

RenDetNet: Weakly-supervised Shadow Detection with Shadow Caster Verification

Nikolina Kubiak^{1,2}, Elliot Wortman¹, Armin Mustafa¹, Graeme Phillipson²,
Stephen Jolly², and Simon Hadfield¹

¹ University of Surrey, Guildford, UK
{n.kubiak,s.hadfield}@surrey.ac.uk

² BBC R&D, Salford, UK

Abstract. Existing shadow detection models struggle to differentiate dark image areas from shadows. In this paper, we tackle this issue by verifying that all detected shadows are real, *i.e.* they have paired shadow casters. We perform this step in a physically-accurate manner by differentially re-rendering the scene and observing the changes stemming from carving out estimated shadow casters. Thanks to this approach, the RenDetNet proposed in this paper is the first learning-based shadow detection model whose supervisory signals can be computed in a self-supervised manner. The developed system compares favourably against recent models trained on our data. As part of this publication, we release our code on github.

1 Introduction

Shadow detection and its twin task of shadow removal are crucial for the development of reliable computer vision systems. Such solutions can be used to solve shadow-related problems in real-life domains such as farming [7], remote sensing [24], autonomous driving [31], medical diagnostics [41], document de-shadowing [15] or appearance correction in casual capture photos [43].

The vast majority of existing shadow detection solutions operate in a supervised manner. The networks take in shadowed images and compare the estimated shadow masks with paired ground truth data. However, this approach has a number of disadvantages: Firstly, the labels are obtained via laborious hand-annotation, and are usually produced by multiple individuals with different labelling behaviours. Secondly, the datasets [30, 32, 46] that are the *de facto* standard in the image-based shadow detection literature lack diversity. They depict predominantly pavements, grass, etc., and contain little texture variation, other objects and/or clutter. This means that shadow detection is often posed simply as the task of detecting a darker region within the scene. Consequently, when applied to general images, shadow detection models trained on such datasets tend to mistake dark-coloured regions or objects for shadows. This is unsurprising as differentiating shadows from dark regions is an ill-posed problem. Even in biological vision, the only way to solve this problem is through context, and an understanding of scene structure. We visualise this in Fig. 1.




Fig. 1: Examples of the common shadow vs dark region problem in shadow detection. The examples come from the ISTD dataset [32] and the estimated shadow masks (SMs) were generated using a SOTA shadow detection model [4]. This shows that even some of the best available networks still struggle to distinguish between shadows and dark image regions, and inspires us to develop a detection model with improved scene understanding.

In this paper, we propose a completely new way to think about shadow detection. Our model, RenDetNet, is trained without any hand-labelled shadow masks. Instead, our system learns from synthetic rendered scenes, where it estimates a shadow mask and a caster mask (*i.e.* a mask of the object casting the shadow). We then use the caster mask to carve out the caster and re-render the scene. If the (binarised) difference between the original render and the carved render matches the estimated shadow and caster regions, it means the masked caster was responsible for the detected shadow. Using this approach, our RenDetNet learns not to predict shadows without credible casters.

In summary, the contributions of this paper are as follows:

1. We propose a weakly-supervised deep learning approach to shadow detection, producing accurate caster masks and shadow masks corresponding to real shadows; this includes cast and self-cast shadows (Fig. 3)
2. We present a self-supervised caster-aware dataset generation pipeline;
3. Our new model compares favourably against recent shadow detection methods on the new shadow-caster datasets presented in this paper.

2 Related work

Some shadow detection models treat the task as a preliminary step for shadow removal [32, 42]. However, more relevant to our work are techniques which attempt to distinguish between dark texture and shadow regions. scGAN [23] is trained with tunable sensitivity and final masks created via multi-scale mask aggregation. Le *et al.* [18] attenuate the strength of shadows in their dataset to train a robust detector. Hu *et al.* [10] learn the direction-aware spatial context (DSC) to estimate the lighting direction and thus the regions likely to be shadows. BDRAR [47] refines the features by combining the contexts in adjacent layers, and uses a bidirectional feature pyramid network to aggregate them across the network. Wang *et al.* [38] use stacked parallel fusion branches to combine the global context with the fine local features. DSDNet [44] actively tackles false

positives and negatives in shadow detection by operating on multi-scale features. Zhu *et al.*'s model [49] detects shadow and non-shadow areas, and leverages this complementary information to reduce the frequency of misdetections. Jie and Zhang [13] use shuffled multi-level features to provide good global and local context simultaneously. FDRNet [48] extracts illumination-invariant and -variant features, and combines them in a way that improves its robustness to brightness changes. In FCSD-Net [28] the estimated masks are verified through shadow removal. Similarly, SDDNet [6] models shadows and the backgrounds upon which they fall (*i.e.* shadow-free areas) separately. Sun *et al.* [27] use adaptive mapping to adjust the detection process for raw images. SILT [40] tackles the noise in shadow masks using shadow counterfeiting and self-teaching. Wang *et al.* expand the problem of detection, and additionally match all shadow instances with their casters [33–35]. The intuition behind these methods is similar to our approach, but the models require full supervision and laborious annotation.

