

Fine-Context Shadow Detection using Shadow Removal

Jeya Maria Jose Valanarasu, and Vishal M. Patel Johns Hopkins University

{jvalana1, vpatel36}@jhu.edu

Abstract

Current shadow detection methods perform poorly when detecting shadow regions that are small, unclear or have blurry edges. In this work, we attempt to address this problem on two fronts. First, we propose a Fine Contextaware Shadow Detection Network (FCSD- Net), where we constraint the receptive field size and focus on low-level features to learn fine context features better. Second, we propose a new learning strategy, called Restore to Detect (R2D), where we show that when a deep neural network is trained for restoration (shadow removal), it learns meaningful features to delineate the shadow masks as well. To make use of this complementary nature of shadow detection and removal tasks, we train an auxiliary network for shadow removal and propose a complementary feature learning block (CFL) to learn and fuse meaningful features from shadow removal network to the shadow detection network. We train the proposed network, FCSD-Net, using the R2D learning strategy across multiple datasets. Experimental results on three public shadow detection datasets (ISTD, SBU and UCF) show that our method improves the shadow detection performance while being able to detect fine context better compared to the other recent methods. Our proposed learning strategy can also be adopted easily as a useful pipeline in future advances in shadow detection and removal.

1. Introduction

There is a high probability of the presence of shadows in any image captured in natural conditions. Shadows are caused by objects occluding the light from the illumination source thus causing color, intensity and texture changes in the surface where the light is obstructed. Presence of a single object with an illumination source is enough to cast a shadow (excluding perpendicular conditions), be it outdoors or indoors. Detecting shadows is of high importance in computer vision because of two main reasons. First, shadows hamper the image scene causing a performance drop for other major vision tasks like semantic segmen-



Figure 1. Our proposed learning strategy R2D. The shadow feature maps learned by the restoration network R() are forwarded to the detection network D() via a complementary feature learning block (CFL) to enhance the detection performance.

tation [11], object detection [29], video surveillance [28] and visual tracking [7]. In these tasks, careful delineation or removal of shadows aids in a performance boost. Second, detecting shadows help in other tasks like inferring the scene geometry [19, 18, 31], camera parameters [42] and light source localization [32, 22].

Traditional methods proposed for shadow detection in images develop physical models [10, 9] or use machine learning models based on hand-crafted features [23, 16, 46]. Following the popularity of deep learning methods in solving computer vision tasks, a lot of methods based on convolutional neural networks (ConvNets) were proposed for the task of shadow detection [44, 45, 47, 40, 39, 41, 6, 4, 38, 37, 30]. ConvNet-based methods are found to be successful for shadow detection as they are able to learn global contexts of the shadow region better than any other previous methods. This key property is specifically demonstrated in recent works [47, 30, 15]. However, these methods perform poorly when shadow regions or small or have blurry boundaries. In this work, we focus on solving the fine-context problem in shadow detection.

First, we propose a new network architecture for the de-



Figure 2. Predicting the shadow segmentation map by finding the residual of restored image from the input image. It can be seen that the residual is noisy while our proposed learning strategy R2D predicts a shadow segmentation as good as the ground truth.

tection to efficiently segment shadow regions of fine context. Although ConvNets work better at capturing global context as the deeper layers focus on bigger objects due to a larger receptive field, the number of filters focusing on extracting information about fine context is still very less as only the shallow layers are responsible for extracting local information. Even though extracting local features were explored by using input at different resolutions [47], it still does not work well in segmenting fine context shadow regions which are small or with unclear boundaries. To this end, we propose a Fine Context aware Shadow Detection Network (FCSD-Net) where we introduce a fine context detector block that constraints the receptive field size from enlarging in the deeper layers thus helping in feature extraction of shadow regions which are small, unclear or with confounding boundaries.

Shadow removal is also a growing area of research in computer vision where the task is to restore the original image by removing the shadow. Traditional methods [27, 43, 3, 2] as well as deep learning-based methods [24, 26, 36, 1, 21, 5] have been proposed to solve this task. Shadow detection and removal can be thought of as complementary tasks as one involves detecting the shadow region and the other involves removing the shadow region and replacing it with the background scenery. A naive solution for shadow detection using shadow removal would be to subtract the restored image from the input image. However the residual is very noisy as the restoration is not perfect even if it looks visually good. An example for this can be found in Fig 2 where we restore the image using the current stateof-the-art shadow removal method-dual hierarchically aggregation network (DHAN) [8]. It can be noted that even though the PSNR and SSIM are reasonable for restoration, the residual is still not the desirable shadow segmentation. Thus in this work, we propose a new way to efficiently use the removal task to enhance the detection task.

It is evident that deep networks trained for shadow removal also learn features of that about the shadow region. This concept in explained in detailed in section 3.3 and also illustrated in Fig 6. We try to make use of this phenomenon and avail these feature maps to enhance shadow detection. A recently published dataset - ISTD [39] consists of triplets of images- image with shadow, shadow segmentation and the clean image making this idea feasible. It can be noted that no previous method has attempted to make use of the restoration task to improve the detection task. The closest work is [39] where a stacked conditional generative adversarial network was designed to solve shadow detection and removal in sequence. It adds the shadow detection mask as an input the removal network however its detection pipeline remains generic. We propose a new learning strategy for shadow detection, called Restore to Detect (R2D), where we have a restoration network R() that learns shadow removal as an auxiliary task and feeds the feature maps learned to the detection network D(). We also propose a complementary feature leaning block (CFL) that learns to feed only the shadow features from R() to D(). The proposed learning framework R2D is illustrated in Fig 1. To the best of our knowledge, this is the first work that attempts to use the shadow removal task to enhance shadow detection. We perform extensive experiments on multiple datasets to show that our proposed network FCSD-Net with the proposed learning strategy R2D performs better then the existing methods both qualitatively and quantitatively.

