
DSC	11.31	6.72	7.83	9.22	6.50	7.10
ARGAN	14.68	5.94	7.65	11.79	6.33	7.57
ARGAN	7.24	4.71	5.74	7.21	5.83	6.68
ARGAN+SS	6.35	4.46	5.31	6.65	5.41	5.89

Table 3. Quantitative comparison results on shadow removal with RMSE metric. The best and second best results are marked in red and blue colors, respectively. In the table, S represents shadow regions, N represents non-shadow regions, and A represents the whole image.

involves multiple progressive steps in a coarse-to-fine fashion. Our model is able to handle shadows with complex scenes and is very flexible to incorporate sufficient unsupervised shadow images to train a powerful model. In the future, we plan to extend our method to solve the challenging video shadow detection and removal problem [62].

Acknowledgement

This work was partly supported by The National Key Research and Development Program of China (2017YFB1002600), the NSFC (No. 61672390), and Key Technological Innovation Projects of Hubei Province (2018AAA062). The corresponding author is Chunxia Xiao.

References

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, volume 3, page 6, 2018.
- [2] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [3] Jan K Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio. Attention-based models for speech recognition. In *NeurIPS*, 2015.
- [4] Yung-Yu Chuang, Dan B Goldman, Brian Curless, David H Salesin, and Richard Szeliski. Shadow matting and compositing. In *ACM Transactions on Graphics (TOG)*, volume 22, pages 494–500, 2003.
- [5] Rita Cucchiara, Costantino Grana, Massimo Piccardi, Andrea Prati, and Stefano Sirotti. Improving shadow suppression in moving object detection with hsv color information. In *Intelligent Transportation Systems*, 2002.
- [6] Arbel Eli and Hel Or Hagit. Shadow removal using intensity surfaces and texture anchor points. *IEEE Transactions on Pattern Analysis Machine Intelligence (T-PAMI)*, 33(6):1202–1216, 2011.
- [7] Liu Feng and Michael Gleicher. Texture-consistent shadow removal. In *ECCV*, 2008.
- [8] Xueyang Fu, Jiabin Huang, Delu Zeng, Yue Huang, Xinghao Ding, and John Paisley. Removing rain from single images via a deep detail network. In *CVPR*, 2017.
- [9] Lianli Gao, Xiangpeng Li, Jingkuan Song, and Heng Tao Shen. Hierarchical lstms with adaptive attention for visual captioning. *IEEE Transactions on Pattern Analysis Machine Intelligence (T-PAMI)*, 2019.
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014.
- [11] Maciej Gryka, Michael Terry, and Gabriel J. Brostow. Learning to remove soft shadows. *ACM Transactions on Graphics (TOG)*, 34(5):1–15, 2015.
- [12] Ruiqi Guo, Qieyun Dai, and Derek Hoiem. Single-image shadow detection and removal using paired regions. In *CVPR*, 2011.
- [13] Kai Ming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [14] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [15] Xiaowei Hu, Chi-Wing Fu, Lei Zhu, Jing Qin, and Pheng-Ann Heng. Direction-aware spatial context features for shadow detection and removal. *arXiv preprint arXiv*, 2018.
- [16] Xiaowei Hu, Lei Zhu, Chi-Wing Fu, Jing Qin, and Pheng-Ann Heng. Direction-aware spatial context features for shadow detection. 2018.
- [17] Gang Hua, Chengjiang Long, Ming Yang, and Yan Gao. Collaborative active learning of a kernel machine ensemble for recognition. In *ICCV*, 2013.
- [18] Gang Hua, Chengjiang Long, Ming Yang, and Yan Gao. Collaborative active visual recognition from crowds: A distributed ensemble approach. *IEEE transactions on pattern analysis and machine intelligence*, 40(3):582–594, 2018.
- [19] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint*, 2017.
- [20] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016.
- [21] Imran N Junejo and Hassan Foroosh. Estimating geotemporal location of stationary cameras using shadow trajectories. In *Proceedings of the ECCV*, 2008.
- [22] Kevin Karsch, Kalyan Sunkavalli, Sunil Hadap, Nathan Carr, Hailin Jin, Rafael Fonte, Michael Sittig, and David Forsyth. Automatic scene inference for 3d object compositing. *ACM Transactions on Graphics (TOG)*, 33(3):1–15, 2014.
- [23] Salman H Khan, Mohammed Bennamoun, Ferdous Sohel, and Roberto Togneri. Automatic shadow detection and removal from a single image. *IEEE Transactions on Pattern Analysis Machine Intelligence (T-PAMI)*, (3):431–446, 2016.
- [24] Douwe Kiela, Changhan Wang, and Kyunghyun Cho. Dynamic meta-embeddings for improved sentence representations. In *Conference on Empirical Methods in Natural Language Processing*, 2018.
- [25] Jean François Lalonde, Alexei A. Efros, and Srinivasa G. Narasimhan. Detecting ground shadows in outdoor consumer photographs. In *ECCV*, 2010.
- [26] Hieu Le, Yago Vicente, F Tomas, Vu Nguyen, Minh Hoai, and Dimitris Samaras. A+d net: Training a shadow detector with adversarial shadow attenuation. In *ECCV*, 2018.
- [27] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint*, 2017.
- [28] Chuan Li and Michael Wand. Precomputed real-time texture synthesis with markovian generative adversarial networks. In *ECCV*, 2016.
- [29] Zhengqi Li and Noah Snavely. Learning intrinsic image decomposition from watching the world. In *CVPR*, 2018.
- [30] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NeurIPS*, 2017.
- [31] Chengjiang Long and Gang Hua. Multi-class multi-annotator active learning with robust gaussian process for visual recognition. In *ICCV*, 2015.
- [32] Chengjiang Long and Gang Hua. Correlational gaussian processes for cross-domain visual recognition. In *CVPR*, 2017.
- [33] Chengjiang Long, Xiaoyu Wang, Gang Hua, Ming Yang, and Yuanqing Lin. Accurate object detection with location relaxation and regionlets re-localization. In *ACCV*, 2014.
- [34] Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *CVPR*, volume 6, page 2, 2017.
- [35] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering. In *NeurIPS*, 2016.