2.1 Reducing the supervision requirements

Unlike the previously discussed works, we strive to reduce the necessary supervision requirements to allow for shadow detection in more complex domains. A few other papers explore this idea: A semi-supervised, multi-task MTMT-Net [4] learns to estimate shadow regions, edges and count in a supervised way, and is later finetuned using unlabelled data and a consistency loss. Chen *et al.* [5] and Lu *et al.* [22] propose video shadow detection models trained on labelled image data and unlabelled video data; Xing *et al.*'s method [39] uses labelled and unlabelled videos. Although these approaches use fewer labelled samples, they do not completely eliminate the need for ground truth. We could not identify any deep learning solutions to shadow detection capable of operating in a fully un- or self-supervised manner. However, some non-learning techniques exist, *e.g.* that of paired regions. Such methods operate by finding a shadow-free match for each shadowed pixel, based on reflectance consistency [8, 37]. The solutions do not require training, but they also do not perform on par with recent deep-learning models. Finally, while the idea of reducing supervision requirements has not been as popular in shadow *detection* works, some examples can be found in the tangential field of shadow *removal* [9, 14, 17, 21, 29].

2.2 Summary

As illustrated above, there exists a wide range of shadow detection models. Unfortunately, the supervision requirements of the available solutions are high - the methods are at least semi-supervised. This necessitates a heavy supervision workload for general applicability or, more realistically, leads to systems which can only operate in a restricted domain. Therefore, in this paper we present a novel weakly-supervised model where the necessary supervisory signal can be obtained in a self-supervised manner and the training is guided by differentiable scene carving and re-rendering in synthetic scenes.

3 Methodology

In the following sections, we discuss the RenDetNet model and our approach to weakly-supervised training. All elements of the proposed system are shown in Fig. 2. At inference, only the parts in the dotted box are required.

3.1 Differentiable shadow caster verification with RenDetNet

The model proposed in this paper is a fully-convolutional network based on [36]. Our RenDetNet has a shared encoder E and 2 decoder heads – one for estimating a shadow mask D_{sm} and one for a caster mask D_{cm} . Given a scene render \mathbf{I} , the process of obtaining a shadow mask \mathbf{SM} and a caster mask \mathbf{CM} can be formalised as

$$\mathbf{SM} = D_{sm}(E(\mathbf{I})) \quad \text{and} \quad \mathbf{CM} = D_{cm}(E(\mathbf{I})). \quad (1)$$

The scene is rendered using the rendering function \mathcal{R} that takes lighting ϕ , camera κ and mesh μ data as arguments, *i.e.* $\mathbf{I} = \mathcal{R}(\phi, \kappa, \mu)$. The mesh data can comprise a number of disconnected objects, not necessarily just one. The mesh μ can be defined as a collection of N triangular faces $\{\mathcal{F}\}$ each of which comprises a triplet of 3D vertices ($\mathbf{v} \in \mathbb{R}^3$), *i.e.*

$$\mu = \{\mathcal{F}_i | \mathcal{F}_i = \{\mathbf{v}_j\}_{j=1}^3\}_{i=1}^N. \quad (2)$$

To render the mesh, we need to project the 3D mesh vertices to 2D. To this end, we use a projection function π such that 2D point $\mathbf{v}' = \pi(\mathbf{v}, \mathbf{K}_\kappa, \mathbf{T})$ where \mathbf{v} , \mathbf{K}_κ and \mathbf{T} signify the 3D point coordinates, the intrinsics of camera κ and camera extrinsics respectively. Therefore, the 2D-projected mesh can be described as $\mu' = \{\mathcal{F}'_i\}_{i=1}^N$ with 2D faces

$$\mathcal{F}'_i = \{\pi(\mathbf{v}, \mathbf{K}_\kappa, \mathbf{T}) | \mathbf{v} \in \mathcal{F}_i\}. \quad (3)$$

Next, for our shadow caster verification, we need to check which vertices of the foreground object(s) lie in the estimated \mathbf{CM} . To do this, we pick all mesh faces that have at least 1 vertex in this mask. More formally, the parts of the mesh inside \mathbf{CM} are defined as

$$\mu'_{cm} = \{\mathcal{F}'_i | (\exists \mathbf{v} \in \mathcal{F}_i : \pi(\mathbf{v}, \mathbf{K}_\kappa, \mathbf{T}) \in \mathbf{CM})\}. \quad (4)$$

Having identified the mesh faces identified by the caster mask, we remove them to obtain a new, carved mesh

$$\hat{\mu} = \mu \setminus \{\mathcal{F}_i | \mathcal{F}'_i \in \mu'_{cm}\}. \quad (5)$$

Finally, we update the scene parameters with the new mesh data and re-render the scene. The resulting image $\hat{\mathbf{I}} = \mathcal{R}(\phi, \kappa, \hat{\mu})$ depicts the scene with some foreground structures masked and correspondingly reduced shadows.