The main contributions of this paper are as follows:

- We propose FCSD-Net that can effectively detect shadow regions which are unclear, blurry or small. It also reduces the false-detection of background regions with similar color and intensity as shadows.
- We propose a new learning strategy called Restore to Detect (R2D) that leverages the feature maps learned during shadow removal for enhancing the performance of shadow detection. This learning strategy can be easily adopted by any current or future shadow detection methods with ease.
- We achieve improvements in performance across three public shadow detection datasets - SBU [38], UCF [46] and ISTD [39] both quantitatively and qualitatively, especially in confounding cases.

2. Related Work

Traditional methods. Early methods for shadow detection focused on hand crafting meaningful features to discriminate the shadow from the background in the image. For example, [23, 16, 46] design hand-crafted features based on edge and pixel information of shadow regions for shadow detection in consumer photographs. Later, various classifiers were explored using the hand-crafted features. In particular, Huang *et al.* [16] used edge features of shadow to train an SVM [13] for shadow detection. Guo *et al.* [12] used a graph-based classifier to segment shadows using illumination features. Vicente *et al.* [37] used MRF to boost the shadow detection performance by using pairwise region context information. Few early works [9, 10] also explored building illumination and color based models to segment

shadows. Most of these methods perform poorly in complex scenarios as hand-crafted features do not discriminate shadow region from the image background well.

Deep learning-based methods. Following the success of deep learning in various computer vision tasks such as semantic segmentation, object detection, image restoration, and image-to-image translation, various ConvNetbased methods for shadow detection have been proposed in the literature [20, 38, 30, 25, 15]. Khan et al. [20] first used a simple 7 layer ConvNet to automatically extract the feature descriptors for shadow detection. The feature learning part was done at the super-pixel level and along the shadow boundaries. Vicente et al. [38] proposed a semantic patch level ConvNet to train efficiently on patches while also using the image level semantic information. Nguyen et al. [30] introduced scGAN where cGAN was tailor-made for the task of shadow detection. scGAN used an adversarial approach to model high-level context and global features better. He et al. [25] proposed another adversarial strategy where two separate networks - attenuation network and shadow detection network were trained together in an adversarial way where the attenuation network tries to fool the detection network's shadow prediction by modifying the input shadow image. This effectively acts as data augmentation to training data as the shadow detection network is now trained on various new instances of the input data.

Hu et al. [15] explored a direction-aware manner to analyze image context for shadow detection. A new module, direction-aware spatial context (DSC), was proposed which uses a spatial RNN to learn spatial context of shadows. Zhu et al. [47] proposed a bidirectional feature pyramid network (BFPN) which combines contexts of two different ConvNets by using a recurrent attention residual (RAR) to refine the context features from deep to shallow layers. Hosseinzadeg et al. [14] proposed obtaining a shadow prior first using multi-class SVM using statistical features and then use it along with the original image to train a patch-level ConvNet. Zheng et al. [45] proposed a Distraction aware Shadow Detection (DSDNet) where a standalone Distraction-aware Shadow (DS) module was introduced to learn discriminative features for robust shadow detection. DS module learns to predict the false positive and false negative maps explicitly which helps in learning and integrating the visual distraction regions to get an efficient shadow detection. Wang et al. [41] proposed a densely cascaded learning method to fuse both global and local details efficiently. Recently, Chen et al. [6] proposed a semi-supervised framework using a teacher-student network where unlabeled shadow data were used to further improve the performance of the network.

3. Proposed Method

In this section, we explain our proposed architecture FCSD-Net and learning strategy R2D in detail.

3.1. Fine Context Shadow Detection

In most of the previous deep learning-based solutions for shadow detection, the ConvNets used are of a generic encoder-decoder network architecture where the input image is taken to a lower dimension. The conv layers are designed such that the receptive field size of the successive layers increases through every conv layer. This forces the network to focus on high level information in deeper layers. Although this actually helps the networks learn shadow structures of people and large objects well, it reduces the focus the network gives to smaller shadow regions and shadows with unclear boundaries. This happens because smaller shadow regions and sharper edges need conv filters with smaller receptive field to learn features that extract them. So, in our proposed solution to this problem we make use of an alternate design of ConvNets to focus more on fine details.

Consider a configuration of two conv layers in succession where *I* be the input image, F_1 and F_2 be the feature maps extracted from the conv layers 1 and 2, respectively. As in a generic ConvNet, let there be max-pooling layer between these conv layers. In this case, the receptive field of conv layer 2 (to which F_1 is forwarded) on the input image would be $2 \times k \times 2 \times k$. However if we have an upsampling layer instead of the max-pooling layer, the receptive field would become $\frac{1}{2} \times k \times \frac{1}{2} \times k$ [34, 35]. This helps in the alternative ConvNet architecture to learn more low-level information like edges and other finer details better.