- [36] Wenhan Luo, Peng Sun, Fangwei Zhong, Wei Liu, Tong Zhang, and Yizhou Wang. End-to-end active object tracking and its real-world deployment via reinforcement learning. *IEEE Transactions on Pattern Analysis Machine Intelligence (T-PAMI)*, 2019.
- [37] Ivana Miki, Pamela C. Cosman, Greg T. Kogut, and Mohan M. Trivedi. Moving shadow and object detection in traffic scenes. In *International Conference on Pattern Recognition (ICPR)*, 2000.
- [38] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv*, 2018.
- [39] A Mohan, J Tumblin, and P Choudhury. Editing soft shadows in a digital photograph. *IEEE Comput Graph Appl*, 27(2):23–31, 2007.
- [40] Vu Nguyen, Tomás F. Yago Vicente, Maozheng Zhao, Minh Hoai, and Dimitris Samaras. Shadow detection with conditional generative adversarial networks. In *ICCV*, 2017.
- [41] Takahiro Okabe, Imari Sato, and Yoichi Sato. Attached shadow coding: Estimating surface normals from shadows under unknown reflectance and lighting conditions. In *ICCV*, 2009.
- [42] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016.
- [43] Rui Qian, Robby T Tan, Wenhan Yang, Jiajun Su, and Jiaying Liu. Attentive generative adversarial network for rain-drop removal from a single image. In *CVPR*, 2018.
- [44] Liangqiong Qu, Jiandong Tian, Shengfeng He, Yandong Tang, and Rynson WH Lau. Dshadownet: A multi-context embedding deep network for shadow removal. In *CVPR*, 2017.
- [45] Erik Reinhard, Michael Adhikhmin, Bruce Gooch, and Peter Shirley. Color transfer between images. *IEEE Computer graphics and applications*, 21(5):34–41, 2001.
- [46] Yael Shor and Dani Lischinski. The shadow meets the mask: Pyramid-based shadow removal. In *Computer Graphics Forum*, volume 27, pages 577–586, 2008.
- [47] Edgar Simo-Serra, Satoshi Iizuka, and Hiroshi Ishikawa. Mastering sketching: adversarial augmentation for structured prediction. *ACM Transactions on Graphics (TOG)*, 37(1):11, 2018.
- [48] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [49] Hanqing Tao, Shiwei Tong, Hongke Zhao, Tong Xu, Binbin Jin, and Qi Liu. A radical-aware attention-based model for chinese text classification. In *AAAI*, 2019.
- [50] Linh Tran, Jean Kossaiji, Yannis Panagakis, and Maja Pantic. Disentangling geometry and appearance with regularised geometry-aware generative adversarial networks. *International Journal of Computer Vision*, pages 1–21, 2019.
- [51] Tomás F Yago Vicente, Minh Hoai, and Dimitris Samaras. Leave-one-out kernel optimization for shadow detection and removal. *IEEE Transactions on Pattern Analysis Machine Intelligence (T-PAMI)*, 40(3):682–695, 2018.
- [52] Tomás F Yago Vicente, Le Hou, Chen-Ping Yu, Minh Hoai, and Dimitris Samaras. Large-scale training of shadow detectors with noisily-annotated shadow examples. In *ECCV*, 2016.
- [53] Jifeng Wang, Xiang Li, and Jian Yang. Stacked conditional generative adversarial networks for jointly learning shadow detection and shadow removal. In *CVPR*, 2018.
- [54] Tai-Pang Wu, Chi-Keung Tang, Michael S Brown, and Heung-Yeung Shum. Natural shadow matting. *ACM Transactions on Graphics (TOG)*, 26(2):8, 2007.
- [55] 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, 2013.
- [56] Chunxia Xiao, Donglin Xiao, Ling Zhang, and Lin Chen. Efficient shadow removal using subregion matching illumination transfer. In *Computer Graphics Forum*, volume 32, 2013.
- [57] Yao Xiao, Efstratios Tsougenis, and Chi-Keung Tang. Shadow removal from single rgb-d images. In *CVPR*, 2014.
- [58] SHI Xingjian, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *NeurIPS*, 2015.
- [59] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015.
- [60] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, XiaoLei Huang, and Dimitris N Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. pages 5907–5915, 2017.
- [61] Ling Zhang, Qing Zhang, and Chunxia Xiao. Shadow remover: Image shadow removal based on illumination recovering optimization. *IEEE Transactions on Image Processing (TIP)*, 24(11):4623–4636, 2015.
- [62] Ling Zhang, Yao Zhu, Bin Liao, and Chunxia Xiao. Video shadow removal using spatio-temporal illumination transfer. In *Computer Graphics Forum*, volume 36, 2017.
- [63] Xiaoning Zhang, Tiantian Wang, Jinqing Qi, Huchuan Lu, and Gang Wang. Progressive attention guided recurrent network for salient object detection. In *CVPR*, 2018.
- [64] Jiejie Zhu, Kegan GG Samuel, Syed Z Masood, and Marshall F Tappen. Learning to recognize shadows in monochromatic natural images. In *CVPR*, 2010.
- [65] Lei Zhu, Zijun Deng, Xiaowei Hu, Chi-Wing Fu, Xuemiao Xu, Jing Qin, and Pheng-Ann Heng. Bidirectional feature pyramid network with recurrent attention residual modules for shadow detection. In *ECCV*, 2018.