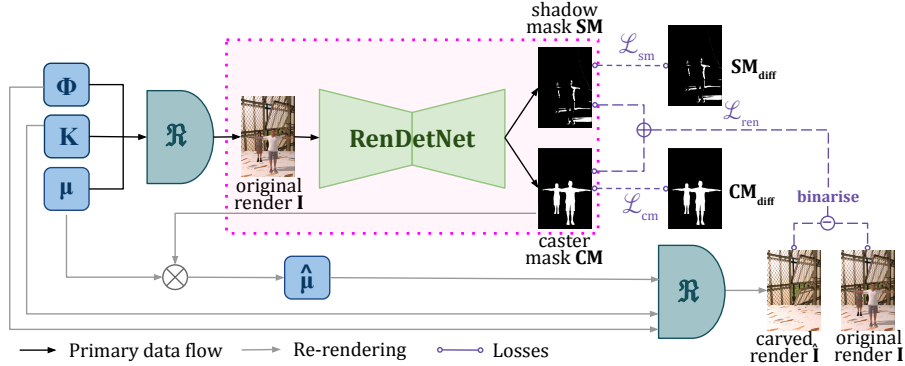


Fig. 2: At train time, we take in lighting ϕ , camera κ and mesh μ parameters and render the scene $\mathbf{I} = \mathcal{R}(\phi, \kappa, \mu)$. We push the render \mathbf{I} through our RenDetNet to obtain the shadow mask \mathbf{SM} and the caster mask \mathbf{CM} . We then use \mathbf{CM} to carve the mesh and re-render the scene, $\hat{\mathbf{I}} = \mathcal{R}(\phi, \kappa, \hat{\mu})$. At test time, only the region inside the dotted pink box is used.

3.2 The RenDetNet approach to supervision

To learn to estimate the shadow mask \mathbf{SM} and the caster mask \mathbf{CM} correctly, we use a number of losses described below. The supervisory signals used for these errors can be obtained at train-time in a self-supervised manner.

The key idea proposed in this paper is the shadow caster verification, *i.e.* ensuring that shadows are only detected in the image if they have plausible sources. To achieve this, we calculate the absolute difference between the original render \mathbf{I} and the carved mesh render $\hat{\mathbf{I}}$. We then apply Otsu thresholding [25] to calculate the threshold τ and turn this difference into a binary mask, represented by the Iverson brackets $\llbracket \cdot \rrbracket$. The change area $\Delta\mathbf{I}$ stemming from the mesh update can be defined as $\Delta\mathbf{I} = \llbracket abs(\mathbf{I} - \hat{\mathbf{I}}) > \tau \rrbracket$. We can then use $\Delta\mathbf{I}$ to calculate the rendering loss \mathcal{L}_{ren} given as

$$\mathcal{L}_{ren} = \|\Delta\mathbf{I} - (\mathbf{SM} \cup \mathbf{CM})\|_1, \quad (6)$$

and use this signal to guide shadow and caster mask generation.

If $\Delta\mathbf{I}$ and the masks differ, there are 2 possible explanations: Either a) the portion of the scene we wished to carve was not a shadow-casting foreground object and, thus, there was no $\Delta\mathbf{I}$ change in this region, or b) the carved caster removed a previously undetected shadow. Notably, as $\Delta\mathbf{I}$ is compared against the sum of the masks, \mathcal{L}_{ren} cannot distinguish between shadows and casters. In other words, \mathcal{L}_{ren} can be trivially satisfied by placing both shadows and casters in \mathbf{SM} and leaving \mathbf{CM} empty, or vice versa. To better control the masks, we add two extra losses.

Firstly, we add an additional constraint on \mathbf{CM} by observing the scene under new lighting ϕ' – constant emitter lighting. The resulting render is equivalent

to the scene’s reflectance, *i.e.* its illumination-invariant characteristics, and objects viewed under such illumination will not cast any shadows or have any self-shadows. If we threshold such a render with τ' , we can very easily obtain the object mask,

$$\mathbf{CM}_{diff} = \llbracket \mathcal{R}(\phi', \kappa, \mu) > \tau' \rrbracket, \quad (7)$$

which can help to ensure that the estimated \mathbf{CM} is fully contained within the object region. We can then use this constraint to calculate the caster loss \mathcal{L}_{cm} given as

$$\mathcal{L}_{cm} = BCE(\mathbf{CM}_{diff}, \mathbf{CM}), \quad (8)$$

where BCE denotes the binary cross-entropy loss.