3.2. FCSD-Net

Using the above concept, we propose FCSD-Net which is an end-to-end trainable network where the input is the image with shadow and the output is shadow segmentation. Similar to [45, 47], we use a ResNeXt backbone network for feature extraction. We select the output of the last conv layer of each conv block in the ResNeXT architecture. Specifically, we choose the features from the layers - conv1, res2c, res3b3, res4b22 and res5c and term them L1, L2, L3, L4 and L5 for simplicity. Note that, out of these feature maps L1 and L2 have low-level feature information while all the others have high-level information. From the backbone network, we pass the features extracted to two blocks - Fine context detector (FCD) and coarse context detector (CCD). We pass all the feature maps (L1, L2, L3, L4 and L5) to CCD while we pass only L1 to FCD as the focus of FCD is to detect low-level shadow information.

In CCD, all these backbone feature maps are first forwarded to a conv block to learn the image features corre-



Figure 3. The proposed architecture for FCSD-Net. The input image is first fed to the ResNeXt backbone network. Backbone features L1 - L5 are forwarded to the CCD and FCD blocks. The feature maps F1 - F6 and DS1 - DS6 extracted from these blocks are forwarded to the fusion block which outputs the shadow segmentation map.



Figure 4. Feature maps collected from the trained FCSD-Net. Top row corresponds to the feature maps collected from CCD while the bottom row corresponds to the feature maps collected from FCD. It can be seen that feature maps in CCD focus on high-level context while the feature maps in FCD focus on fine context.

sponding to shadow detection. We term these image features as F1, F2, F3, F4 and F5, respectively. Then at each scale, we pass these image features to the Distraction Aware Shadow (DS) Block [45]. We note from [45] that capturing the distraction features using the DS module helps in improving the shadow detection performance. It is designed such that the false positive distraction and the false negative distraction of the shadow are learned separately and then efficiently integrated. More information about the DS module can be found in [45]. We term these distraction aware features DS1, DS2, DS3, DS4 and DS5.

In FCD, input L1 is forwarded to an alternative ConvNet architecture. Here, we have a set of 4 conv layers with upsampling layers in between them. Note that upsampling can be done either by interpolation techniques or transpose convolution. After some early experiments, we found that the performance of both methods were similar. So we use a simple bilinear interpolation in this work to reduce the computation. The interpolation factor chosen at each layer is such that feature maps grow 50 more pixels in terms of height and width when compared to the feature map dimensions before the upsampling layer. The receptive field gets constrained after each upsampling layer and each successive conv layer. We term the fine context features extracted from the last conv layer as F6. Similar to CCD, we use the DS module here as well as it helps in capturing the shadow distraction features. We pass F6 to the DS module to get the distraction features DS6.

The feature maps F1 to F6, DS1 to DS6 taken from FCD and CCD are forwarded to the fusion block. In the fusion block, we interpolate all the feature maps to the same size as of the input image. After this we concatenate all the feature maps and pass them through a 1×1 conv layer. This output is forwarded to a sigmoid activation to get the binary shadow map as output. The architecture of FCSD-Net is illustrated in Fig 3. The fusion block figure can be found in the supplementary material. Sample features extracted from FCD and CCD blocks are illustrated in Fig 4.

3.3. Complementary Nature of Shadow Removal and Detection

Shadow detection is the task of predicting the shadow binary mask from the input image with shadow. Shadow removal is the task of removing the shadow from the image and produce a restored image where the shadow part is replaced with the background. Both of these tasks need information about the shadow region as they involve careful delineation of the shadow. To show the complementary nature of shadow removal and detection tasks, let us assume that the input image with shadow is denoted by I, shadow segmentation mask that has the detected shadow region is denoted d and the restored image is denoted by r. Let the



Figure 5. A detailed illustration of the proposed method- R2D. The input image is passed to both restoration network (R()) - U-Net and detection network (D()) - FCSD-Net. The features R1 and R2 are forwarded to CFL. The output of CFL - C1 and C2 are forwarded to FCSD-Net where they are added to L2 and L5, respectively. The output is taken from FCSD-Net.

task of shadow detection be learned by a ConvNet D() and the task of shadow removal be learned by a ConvNet R(). Then, the shadow detection and removal tasks can be formulated as follows:

$$d = D(I), \quad r = R(I).$$
 (1)
The shadow removal task can be further formulated as

$$r = I - d,$$
 (2)

where d represents the shadow pixels that need to be removed. One can clearly see that learning R() involves learning d. Hence, we try to use this complementary information when we learn to restore and use it to enhance the performance of detection. This phenomenon can be empirically observed in Fig 6.



Figure 6. Feature maps collected while training U-Net for shadow removal. Even though the task here is shadow removal, the ConvNet still learns a lot of features to detect the shadow region.

Although the standard datasets used for shadow detection consists of only pairs of shadow images with their corresponding segmentation masks, the recently released ISTD dataset consists of triplets of images, image with shadow, shadow segmentation and the clean restored image. This opens the door for us to explore the idea of availing the information learned while learning to remove a shadow and use it for detection.