We also restrict the \mathbf{SM} generation process in a similar, self-supervised manner; this consists of two steps. Firstly, we consider the scene under 2 different lighting conditions, with a point light flipped along the x-axis (denoted as $-\phi$). As before, \mathbf{I} denotes the render with original illumination ϕ and now \mathbf{I}' represents the image with the new light setting, *i.e.* $\mathbf{I}' = \mathcal{R}(-\phi, \kappa, \mu)$. Under these conditions, the shadow area in \mathbf{I} will have a corresponding lighter, non-shadowed region in \mathbf{I}' (and vice versa). If we subtract \mathbf{I} from \mathbf{I}' , we will get a positive difference in the \mathbf{I} shadow area and a negative difference in the \mathbf{I}' shadow area. Since we are only interested in the former, we only consider at the positive part of the difference. We turn the described difference into a binary mask \mathbf{SM}_{diff_1} defined as

$$\mathbf{SM}_{diff_1} = \llbracket (\mathcal{R}(-\phi, \kappa, \mu) - \mathbf{I}) > 0 \rrbracket. \quad (9)$$

Unfortunately, this way of supervising the \mathbf{SM} only works if there is no overlap between the shadows cast under ϕ and $-\phi$. To remedy this, we add a second constraint on this self-supervision method. We render our scene with the original mesh μ and again with just the background (*i.e.* with an empty foreground mesh $\mathcal{R}(\phi, \kappa, \emptyset)$). If we consider the difference between this background render and the original, the region of change will correspond to the objects represented by the mesh μ and their shadows. Knowing the caster mask estimate \mathbf{CM}_{diff} (Eq. 7), we can obtain the binary mask of just the shadow region,

$$\mathbf{SM}_{diff_2} = \llbracket abs(\mathcal{R}(\phi, \kappa, \emptyset) - \mathbf{I}) \rrbracket - \mathbf{CM}_{diff}. \quad (10)$$

Unlike the previous \mathbf{SM}_{diff_1} estimation method, the \mathbf{SM}_{diff_2} approach will not show self-cast shadows (as they lie inside \mathbf{CM}_{diff}). However, this method addresses the shadow overlap issue of the former solution. Therefore, we can combine both techniques and get the full shadow supervision signal defined as

$$\mathbf{SM}_{diff} = (\llbracket abs(\mathcal{R}(\phi, \kappa, \emptyset) - \mathbf{I}) \rrbracket - \mathbf{CM}_{diff}) + \llbracket (\mathcal{R}(-\phi, \kappa, \mu) - \mathbf{I}) > 0 \rrbracket. \quad (11)$$

Given \mathbf{SM}_{diff} , we can constrain the shadow mask generated by our model. We perform this using the shadow mask loss \mathcal{L}_{sm} defined as

$$\mathcal{L}_{sm} = BCE(\mathbf{SM}_{diff}, \mathbf{SM}). \quad (12)$$

The total loss used to guide the training of RenDetNet can be described as

$$\mathcal{L}_{total} = \mathcal{L}_{ren}\lambda_{ren} + \mathcal{L}_{cm}\lambda_{cm} + \mathcal{L}_{sm}\lambda_{sm}. \quad (13)$$

where the λ s represent the weightings on all loss components. Empirically, we set $\lambda_{cm} = \lambda_{sm} = 1$; λ_{ren} is initially set to 0 and then its value increases to 1 after 1000 iterations.

To evaluate the performance of our model, only the RenDetNet part of the described system (pink box in Fig. 2) is needed. This also means that at test time no 3D data or other scene parameters are required – the inference is performed solely on 2D image data.

4 Experiments

In the following sections, we discuss the datasets created as part of this publication and use them to evaluate the performance of a number of models.

4.1 The dataset

As previously discussed, capturing training datasets for shadow detection is a challenge. Recently, Inoue & Yamasaki [11] proposed a pipeline for synthetic shadow removal/detection dataset creation. In their setup, shadow masks are used to cast shadows of varying intensity onto previously shadow-free scenes, thus generating shadowed samples. While interesting, this approach does not help to resolve the ‘dark vs shadow’ issue faced by the existing solutions.

To address this, we set out a few objectives for our data generation pipeline: We want to include information about shadows and their casters. We do not want to rely on any existing shadow masks (or shadow-filtered data) to run the generation pipeline. Finally, we want to create the dataset in a fully automated, self-supervised manner, without any human annotation. Hand-labelling datasets limits their scalability, so learning with imperfect labels should be prioritised. A number of learning-based solutions already get their (imperfect) masks from previously trained models, so we consider our physics-backed estimation approach to be similarly viable while also being capable of adapting to new data.

To develop and demonstrate our idea, we rendered 2 datasets using *Mitsuba 3* [12]; we refer to them as Datasets #1 & #2. Due to copyright reasons, the meshes used to create Dataset #1 cannot be released. In the interest of reproducibility, we created Dataset #2 using only publicly available components and we provide the code used to create it.

Each dataset scene features between 1 and 3 textured meshes – human meshes from 3D Virtual Humans (3DVH) [3] for Dataset #1 or a mix of BEDLAM [2] humans and clutter objects [16,20] for Dataset #2. Dataset #1 features humans in fixed poses. To maximise the cast shadow diversity, the humans in Dataset #2 are posed in a variety of ways using AGORA [26]. The objects (human or not) are placed in a range of locations on patterned floors and in front of heavily textured

backgrounds, the textures of which come from [1] and [19, 45] respectively. We additionally vary the scene illumination, changing both the magnitude of the constant emitter as well as the positioning and brightness of a point light. The resulting datasets feature approx. 4.2k train / 900 test samples (Dataset #1) and 9k train / 1k test samples (Dataset #2).

4.2 Evaluation metrics and experimental setup

The RenDetNet system was developed in Pytorch. The first version, with Dataset #1, was trained for approx. 8k steps on a GeForce RTX 3090. More compact and diverse Dataset #2 required approx. 12k steps on a single GeForce RTX 2080. All variants of RenDetNet were trained in the Lab colour space unless otherwise stated. The evaluations in the following sections are performed in terms of the BER (Balanced Error Rate) score, defined as

$$BER = \left(1 - 0.5 \times \left(\frac{TP}{N_s} + \frac{TN}{N_{ns}} \right) \right) \times 100, \quad (14)$$

where TP/TN symbolise the true positives/negatives and N_s/N_{ns} describe the total number of shadow/non-shadow pixels. We also cite the average shadow region (S) and non-shadow region (NS) BER scores. In all cases, a lower BER score signifies a better result. When quantitative data is reported, the best score is shown in **bold** and the runner-up is underlined. All training and evaluation runs were performed on full-size images, *i.e.* 512×768 pixels.

4.3 Ablation study

First, we conduct an ablation study to demonstrate the impact of each of our design choices. Specifically, we look at the impact of different training losses and discuss the importance of choosing the right colour space for training a system for lighting-related problems. At the end, we also demonstrate the shadow-caster matching performance of RenDetNet.

Model losses study. In this study, we gradually add losses to our model trained on Dataset #1 and measure their impact on RenDetNet’s performance. The quantitative results of this evaluation are shown in Table 1. During the experiments, we first train the most basic version of our model. As the rendering loss \mathcal{L}_{ren} does not provide enough guidance to disentangle **SM** and **CM** on its own, we consider our “bare minimum” scenario to operate based on \mathcal{L}_{ren} and \mathcal{L}_{sm} (line #1). We then add \mathcal{L}_{cm} (line #2) to control both estimated masks, which leads to slight improvements, particularly in the shadowed region. Next, we consider the network trained with just \mathcal{L}_{cm} and \mathcal{L}_{sm} (line #3) and no shadow caster verification. The resulting model is good at spotting shadows (lowest BER(S)) but it struggles in the non-shadow regions (highest BER(NS)). We attribute this to the dark/textured regions that are misclassified as shadows. Finally, to combine the success of models #2 and #3, we adopt a staged training mechanism: We start training the model with just the mask losses and add

\mathcal{L}_{ren} after 1000 iterations. Around this time the **SM** and **CM** learning process slows down, but the masks have become plausible enough to be constrained in a physics-backed manner. The resulting system (line #4) has the lowest overall BER score, which confirms the complementary nature of the previously discussed losses.

Table 1: Loss ablation study - Dataset #1.

#	\mathcal{L}_{ren}	\mathcal{L}_{sm}	\mathcal{L}_{cm}	staged	BER ↓	BER (S) ↓	BER (NS) ↓
1	✓	✓			7.324	13.526	1.123
2	✓	✓	✓		7.182	13.103	<u>1.261</u>
3		✓	✓		<u>6.598</u>	8.939	4.257
4	✓	✓	✓	✓	6.344	<u>11.010</u>	1.679

Colour space study. Most shadow detection models are trained in the RGB colour space. However, since shadows are lighting-related phenomena, we experimented with training our model in the *Lab* colour space. In *Lab*, the *L* channel corresponds to illumination and the colour information is captured by the *a* and *b* channels. As demonstrated in Table 2, training RenDetNet on data with such a representation is clearly advantageous.