3.4. R2D: Restore to Detect

In the proposed Restore to Detect (R2D) method, we have two networks D() and R() learning to do shadow detection and shadow removal, respectively. For our network R(), we use a standard U-Net [33] architecture as it is lightweight and effective compared to the other standard networks. U-Net consists of an encoder-decoder architecture where both the encoder and decoder have 5 conv blocks each. There are skip connections between the conv blocks in the encoder to the conv blocks of the decoder. We select the output feature maps from the 2nd layer and the 5th layer of the encoder to pass them to network D(). Note that though all the feature maps at each level can be forwarded to the network D(), we choose to forward the feature maps at only 2 levels as from our experiments we found that it was enough in terms of performance and also helped to reduce the complexity of our method. We choose the 2nd layer and the 5th layer specifically as we could leverage both local information and global information learned by the network R(), respectively. We term these feature maps R1 and R2. More explanation on why we exactly choose the 2nd and the 5th layer can be found in the supplementary material. Note that we did not use any state-of-the-art shadow removal network in R() as it increases the complexity by a lot and hinders while optimizing D() in the R2D framework.

We propose a Complementary Feature Learning (CFL) block to efficiently learn and forward only the features that correspond to shadow regions. In CFL, we have conv layers acting at each feature map level, R1 and R2. We then interpolate them to match the size of the feature maps where it is going to be added in the D() network. We term these features C1 and C2. We fuse C1 and C2 with L2 and L5, respectively in FCSD-Net. The R2D learning strategy is illustrated in Fig 5.

3.5. Training Strategy

For training R2D, we jointly optimize the shadow prediction, false positive (FP) and false negative (FN) maps at all scales. Note that all the datasets provide ground truth shadow segmentation masks. For FP and FN ground truth maps, we use the data provided by [45] which was created based on the differences between the existing shadow detection and their ground truths. We first find a weighted binary cross entropy loss \mathcal{L}_W as follows:

$$\mathcal{L}_W = -\left(\sum_i (a_n y_i \log(x_i) + b_n (1 - y_i) \log(1 - x_i))\right),$$

where x_i corresponds to the prediction, y_i corresponds to the ground-truth, $a_n = \frac{N_n}{N_n + N_p}$, $b_n = \frac{N_p}{N_n + N_p}$. N_n and N_p are the number of positive pixels and negative pixels, respectively. Note that this weighted loss is applied pixel wise and summed up over all the pixels. We calculate \mathcal{L}_W for FP and FN maps and term them \mathcal{L}_{FP} and \mathcal{L}_{FN} , respec-

| Method | | UCF [46] SBU [38 | | 8] | ISTD [39] | | | Mean | | |
|----------------------------|-------|------------------|-----------|------|-----------|-----------|-------|--------|-----------|--------|
| | BER | Shadow | Non shad. | BER | Shadow | Non shad. | BER | Shadow | Non shad. | BER |
| MTMT-Net (CVPR 2020) [6] | 7.47 | 10.31 | 4.63 | 3.15 | 3.73 | 2.57 | 1.72 | 1.36 | 2.08 | 4.11 |
| stacked-CNN (ECCV 16) [38] | 13.00 | 9.00 | 17.1 | 10.8 | 8.84 | 12.76 | 8.6 | 7.69 | 9.23 | 10.8 |
| SRM (ICCV 17) [40] | 12.51 | 21.41 | 3.6 | 6.51 | 10.52 | 2.50 | 7.92 | 13.97 | 1.86 | 8.98 |
| scGAN (ICCV 17) [30] | 11.52 | 7.74 | 15.3 | 9.04 | 8.39 | 9.69 | 4.7 | 3.22 | 6.18 | 8.42 |
| ST-CGAN (CVPR 18) [39] | 11.23 | 4.94 | 17.52 | 8.14 | 3.75 | 12.53 | 3.85 | 2.14 | 5.55 | 7.74 |
| DSC (CVPR 18) [15] | 10.54 | 18.08 | 3.00 | 5.59 | 9.76 | 1.42 | 3.42 | 3.85 | 3.00 | 6.51 |
| ADNet (ECCV 18) [25] | 9.25 | 8.37 | 10.14 | 5.37 | 4.45 | 6.3 | - | - | - | 7.31 † |
| RAS (ECCV 18) [4] | 13.62 | 23.06 | 4.18 | 7.31 | 12.13 | 2.48 | 11.14 | 19.88 | 2.41 | 10.69 |
| BDRAR (ECCV 18) [47] | 7.81 | 9.69 | 5.94 | 3.64 | 3.40 | 3.89 | 2.69 | 0.50 | 4.87 | 4.71 |
| DSDNet (CVPR 19) [45] | 7.59 | 9.74 | 5.44 | 3.45 | 3.33 | 3.58 | 2.17 | 1.36 | 2.98 | 4.40 |
| DC-DSPF (IJCAI 19) [41] | 7.90 | 6.50 | 9.30 | 4.90 | 4.70 | 5.10 | - | - | - | 6.40 † |
| EGNet (ICCV 19) [44] | 9.20 | 11.28 | 7.12 | 4.49 | 5.23 | 3.75 | 1.85 | 1.75 | 1.95 | 5.11 |
| R2D (Ours) | 6.96 | 8.32 | 5.60 | 3.15 | 2.74 | 3.56 | 1.69 | 0.59 | 2.79 | 3.93 |