Table 2: Colour space ablation study - Dataset #1.

Colour space	BER ↓	BER (S) ↓	BER (NS) ↓
RGB	<u>46.823</u>	<u>87.269</u>	<u>6.377</u>
Lab	6.344	11.010	1.679



Fig. 3: Examples of RenDetNet’s caster & shadow identification abilities - Dataset #2

Shadow-caster pairing. As demonstrated in Fig. 3, aside from detecting shadows, the proposed RenDetNet is also capable of identifying the shadows’ casters within the scene with high certainty and a good amount of detail.

4.4 Comparisons with existing models

For fair comparison with state-of-the-art solutions, we rendered all of our training images and generated a shadow mask for each sample to create supervised 2D versions of our datasets. Some of the recent shadow detection models rely on false-negatives and false-positives pre-calculated based on the outputs of existing models, *e.g.* [44], or require more data for different multi-task problems [4] or sub-problems [33–35], which unfortunately makes them infeasible for our evaluation. Therefore, for our comparisons, the compatible models – DSC [10], BDRAR [47] and FDRNet [48] – were trained on our data using their publicly available code. The results of this study are presented in Table 3 and Fig. 4.

The qualitative data shows the advantage of our caster verification step. In the top row of Fig. 4 (Dataset #1) we can see that RenDetNet manages to find the shadows, both cast on the ground as well as on the caster itself. FDRNet performs similarly well yet loses some detail in the shadow cast on the floor. BDRAR struggles to find any shadows while DSC finds the caster but also estimates a large portion of the floor as shadow. We speculate this is due to the clutter and variety of colours and textures in the scene. In the next two rows (Dataset #2), DSC largely confuses the caster and the shadow, loosely covering both areas. BDRAR correctly identifies the cast shadows, but completely ignores any self-cast shadows. FDRNet captures more details and self-cast shadows than BDRAR, yet, similarly to DSC, its shadow boundaries tend to spill. Our model reconstructs the ground shadows to a similar degree as BDRAR, and preserves more detail than FDRNet and DSC.

Numerically, our RenDetNet achieves a good balance in terms of BER(S) vs BER(NS). This places it as the top performer (Dataset #1) or a close runner-up (Dataset #2); nevertheless, visually, our method clearly outperforms the top-performing FDRNet. Conservative in its shadow estimation, BDRAR tends to

Table 3: Performance in terms of BER - Dataset #1 (top) & #2 (bottom)

Dataset	Model	BER ↓	BER (S) ↓	BER (NS) ↓
#1	DSC [10]	14.716	19.256	10.176
	BDRAR [47]	8.619	15.887	1.352
	FDRNet [48]	<u>7.276</u>	10.714	3.839
	RenDetNet (ours)	6.344	<u>11.010</u>	<u>1.679</u>
#2	DSC [10]	15.029	<u>16.599</u>	13.459
	BDRAR [47]	15.415	30.359	0.472
	FDRNet [48]	11.634	14.449	8.820
	RenDetNet (ours)	<u>11.992</u>	21.347	<u>2.638</u>



Fig. 4: Qualitative results – row 1: Dataset #1; rows 2-3: Dataset #2.

be superior in terms of $BER(NS)$, yet the system tends to underestimate the shadows (high $BER(S)$). In contrast, FDRNet tends to overestimate the shadows, which then leads to good precision in terms of shadow-area detection, but worse $BER(NS)$ performance. RenDetNet offers a happy medium.

5 Conclusions

In this paper, we have proposed a new weakly-supervised approach to physically-verified shadow caster matching. The method allows us to accurately estimate shadows, both cast on the external environment as well as self-cast shadows which overlap the caster itself. Additionally, the proposed RenDetNet model can be trained in a fully-differentiable setup with the supervisory signals calculated on the fly, and requires no hand annotation or other human input. Finally, we have presented a new self-supervised caster-aware dataset generation pipeline. While developed on a particular set of images, meshes and textures, the same idea could be applied to other sets of data.

As future work, it would be interesting to explore alternative training regimes which can jointly exploit existing strongly supervised datasets alongside our novel self-supervised data generation approach. It would also be productive to

obtain a set of real-life images, perhaps with hand-annotated masks, to create an improved validation set for shadow/caster estimation.

Acknowledgements

This work was partially supported by the British Broadcasting Corporation (BBC) and the Engineering and Physical Sciences Research Council's (EPSRC) industrial CASE project "Computational lighting in video" (voucher number 19000034).