Table 1. Comparison of quantitative results in terms of BER of shadow region, non-shadow region and average BER for the UCF, SBU and ISTD datasets with the state-of-the-art shadow detectors. Note that all the numbers are in terms of error, so the lesser the better. †Mean found only using 2 datasets. Our results are in **Bold** and we acheive state-of-the art across all datasets with least BER.

tively. We also adopt the distraction-aware cross entropy loss \mathcal{L}_{DS} from [45] which forces the predictions to be less prone to false detections. \mathcal{L}_{DS} is formulated as follows:

$$\mathcal{L}_{DS} = -\left(\sum_{i} (a_n y_i^{fnd} y_i \log(x_i) + b_n y_i^{fpd} (1 - y_i) \log(1 - x_i))\right),$$
(3)

 $+ b_n g_i$ $(1 - g_i) \log(1 - x_i))),$ where y_i^{fpd} is the ground truth of FP pixel and y_i^{fnd} is the ground truth of FN pixel. We define the shadow loss \mathcal{L}_{shadow} calculated on the final shadow prediction as:

$$\mathcal{L}_{shadow} = \mathcal{L}_W + \mathcal{L}_W \mathcal{L}_{DS},\tag{4}$$

which is calculated between the shadow predictions and the ground truth segmentation map. The total loss used to train the detection network is defined as follows:

$$\mathcal{L}_{det} = \sum_{k} \alpha \mathcal{L}_{shadow}^{k} + \beta \mathcal{L}_{FP}^{k} + \gamma \mathcal{L}_{FN}^{k}, \qquad (5)$$

where k represents the shadow predictions found across different scales in the fusion block. The values of α , β and γ are set equal to 1, 2 and 2, respectively in our experiments. The L_2 loss is used to pretrain the restoration network (U-Net). It is defined as: $\mathcal{L}_{res} = \sum_i (y_i - p_i)^2$, where y_i corresponds to the target clean image pixel and p_i corresponds to the prediction at the i^{th} pixel location in the restored image. During the fine-tuning stage, we train the entire framework using the following total loss:

$$\mathcal{L}_{R2D} = \mathcal{L}_{det} + \mathcal{L}_{res}.$$
 (6)

While using the R2D strategy, we first pretrain the restoration network R() on the ISTD dataset using \mathcal{L}_{res} for 500 epochs. R() performs restoration reasonably well as it achieves an RMSE of 11.21 while evaluated on ISTD test dataset. For training the images we first rescale them to 320×320 resolution. We use a batch size of 16 and a learning rate of 0.001. We use SGD as the optimizer with momentum set at 0.9 and weight decay at 0.0001. We use the PyTorch framework for training using a Nvidia Quadro RTX-8000 GPU. Note that the ResNext-101 was pre-trained on ImageNet and all the other parameters are randomly initialized. We also use data augmentation by randomly flipping both the input and the ground truth.

Note that only the ISTD dataset consists of triplets of images while the SBU and UCF datasets contain only image pairs. The pretraining stage is common for experiments on all the datasets where the restoration network is trained on the ISTD dataset. During the finetuning phase while training ISTD dataset, we still have access to the clean images so we train the entire network with \mathcal{L}_{R2D} . For the other datasets, during the fine tuning stage we only use \mathcal{L}_{det} . Note that even then, we have feature maps flowing from the restoration network to the detection network which enhance the quality of shadow predictions. The weights are also back-propagated through the CFL to the encoder of the restoration network which makes our training strategy not dependent on the availability of clean images after pretraining. The fine-tuning is performed for 6000 iterations.

3.6. Inference Strategy

During inference , we feed forward the image to both FCSD-Net and U-Net (D() and R() respectively). The output shadow segmentation map is taken from D(). We do stochastic weight averaging [17] on the models saved at 4000, 5000, and 6000 iterations and use it to perform the inference.

4. Experimental and Results

In this section, we give details about the experiments we conduct to compare our method against the recent deep learning-based methods: stacked-CNN [37], SRM [40], sc-GAN [30], ST-GAN [39], BDRAR [47], DSC [15], ADNet [25], DC-DSPF [41], RAS [4], MTMT-Net [6] and DSDNet [45]. We illustrate the qualitative results as well as compute the performance metrics that are widely used in the shadow detection literature to quantitatively compare our proposed methods with the recent methods.

4.1. Datasets and Evaluation Metrics

The following three datasets are used to conduct experiments - UCF [46], SBU [38] and ISTD [39] datasets. In



Figure 7. Comparison of predictions of our proposed methods with leading shadow detection methods. The first and last columns correspond to the input and ground-truth, respectively. Other columns correspond to the predictions obtained using different methods.

both the SBU and UCF datasets, there are images with shadows and corresponding shadow segmentation masks. The SBU dataset consists of 4089 training images and 638 testing images. In the UCF dataset, we only use the testing set which contains 110 images similar to the previous works. We train the network on the SBU training set and test on both SBU testing set and UCF testing set. Unlike the SBU and UCF datasets, ISTD dataset contains triplets of imagesimage with shadow, shadow segmentation mask and clean image without shadow. The number of training and testing sets of images in ISTD is 1870 and 540 respectively.