References

1. Tiles direct. <https://www.tiles-direct.com>
2. Black, M.J., Patel, P., Tesch, J., Yang, J.: Bedlam: A synthetic dataset of bodies exhibiting detailed lifelike animated motion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
3. Caliskan, A., Mustafa, A., Imre, E., Hilton, A.: Multi-view consistency loss for improved single-image 3d reconstruction of clothed people. In: Proceedings of the Asian Conference on Computer Vision (ACCV) (2020)
4. Chen, Z., Zhu, L., Wan, L., Wang, S., Feng, W., Heng, P.A.: A multi-task mean teacher for semi-supervised shadow detection. In: Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition (CVPR). pp. 5611–5620 (2020)
5. Chen, Z., Lu, X., Zhang, L., Xiao, C.: Semi-supervised video shadow detection via image-assisted pseudo-label generation. In: Proceedings of the 30th ACM International Conference on Multimedia. p. 2700–2708 (2022)
6. Cong, R., Guan, Y., Chen, J., Zhang, W., Zhao, Y., Kwong, S.: Sddnet: Style-guided dual-layer disentanglement network for shadow detection. In: Proceedings of the 31st ACM International Conference on Multimedia. pp. 1202–1211 (2023)
7. Dornadula, S.P.K., Brunet, P., Elias, D., et al.: Ai driven shadow model detection in agropv farms. arXiv preprint arXiv:2304.07853 (2023)
8. Guo, R., Dai, Q., Hoiem, D.: Single-image shadow detection and removal using paired regions. In: Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition (CVPR). pp. 2033–2040 (2011)
9. Hu, X., Jiang, Y., Fu, C.W., Heng, P.A.: Mask-ShadowGAN: Learning to Remove Shadows from Unpaired Data. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2019)
10. Hu, X., Zhu, L., Fu, C.W., Qin, J., Heng, P.A.: Direction-aware spatial context features for shadow detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7454–7462 (2018)
11. Inoue, N., Yamasaki, T.: Learning from synthetic shadows for shadow detection and removal. *IEEE Transactions on Circuits and Systems for Video Technology* **31**(11), 4187–4197 (2020)
12. Jakob, W., Speierer, S., Roussel, N., Nimier-David, M., Vicini, D., Zeltner, T., Nicolet, B., Crespo, M., Leroy, V., Zhang, Z.: Mitsuba 3 renderer (2022), <https://mitsuba-renderer.org>
13. Jie, L., Zhang, H.: Rmlanet: Random multi-level attention network for shadow detection. In: 2022 IEEE International Conference on Multimedia and Expo (ICME). pp. 1–6 (2022)
14. Jin, Y., Sharma, A., Tan, R.T.: Dc-shadownet: Single-image hard and soft shadow removal using unsupervised domain-classifier guided network. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5027–5036 (2021)
15. Jung, S., Hasan, M.A., Kim, C.: Water-filling: An efficient algorithm for digitized document shadow removal. In: Proceedings of the Asian Conference on Computer Vision (ACCV) (2018)
16. Kaskman, R., Zakharov, S., Shugurov, I., Ilic, S.: Homebreweddb: Rgb-d dataset for 6d pose estimation of 3d objects. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (ICCVW) (2019)