We use the balanced error rate (BER) as the performance metric for quantitative comparisons with the recent methods. BER is calculated as follows:

$$BER = 1 - \frac{1}{2} \left(\frac{TP}{TP + FN} + \frac{TN}{TN + FP} \right), \quad (7)$$

where TP, TN, FP and FN correspond to the number of pixels which are true positives, true negatives, false positives and false negatives, respectively. We calculate the BER of shadow and the non-shadow regions separately and then calculate the average BER. We report the individual shadow, non-shadow BER as well as the average BER. We also report the mean BER found across all 3 datasets.

4.2. Quantitative Results

In Table 1, we summarize the results of our experiments. It can be observed from this table that our methods are observed to perform the best across all 3 datasets. We note that overall MTMT-Net [6] and DSDNet [45] are the second best performing methods. When compared to DSDNet, our method has an 8.30 %, 8.69 %, 22.11 % improvement in terms of BER on the UCF, SBU and ISTD datasets, respectively. We also calculate the mean BER across all datasets where we achieve an improvement of 10.68 %, 4.3% over DSDNet and MTMT-Net, respectively. In addition, we also conducted an experiment using the residuals as the shadow

segmentation masks as explained in Section 1. The performance was not even comparable with any of the baseline methods as the BER on the ISTD dataset was 27.32.

4.3. Qualitative Results

We visualize the predictions of our method and the best performing methods for comparison in Fig 7. It can be observed that our method's predictions are better and more closer to the ground truth when compared to other methods. In the first row, other methods mistake the lines between the wooden strips as shadows as they are of the same color of the shadows. However, our method does not falsely detect them as shadows as we focus on extracting fine context shadow regions better. Similar observations can be made in the third row. In the second row, our method does not falsely detect the smaller shadow regions found in the shore while all the other methods detect them as shadows.

5. Discussion

Ablation Study: We conduct an ablation study to show the importance of each individual component in our proposed method. We start with just using the backbone features (BB) for shadow detection. Then, we add the CCD block and the FCD block separately and conduct experiments. Then we use both blocks together and use the fusion block to fuse the features learned by CCD and FCD. This configuration corresponds to FCSD-Net. We then add the U-Net based restoration network and directly fuse the feature maps to FCSD-Net. Then, we add the CFL block to specifically learn shadow features to be forwarded from the restoration network to the detection network. This configuration corresponds to R2D. The results corresponding to these experiments are shown in Table 2 and Figure 8.

Using R2D with other networks: Also, to show the adaptability of R2D, we conduct an experiment where we use



Input Image BB BB+CCD BB+FCD BB+CCD+FCD R2D (w/o CFL) R2D (w CFL) Ground Truth Figure 8. Comparison of qualitative results for ablation study. The first and last column correspond to the input and ground truth respectively. Other columns correspond to the predictions obtained using different configurations in the ablation study.



Figure 9. The main differences between ST-cGAN and R2D.

R2D for DSDNet. The results can be seen in Table 3. It can be observed that we get an improvement of 5.90 % while using R2D learning strategy to train DSDNet when compared to training it normally.

Difference from ST-cGAN: ST-cGAN [39] uses a stacked conditional generative adversarial network to solve shadow detection and removal in sequence. It uses the shadow mask output as an additional information for the shadow removal network. Note that the detection pipeline in ST-cGAN is generic. In contrast, R2D learning strategy uses complementary feature maps information from the removal network to improve the detection performance. This difference is also visualized in Figure 9.

Limitations: We acknowledge that our method is more computationally complex $(1.3 \times \text{DSDNet})$ in terms of FLOPs) than previous methods. Our inference time is 0.26 seconds per image of resolution 320×320 which is more than MTMT-Net. In the R2D framework, we used U-Net as our restoration network. A better strategy would be to use powerful shadow removal networks like DHAN [8]. From our experiments, we were unable to train R() and D() efficiently when we chose R() as DHAN instead of U-Net due to the huge number of parameters and difficulty in optimizing both networks in parallel. With better engineering strategy, using DHAN should be able to further improve the effectivess of R2D. Using FCSD-Net's shadow detection features to boost shadow removal performance is also a possible direction left unexplored in this work due to lack

| Method | UCF [46] | SBU [38] | ISTD [39] | Mean |
|---------------|----------|----------|-----------|------|
| BB | 9.64 | 4.48 | 3.41 | 5.84 |
| BB + CCD | 7.63 | 3.48 | 2.18 | 4.43 |
| BB + FCD | 7.52 | 3.40 | 1.82 | 4.24 |
| BB +CCD + FCD | 7.08 | 3.30 | 1.71 | 4.03 |
| R2D (w/o CFL) | 7.04 | 3.26 | 1.70 | 4.00 |
| R2D (w CFL) | 6.96 | 3.15 | 1.69 | 3.93 |

Table 2. Comparison of quantitative results in terms of average BER for ablation study.

| Method | UCF [46] | SBU [38] | ISTD [39] | Mean |
|--------------|-----------|----------|--------------|------|
| DSDNet | 7.59 | 3.45 | 2.17 | 4.40 |
| DSDNet + R2D | 7.30 | 3.39 | 1.78 | 4.15 |
| 11 2 0 . | · · · · · | | . • . | C |

Table 3. Comparison of quantitative results in terms of average BER for ablation study on using R2D for DSD-Net.

of compute power. Optimizing the networks in parallel and leveraging shadow removal for shadow detection and viceversa is another possible setup not explored in this paper. Although these setups are theoretically possible, the roadblocks are compute power and unstable training.

6. Conclusion

In this work, we explored a new direction for shadow detection. We propose a new method, R2D, in which we leverage the shadow features learned during shadow removal to improve the shadow detection performance. We also propose a new network architecture, FCSD-Net, that serves as the detection network architecture in our R2D framework. It mainly focuses on fine context feature extraction for shadow detection. We do this by designing a fine context detector block where we constrain the receptive field size to focus more on the local features which improves the detection performance especially in confounding cases where the shadow region and background have similar color intensities. R2D can be easily adopted as the learning strategy to enhance any shadow detection network. We conduct extensive experiments to show the effectiveness of our methods. Using shadow detection as an auxiliary task for shadow removal is considered as a future direction of this work.

References

- M Anand, A Ashwin Natraj, V Jeya Maria Jose, K Subramanian, Priyanka Bhardwaj, R Pandeeswari, and S Deivalakshmi. Tackling multiple visual artifacts: Blind image restoration using conditional adversarial networks. In *International Conference on Computer Vision and Image Processing*, pages 331–342. Springer, 2019.
- [2] Eli Arbel and Hagit Hel-Or. Shadow removal using intensity surfaces and texture anchor points. *IEEE transactions on pattern analysis and machine intelligence*, 33(6):1202–1216, 2010.
- [3] Masashi Baba, Masayuki Mukunoki, and Naoki Asada. Shadow removal from a real image based on shadow density. In ACM SIGGRAPH 2004 Posters, page 60. 2004.
- [4] Shuhan Chen, Xiuli Tan, Ben Wang, and Xuelong Hu. Reverse attention for salient object detection. In *Proceedings* of the European Conference on Computer Vision (ECCV), pages 234–250, 2018.
- [5] Zhihao Chen, Liang Wan, Lei Zhu, Jia Shen, Huazhu Fu, Wennan Liu, and Jing Qin. Triple-cooperative video shadow detection, 2021.
- [6] Zhihao Chen, Lei Zhu, Liang Wan, Song Wang, Wei Feng, and Pheng-Ann Heng. A multi-task mean teacher for semi-supervised shadow detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5611–5620, 2020.
- [7] Rita Cucchiara, Costantino Grana, Massimo Piccardi, Andrea Prati, and Stefano Sirotti. Improving shadow suppression in moving object detection with hsv color information. In *ITSC 2001. 2001 IEEE Intelligent Transportation Systems. Proceedings (Cat. No. 01TH8585)*, pages 334–339. IEEE, 2001.
- [8] Xiaodong Cun, Chi-Man Pun, and Cheng Shi. Towards ghost-free shadow removal via dual hierarchical aggregation network and shadow matting gan. In AAAI, pages 10680– 10687, 2020.
- [9] Graham D Finlayson, Mark S Drew, and Cheng Lu. Entropy minimization for shadow removal. *International Journal of Computer Vision*, 85(1):35–57, 2009.
- [10] 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 intelli*gence, 28(1):59–68, 2005.
- [11] Ye-Peng Guan. Wavelet multi-scale transform based foreground segmentation and shadow elimination. *The Open Signal Processing Journal*, 1(1), 2008.
- [12] Ruiqi Guo, Qieyun Dai, and Derek Hoiem. Single-image shadow detection and removal using paired regions. In CVPR 2011, pages 2033–2040. IEEE, 2011.
- [13] Marti A. Hearst, Susan T Dumais, Edgar Osuna, John Platt, and Bernhard Scholkopf. Support vector machines. *IEEE Intelligent Systems and their applications*, 13(4):18–28, 1998.
- [14] Sepideh Hosseinzadeh, Moein Shakeri, and Hong Zhang. Fast shadow detection from a single image using a patched convolutional neural network. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 3124–3129. IEEE, 2018.

- [15] Xiaowei Hu, Lei Zhu, Chi-Wing Fu, Jing Qin, and Pheng-Ann Heng. Direction-aware spatial context features for shadow detection. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 7454– 7462, 2018.
- [16] Xiang Huang, Gang Hua, Jack Tumblin, and Lance Williams. What characterizes a shadow boundary under the sun and sky? In 2011 international conference on computer vision, pages 898–905. IEEE, 2011.
- [17] Pavel Izmailov, Dmitrii Podoprikhin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. arXiv preprint arXiv:1803.05407, 2018.
- [18] Imran N Junejo and Hassan Foroosh. Estimating geotemporal location of stationary cameras using shadow trajectories. In *European conference on computer vision*, pages 318–331. Springer, 2008.
- [19] Kevin Karsch, Varsha Hedau, David Forsyth, and Derek Hoiem. Rendering synthetic objects into legacy photographs. ACM Transactions on Graphics (TOG), 30(6):1–12, 2011.
- [20] Salman Hameed Khan, Mohammed Bennamoun, Ferdous Sohel, and Roberto Togneri. Automatic feature learning for robust shadow detection. In 2014 IEEE Conference on Computer Vision and Pattern Recognition, pages 1939– 1946. IEEE, 2014.
- [21] Salman H Khan, Mohammed Bennamoun, Ferdous Sohel, and Roberto Togneri. Automatic shadow detection and removal from a single image. *IEEE transactions on pattern* analysis and machine intelligence, 38(3):431–446, 2015.
- [22] Jean-François Lalonde, Alexei A Efros, and Srinivasa G Narasimhan. Estimating natural illumination from a single outdoor image. In 2009 IEEE 12th International Conference on Computer Vision, pages 183–190. IEEE, 2009.
- [23] Jean-François Lalonde, Alexei A Efros, and Srinivasa G Narasimhan. Detecting ground shadows in outdoor consumer photographs. In *European conference on computer* vision, pages 322–335. Springer, 2010.
- [24] Hieu Le and Dimitris Samaras. Shadow removal via shadow image decomposition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8578–8587, 2019.
- [25] Hieu Le, Tomas F Yago Vicente, Vu Nguyen, Minh Hoai, and Dimitris Samaras. A+ d net: Training a shadow detector with adversarial shadow attenuation. In *Proceedings* of the European Conference on Computer Vision (ECCV), pages 662–678, 2018.
- [26] Yun-Hsuan Lin, Wen-Chin Chen, and Yung-Yu Chuang. Bedsr-net: A deep shadow removal network from a single document image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12905–12914, 2020.
- [27] Feng Liu and Michael Gleicher. Texture-consistent shadow removal. In *European Conference on Computer Vision*, pages 437–450. Springer, 2008.
- [28] F Matusek, G Pujolle, and R Reda. Shadow detection for increased accuracy of privacy enhancing methods in video surveillance edge devices. In *Proceedings of the Interna-*