17. Kubiak, N., Mustafa, A., Phillipson, G., Jolly, S., Hadfield, S.: S3r-net: A single-stage approach to self-supervised shadow removal. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (2024)
18. Le, H., Vicente, T.F.Y., Nguyen, V., Hoai, M., Samaras, D.: A+ d net: Training a shadow detector with adversarial shadow attenuation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 662–678 (2018)
19. Li, J., Zhang, J., Maybank, S.J., Tao, D.: Bridging composite and real: towards end-to-end deep image matting. *International Journal of Computer Vision (IJCV)* **130**(2) (2022)
20. Lin, Y., Tremblay, J., Tyree, S., Vela, P.A., Birchfield, S.: Multi-view fusion for multi-level robotic scene understanding. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 6817–6824 (2021)
21. Liu, Z., Yin, H., Mi, Y., Pu, M., Wang, S.: Shadow removal by a lightness-guided network with training on unpaired data. *IEEE Transactions on Image Processing* **30**, 1853–1865 (2021)
22. Lu, X., Cao, Y., Liu, S., Long, C., Chen, Z., Zhou, X., Yang, Y., Xiao, C.: Video shadow detection via spatio-temporal interpolation consistency training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3116–3125 (2022)
23. Nguyen, V., Vicente, T.F.Y., Zhao, M., Hoai, M., Samaras, D.: Shadow detection with conditional generative adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 4520–4528 (2017)
24. Ni, X., Zhang, X., Zhan, X., Yang, Z., Shi, J., Wei, S., Zeng, T.: Shadow-oriented tracking method for multi-target tracking in video-sar. *arXiv preprint arXiv:2211.15995* (2022)
25. Otsu, N.: A threshold selection method from gray-level histograms. *IEEE transactions on systems, man, and cybernetics* **9**(1) (1979)
26. Patel, P., Huang, C.H.P., Tesch, J., Hoffmann, D.T., Tripathi, S., Black, M.J.: Agora: Avatars in geography optimized for regression analysis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13468–13478 (2021)
27. Sun, J., Xu, K., Pang, Y., Zhang, L., Lu, H., Hancke, G., Lau, R.: Adaptive illumination mapping for shadow detection in raw images. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 12709–12718 (2023)
28. Valanarasu, J.M.J., Patel, V.M.: Fine-context shadow detection using shadow removal. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 1705–1714 (2023)
29. Vasluianu, F.A., Romero, A., Gool, L.V., Timofte, R.: Shadow removal with paired and unpaired learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (2021)
30. Vicente, T.F.Y., Hou, L., Yu, C.P., Hoai, M., Samaras, D.: Large-scale training of shadow detectors with noisily-annotated shadow examples. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 816–832 (2016)
31. Wang, C., Xu, H., Zhou, Z., Deng, L., Yang, M.: Shadow detection and removal for illumination consistency on the road. *IEEE Transactions on Intelligent Vehicles* **5**(4) (2020)
32. Wang, J., Li, X., Yang, J.: 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 (CVPR). pp. 1788–1797 (2018)
33. Wang, T., Hu, X., Fu, C.W., Heng, P.A.: Single-stage instance shadow detection with bidirectional relation learning. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1–11 (2021)
 34. Wang, T., Hu, X., Heng, P.A., Fu, C.W.: Instance shadow detection with a single-stage detector. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(3), 3259–3273 (2023)
 35. Wang, T., Hu, X., Wang, Q., Heng, P.A., Fu, C.W.: Instance shadow detection. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1877–1886 (2020)
 36. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2018)
 37. Wang, X., Yao, S., Dai, P., Wang, R., Cao, X.: Updated paired regions for shadow detection from single image. In: Proceedings of the British Machine Vision Conference (BMVC) (2021)
 38. Wang, Y., Zhao, X., Li, Y., Hu, X., Huang, K., CRIPAC, N.: Densely cascaded shadow detection network via deeply supervised parallel fusion. In: *IJCAI*. vol. 5, p. 6 (2018)
 39. Xing, Z., Wang, T., Hu, X., Wu, H., Fu, C.W., Heng, P.A.: Video instance shadow detection. arXiv preprint arXiv:2211.12827 (2022)
 40. Yang, H., Wang, T., Hu, X., Fu, C.W.: Silt: Shadow-aware iterative label tuning for learning to detect shadows from noisy labels. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 12687–12698 (2023)
 41. Yasutomi, S., Arakaki, T., Hamamoto, R.: Shadow detection for ultrasound images using unlabeled data and synthetic shadows. In: *Medical Imaging with Deep Learning* (2019)
 42. Yücel, M.K., Dimaridou, V., Manganelli, B., Ozay, M., Drosou, A., Saà-Garriga, A.: Lra&ldra: Rethinking residual predictions for efficient shadow detection and removal. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 4925–4935 (2023)
 43. Zhang, X.C., Barron, J.T., Tsai, Y.T., Pandey, R., Zhang, X., Ng, R., Jacobs, D.E.: Portrait Shadow Manipulation. *ACM Transactions on Graphics* (2020)
 44. Zheng, Q., Qiao, X., Cao, Y., Lau, R.W.: Distraction-aware shadow detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5167–5176 (2019)
 45. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2017)
 46. Zhu, J., Samuel, K.G.G., Masood, S.Z., Tappen, M.F.: Learning to recognize shadows in monochromatic natural images. In: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 223–230 (2010)
 47. Zhu, L., Deng, Z., Hu, X., Fu, C.W., Xu, X., Qin, J., Heng, P.A.: Bidirectional feature pyramid network with recurrent attention residual modules for shadow detection. In: Proceedings of the European Conference on Computer Vision (ECCV) (2018)
 48. Zhu, L., Xu, K., Ke, Z., Lau, R.W.: Mitigating intensity bias in shadow detection via feature decomposition and reweighting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4702–4711 (2021)

49. Zhu, Y., Fu, X., Cao, C., Wang, X., Sun, Q., Zha, Z.J.: Single image shadow detection via complementary mechanism. In: Proceedings of the 30th ACM International Conference on Multimedia. p. 6717–6726 (2022)