tional Conference on Pattern Recognition and Computer Vision. Citeseer, 2008.

- [29] Ivana Mikic, Pamela C Cosman, Greg T Kogut, and Mohan M Trivedi. Moving shadow and object detection in traffic scenes. In *Proceedings 15th International Conference on Pattern Recognition. ICPR-2000*, volume 1, pages 321–324. IEEE, 2000.
- [30] Vu Nguyen, Tomas F Yago Vicente, Maozheng Zhao, Minh Hoai, and Dimitris Samaras. Shadow detection with conditional generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4510–4518, 2017.
- [31] Takahiro Okabe, Imari Sato, and Yoichi Sato. Attached shadow coding: Estimating surface normals from shadows under unknown reflectance and lighting conditions. In 2009 IEEE 12th International Conference on Computer Vision, pages 1693–1700. IEEE, 2009.
- [32] Alexandros Panagopoulos, Dimitris Samaras, and Nikos Paragios. Robust shadow and illumination estimation using a mixture model. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 651–658. IEEE, 2009.
- [33] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. Unet: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [34] Jeya Maria Jose Valanarasu, Vishwanath A Sindagi, Ilker Hacihaliloglu, and Vishal M Patel. Kiu-net: Overcomplete convolutional architectures for biomedical image and volumetric segmentation. arXiv preprint arXiv:2010.01663, 2020.
- [35] Jeya Maria Jose Valanarasu, Vishwanath A Sindagi, Ilker Hacihaliloglu, and Vishal M Patel. Kiu-net: Towards accurate segmentation of biomedical images using over-complete representations. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 363–373. Springer, 2020.
- [36] Florin-Alexandru Vasluianu, Andres Romero, Luc Van Gool, and Radu Timofte. Self-supervised shadow removal. arXiv preprint arXiv:2010.11619, 2020.
- [37] Tomas F Yago Vicente, Minh Hoai, and Dimitris Samaras. Leave-one-out kernel optimization for shadow detection and removal. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(3):682–695, 2017.
- [38] 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 *European Conference on Computer Vision*, pages 816–832. Springer, 2016.
- [39] Jifeng Wang, Xiang Li, and Jian Yang. Stacked conditional generative adversarial networks for jointly learning shadow detection and shadow removal. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1788–1797, 2018.
- [40] Tiantian Wang, Ali Borji, Lihe Zhang, Pingping Zhang, and Huchuan Lu. A stagewise refinement model for detecting salient objects in images. In *Proceedings of the IEEE Inter-*

national Conference on Computer Vision, pages 4019–4028, 2017.

- [41] Yupei Wang, Xin Zhao, Yin Li, Xuecai Hu, Kaiqi Huang, et al. Densely cascaded shadow detection network via deeply supervised parallel fusion. 2019.
- [42] Lin Wu, Xiaochun Cao, and Hassan Foroosh. Camera calibration and geo-location estimation from two shadow trajectories. *Computer Vision and Image Understanding*, 114(8):915–927, 2010.
- [43] Mei Xiao, Chong-Zhao Han, and Lei Zhang. Moving shadow detection and removal for traffic sequences. *International Journal of Automation and Computing*, 4(1):38–46, 2007.
- [44] Jia-Xing Zhao, Jiang-Jiang Liu, Deng-Ping Fan, Yang Cao, Jufeng Yang, and Ming-Ming Cheng. Egnet: Edge guidance network for salient object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8779–8788, 2019.
- [45] Quanlong Zheng, Xiaotian Qiao, Ying Cao, and Rynson WH Lau. Distraction-aware shadow detection. In *Proceedings* of the IEEE Conference on Computer Vision and Pattern Recognition, pages 5167–5176, 2019.
- [46] Jiejie Zhu, Kegan GG Samuel, Syed Z Masood, and Marshall F Tappen. Learning to recognize shadows in monochromatic natural images. In 2010 IEEE Computer Society conference on computer vision and pattern recognition, pages 223–230. IEEE, 2010.
- [47] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 121–136, 2018.